

18 Reviews of individual patient data

Lesley A Stewart, Jayne F Tierney and Mike Clarke on behalf of the Cochrane Individual Patient Data Meta-analysis Methods Group

Key Points

- In an individual patient data (IPD) meta-analysis, the original research data for each participant in each study are sought directly from the researchers responsible for that study.
- Having access to the ‘raw’ data for each study enables data checking, thorough exploration, and re-analysis of the data in a consistent way.
- IPD meta-analysis has particular benefits when the published information does not permit a good quality review, or where particular types of analyses are required that are not feasible using summary data.
- Most IPD meta-analyses are carried out and published by a collaborative group, comprising a project team or secretariat, the researchers who contribute their study data, and often also an advisory group.
- An IPD approach usually takes longer and costs more than a conventional systematic review relying on published or aggregate data.
- There may be circumstances where the benefits of obtaining IPD are marginal; others where it could be vital.

18.1 Introduction

18.1.1 What is an IPD review?

Individual patient data (IPD) meta-analysis is a specific type of systematic review. Instead of extracting data from study publications, the original research data for each participant in an included study are sought directly from the researchers responsible for that study. These data can then be re-analysed centrally and, if appropriate, combined in meta-analyses. Cochrane reviews can be undertaken as IPD reviews, but IPD reviews usually require dedicated staff and would be difficult to conduct in 'free time'. The approach requires particular skills and usually takes longer and costs more than a conventional systematic review relying on published or aggregate data. However, IPD reviews offer benefits related particularly to the quality of data and the type of analyses that can be done (Stewart 1995, Stewart 2002). For this reason they are considered to be a 'gold standard' of systematic review. This chapter aims to provide an overview of the IPD approach to systematic review, to help authors decide whether collecting IPD might be useful and feasible in their review. It does not provide detailed methodology, and anyone contemplating carrying out their first IPD review should seek appropriate advice and guidance from experienced researchers through the IPD Meta-analysis Methods Group (see Box 18.6.a).

18.1.2 When should an IPD review be done?

IPD reviews should be considered in circumstances where the published information does not permit a good quality review, or where particular types of analyses are required that are not feasible using standard approaches. There are situations where the IPD approach will not be feasible, because data have been destroyed or lost or, despite every effort, researchers do not wish to collaborate. There may also be circumstances where it may not be necessary, for example if all the required data are readily available in a suitable format within publications. Further details of when IPD may be beneficial are given in Box 18.1.a.

18.1.3 How are IPD review methods different?

The general approach to IPD meta-analysis is the same as for any other systematic review, and the methods used should differ substantially only in the data collection, checking and analysis stages. Just as for any Cochrane review, a detailed protocol should be prepared, setting out the objective for the review, the specific questions to be addressed, study inclusion and exclusion criteria, the reasons why IPD are sought, the methods to be used and the analyses that are planned. Similarly, the methods used to identify and screen studies for eligibility should be the same irrespective of whether IPD will be sought, although the close involvement of the original researchers in the project might make it easier to find other studies done by them or known to them. The

Box 18.1.a Potential benefits of IPD**IPD may be beneficial in the following situations.**

- Many studies are unpublished or published only in the grey literature.
- There is poor reporting of studies (e.g. information presented is inadequate, selective or ambiguous).
- A high proportion of individuals has been excluded from published analyses.
- Obtaining additional longer-term outcome data beyond that reported may provide useful insights (e.g. for mortality or child development outcomes).
- Outcome measures have been defined differently across studies.
- Time-to-event outcome measures are required.
- Multivariate or other complex analyses are required.
- Exploration of interactions between interventions and patient-level characteristics is important.

project should culminate in the preparation and dissemination of a structured report. An IPD review might also include a meeting at which results are presented and discussed with the collaborating researchers.

18.1.4 How are IPD reviews organized?

IPD reviews are usually carried out as collaborative projects whereby all researchers contributing information from their studies, together with those managing the project, become part of an active collaboration. The projects are managed by a small local project group or secretariat, which may be aided in important and strategic decision-making by a larger advisory group. Results are usually published in the name of the collaborative group. The secretariat may also be responsible for organizing meetings of collaborators, to bring individuals together to discuss the preliminary results.

18.1.5 What healthcare areas have used the IPD approach?

IPD meta-analyses have an established history in cardiovascular disease and cancer, where the methodology has been developing steadily since the late 1980s. In cancer, for example, there are now more than 50 IPD meta-analyses of screening and treatment across a wide range of solid tumour sites and haematological malignancies (Clarke 1998). IPD have also been used in systematic reviews in many other fields (Simmonds 2005), including HIV infection, dementia, epilepsy, depression, malaria, hernia and asthma. The Cochrane Collaboration Individual Patient Data Meta-analysis Methods Group web site includes a database of ongoing and completed IPD reviews where further information can be found (see Box 18.6.a).

18.1.6 If I am thinking about doing an IPD review, what should I do first?

Before embarking on an IPD review, the skills and funding required for the success of the project should be considered carefully and training and advice should be sought. The Cochrane Collaboration Individual Patient Data Meta-analysis Methods Group is a good first point of contact (Box 18.6.a).

18.2 The collaborative nature of IPD meta-analyses

18.2.1 Collaborative groups

Most IPD meta-analyses are carried out and published by collaborative groups. These groups comprise the project team or secretariat managing the IPD review, members of the advisory group (if there is one) and the researchers who contribute their study data for re-analysis.

18.2.2 Negotiating collaboration

Establishing collaboration takes considerable time and effort. It can be difficult to trace the people responsible for eligible studies and they may be initially reluctant to participate in the meta-analysis. Often the first approach will be by letter, inviting collaboration, explaining the project, describing what participation will entail and how the meta-analysis will be managed and published. The letter is often from the project team and might be sent on behalf of the advisory group for the review. A protocol is generally supplied at this stage to provide further information, but data are not usually sought in the first correspondence. It may also be necessary to establish separate contact with the data centre or research organization who are (or have been) responsible for management of the study data, and to whom data queries will need to be sent. In encouraging the original investigators to take part in the IPD review, it is important to be as supportive and flexible as possible, to take the time required to build relationships and to keep all collaborators involved and informed of progress. Regular newsletters and e-mail updates can be useful ways of keeping the collaborative group up to date and involved, especially if the project will take place over a prolonged period.

18.2.3 Confidentiality

Researchers naturally require safeguards on the use of their study data and wish to ensure that it will be stored securely and used appropriately. For this reason, a signed confidentiality agreement is often used as a 'contract' between the original investigators and the IPD review team. The details of such agreements will vary, but most will state that data will be held securely, be accessed only by authorized members of the project

team and will not be copied or distributed elsewhere. It is also good practice to request that individual participants are de-identified in supplied data, such that individuals are identified only by a study identifier code and not by name. This seems to be an increasing requirement for obtaining IPD from some countries where data protection legislation requires that a participant cannot be identified from the data supplied. Data sent by email should be encrypted wherever possible.

18.3 Dealing with data

18.3.1 Deciding what data to collect

The protocol should specify what outcomes and patient characteristics are to be analysed. However, before embarking on data collection it is sensible to ask the original investigators about what data are actually available. When deciding which variables to collect, it is often sensible to start by considering carefully what analyses are planned and what data will be needed to do them. This minimizes the possibility that essential information will not be sought or that unnecessary data will be collected. Understandably, investigators can get upset or suspicious if they have gone to the trouble of providing data that are not subsequently analysed and reported.

Although in many cases it will be possible to collect specific variables for outcomes and characteristics as defined in the individual studies, it may be necessary to consider whether there are any data items for which further or constituent variables may be required. For instance, if studies have used different definitions of outcomes it may be desirable to redefine these for each patient in a consistent way across studies, and additional variables may be needed. For example, to redefine pre-eclampsia, data on systolic and diastolic blood pressure and proteinuria would need to be collected.

18.3.2 Data format

Once original investigators have agreed to collaborate, the next step is to provide clear instructions on what data they need to supply and on any preferred data format. The project team should be prepared to accept data in whatever format is most convenient for those supplying it, whether that is electronically, as printouts, or on paper forms, and should be prepared to recode information as necessary. However, although the early IPD meta-analyses in the 1980s relied heavily on data being supplied on paper, most information is now supplied by email or on disk, and investigators are often willing to transform or code their data according to the specified format.

18.3.3 Re-coding and re-defining supplied variables

Collecting data at the level of the individual participant enables translation between different staging, grading, ranking or other scoring systems, and may therefore allow pooling of data from studies that would not otherwise be possible, because of differences

between the data collection tools. To allow this, it is important that the appropriate data are sought (see Section 18.3.1) and that the data supplied are recoded or transformed to reflect common definitions. For example, if the outcome of interest is pre-eclampsia, data on blood pressure and proteinuria would need to be collected and considered together to define whether the pre-eclampsia (according to the review protocol definition) had been observed.

18.3.4 Checking data supplied

The aims of checking data are to increase the probability that data supplied are accurate, to confirm that trials are appropriately randomized, and where appropriate to make sure that, as far as possible, the data are up to date. The exact checking procedures to be carried out will depend on the healthcare area and question addressed, as well as the nature of the data supplied, but four main areas are typical:

18.3.4.1 Checking for missing or duplicated data

When data are received, it is important to check these as soon as possible to ensure that they can be read and loaded into the central analysis system. For example, if the data arrive as email attachments, it should be checked that the files can be opened and that the information is for the correct study. At this stage it is useful to confirm that data have been received for all appropriate (usually all randomized) individuals, checking that the numbers supplied are consistent with any publications or other information and that, for example, there are no obvious omissions or duplicates in the sequence of patient record or study identifier numbers.

18.3.4.2 Checking plausibility

Plausibility checks should include range checks on variables supplied, asking the original investigators to confirm any extreme outliers or unusual values: for example, confirming that records of unusually old or young patients or those with abnormally high or low cholesterol levels are indeed correct. Information supplied should also be checked against any relevant study publications, for example by confirming that the distribution of baseline characteristics, the number of participants and outcome results are consistent (bearing in mind that continued enrolment or additional follow-up may have altered information subsequent to publication).

18.3.4.3 Checking randomization

It is often helpful to check that randomization appears to have been done appropriately. Where dates of randomization are available, this can be explored by looking at plots of cumulative accrual over time; one would expect numbers enrolled to each intervention to be similar and for enrolment curves to cross frequently. It can also be informative

to look at the distribution of randomizations by day of the week. Here, provided that reasonable numbers of individuals have been randomized, one would expect to see roughly the same numbers randomized to each intervention on any given weekday, and that trials randomizing during normal clinic hours have few, if any, participants enrolled on unexpected days. It is also useful to check that the intervention groups are balanced for important baseline characteristics and within important participant subgroups, but bearing in mind that statistically significant imbalances can occur by chance.

18.3.4.4 Checking information is up to date

For outcomes where events are observed over a prolonged period, for example survival in cancer trials, it is important to check that follow-up is as up to date as possible and that it is consistent for each of the intervention groups. Producing a ‘reverse’ Kaplan Meier curve, based on just those patients who have not experienced the event of interest, with censoring then used as the event, can provide a useful check on the balance of follow-up across the groups.

For any individual study, the results of all these checks should be considered together to build up an overall picture of the study and the quality of the data that have been supplied, and any potential problems. Any concerns should be brought diplomatically to the attention of the researchers responsible. Usually, problems turn out to be simple errors or misunderstandings, which can be resolved through discussion. Major problems that cannot be resolved are rare.

A copy of the data as supplied should be archived before carrying out conversions or modifications to the data. Throughout the data checking processes, it is important that any changes and alterations made to the supplied data are properly logged.

18.4 Analysis

18.4.1 Analysis advantages

Having access to the ‘raw’ data for each study enables checking, thorough exploration, and re-analysis of the data in a consistent way. Thus, one does not have to rely on interpreting information and analyses presented in published reports, be constrained by summary data provided in tabular format, or be forced to consider combining the summary statistics from studies that have been calculated in different ways. It also avoids problems with the original analyses; for example it might be possible to carry out analyses according to intention-to-treat principles, even if the original trial analyses did not do this.

18.4.2 General approach

Most IPD meta-analyses to date have used a two-stage approach to analysis. In the first stage, each individual study is analysed in the same way, as set out in the

meta-analysis protocol or analysis plan. In the second step, the results, or summary statistics, of each of these individual study analyses are combined to provide a pooled estimate of effect in the same way as for a conventional systematic review (Simmonds 2005). More complex approaches using multilevel modelling have been described for binary data (Turner 2000), continuous data (Higgins 2001), ordinal data (Whitehead 2001) and time-to-event data (Tudor Smith 2005b) but, currently, their application is less common. When there is no heterogeneity between trials, a stratified log-rank two-stage approach for time-to-event data may be best avoided for estimating larger intervention effects (Tudor Smith 2005a).

18.4.3 Time-to-event analyses

Collecting IPD that include the time interval between the randomization and the event of interest enables time-to-event analyses to be conducted. These include, for example, time to recovery, time free of seizures, time to conception and time to death. Indeed, one of the main reasons that IPD meta-analyses have been so important in the cancer field is that time-to-event analysis of survival is vital in evaluating therapies. Most interventions are more likely to lead to a prolongation of survival rather than a cure. Therefore, it is important to measure not only whether a death happens, but also the time at which it takes place. To allow this type of analysis one needs to know the time that each individual spends 'event-free'. This is usually collected as the date of randomization, the event status (i.e. whether the event was observed or not) and the date of last evaluation for the event. Sometimes, it will be collected as the interval in days between randomization and the most recent evaluation for the event. Time-to-event analyses are performed for each trial to calculate hazard ratios, which are then pooled in the meta-analysis (see Section 9.4.9).

18.4.4 Bringing analyses up to date: long-term outcomes

For outcomes such as survival, where events can continue to take place over time, IPD meta-analyses can provide an important opportunity to examine the effects of interventions over a prolonged period. They can also provide an opportunity for researchers to provide more up-to-date data for relevant outcomes such as mortality than they have published for their study.

18.4.5 Subgroup analysis

Collecting IPD is also the most practical way to carry out analyses to investigate whether any observed effect of an intervention is consistent across well-defined types of participants, for example whether women gain a smaller or larger benefit from treatment than men. In conventional analyses using aggregate data from publications, it is usually very difficult to extract sufficient compatible data to undertake meaningful subgroup analyses, and especially difficult to characterize individuals by more than one

factor at a time. In contrast, IPD permit straightforward categorization of individuals for subgroup analysis (stratified by study) defined by single or multiple factors. The collection of IPD will also allow more complex analyses, such as multilevel modelling, to explore associations between intervention effects and patient characteristics.

18.4.6 Additional analyses

Access to the IPD also permits an in-depth exploration of patient characteristics themselves, irrespective of the intervention. For example, the large datasets collected can be used in the construction of prognostic indices that may be able to predict outcome based on patient characteristics. (International Germ Cell Cancer Collaborative Group 1997).

18.4.7 Software

IPD cannot be analysed directly in RevMan. The data need to be first analysed outside of this software, and summary statistics for each study may be entered into RevMan if a two-stage approach is used. For dichotomous and continuous outcomes, data may be entered in the usual way. For time-to-event outcomes, the observed-minus-expected number of events and variance may be entered using the 'O – E and Variance' option. Alternatively the generic inverse-variance option may be used to analyse effect estimates such as hazard ratios, rate ratios or adjusted estimates.

Although many standard statistical packages can perform the necessary analyses of IPD from the individual studies, it can be unwieldy and time-consuming to have to analyse each outcome in each study one at a time, and commercially available software is not currently available that supports the direct analysis, pooling and plotting of IPD in a meta-analysis. A non-commercial analysis package, 'SCHARP', which analyses each study, pools results and outputs tabulated results and forest plots for dichotomous, continuous and time-to-event IPD, is available free of charge to not-for-profit organizations. This SAS-based package has been developed by the Meta-analysis Group of the UK Medical Research Council Clinical Trials Unit. It is available from the authors, who can be contacted through the IPD Meta-analysis Methods Group (see Box 18.6.a).

18.5 Limitations and caveats

18.5.1 What an IPD review cannot fix

Although the IPD approach can help avoid problems associated with the analyses and reporting of studies, it cannot, generally, help avoid bias associated with study design or conduct. If there are such problems (which would also be reflected in study publications and any systematic reviews based upon them), the study may need to be excluded from the meta-analysis.

18.5.2 Unavailable studies

Obtaining IPD often enables inclusion of studies that could not be included in a standard systematic review because they are either unpublished or do not report sufficient information to allow them to be included in the analyses. This may help avoid many types of publication bias (Stewart 2002). However, one must ensure that by restricting analyses to those studies that can supply IPD, bias is not introduced through selective availability of study data.

The success and validity of the IPD approach requires that data from all or nearly all studies will be available. If unavailability is related to the study results, for example if investigators are keen to supply data from studies with promising results but reluctant to provide data from those that were less encouraging, then ignoring the unavailable studies could bias the results of the IPD review. If a large proportion of the data have been obtained, perhaps 90% or more of individuals randomized, we can be relatively confident of the results. However, with less information we need to be suitably circumspect in drawing conclusions. Sensitivity analysis combining the results of any unavailable studies (as extracted from publications or obtained in tabular form) and comparing these with the main IPD results are a useful aid to interpreting the data. Reports of IPD reviews that were unable to obtain IPD from all studies should state reasons why IPD were not available, and the likelihood of ensuing bias.

As for other types of Cochrane review, IPD meta-analyses should clearly state what studies were not included and the reasons why. If only a limited number of studies are able to provide IPD for analysis, then the value of the approach is questionable. Experiences in cancer have been good and in most cases perseverance has led to data being available from a high proportion of eligible trials. This can make it especially important to explore the ability and willingness of the primary investigators to supply IPD at an early stage in the project.

18.5.3 Deciding when an IPD review is appropriate

When initiating any systematic review it is useful to consider carefully which approach and which type of data will be most appropriate at the outset. Particular thought should be given to factors that are likely to introduce bias to the review. There may be cases where the benefits of obtaining IPD turn out to be marginal, and others where it could be vital.

18.6 Chapter information

Authors: Lesley A Stewart, Jayne F Tierney and Mike Clarke on behalf of the Cochrane Individual Patient Data Meta-analysis Methods Group.

This chapter should be cited as: Stewart LA, Tierney JF, Clarke M. Chapter 19: Reviews of individual patient data. In: Higgins JPT, Green S (editors), *Cochrane Handbook for Systematic Reviews of Interventions*. Chichester (UK): John Wiley & Sons, 2008.

Acknowledgements: We thank Paula Williamson for helpful comments on an earlier draft.

Box 18.6.a The Cochrane Individual Patient Data Meta-analysis Methods Group

The Individual Patient Data Meta-analysis Methods Group (IPD MA MG) comprises individuals who are involved or interested in the conduct of systematic reviews that include IPD and related methodological research. The Group aims to provide guidance to those undertaking IPD meta-analyses within Cochrane reviews.

Activities of IPD MA MG members include the following:

- Undertaking IPD meta-analyses.
- Undertaking empirical research, for example in the relative benefits of IPD meta-analyses compared with other forms of systematic review, and using information collected for IPD meta-analyses to explore whether aspects of design, analysis and reporting of randomized trials and systematic reviews may be sources of bias and heterogeneity.
- Helping authors of Cochrane reviews decide whether it would be appropriate for their systematic review to be conducted using IPD and, if so, to offer advice on how to do so.
- Offering training workshops at Cochrane Colloquia and disseminating training materials from these.
- Maintaining a register of reviews that have used (or will use) IPD and a database of methodological research projects and meta-analyses.

Web site: www.ctu.mrc.ac.uk/cochrane/ipdmg

18.7 References

Clarke 1998

Clarke M, Stewart L, Pignon JP, Bijnens L. Individual patient data meta-analysis in cancer. *British Journal of Cancer* 1998; 77: 2036–2044.

Higgins 2001

Higgins JPT, Whitehead A, Turner RM, Omar RZ, Thompson SG. Meta-analysis of continuous outcome data from individual patients. *Statistics in Medicine* 2001; 20: 2219–2241.

International Germ Cell Cancer Collaborative Group 1997

International Germ Cell Cancer Collaborative Group. International Germ Cell Consensus Classification: a prognostic factor-based staging system for metastatic germ cell cancers. *Journal of Clinical Oncology* 1997; 15: 594–603.

Simmonds 2005

Simmonds MC, Higgins JPT, Stewart LA, Tierney JF, Clarke MJ, Thompson SG. Meta-analysis of individual patient data from randomized trials: a review of methods used in practice. *Clinical Trials* 2005; 2: 209–217.

Stewart 1995

Stewart LA, Clarke MJ. Practical methodology of meta-analyses (overviews) using updated individual patient data. *Statistics in Medicine* 1995; 14: 2057–2079.

Stewart 2002

Stewart LA, Tierney JF. To IPD or not to IPD? Advantages and disadvantages of systematic reviews using individual patient data. *Evaluation in the Health Professions* 2002; 25: 76–97.

Tudor Smith 2005a

Tudor Smith C, Williamson PR. Meta-analysis of individual patient data with time to event outcomes. *International Conference of the Royal Statistical Society*, Cardiff (UK), 2005.

Tudor Smith 2005b

Tudor Smith C, Williamson PR, Marson AG. Investigating heterogeneity in an individual patient data meta-analysis of time to event outcomes. *Statistics in Medicine* 2005; 24: 1307–1319.

Turner 2000

Turner RM, Omar RZ, Yang M, Goldstein H, Thompson SG. A multilevel model framework for meta-analysis of clinical trials with binary outcomes. *Statistics in Medicine* 2000; 19: 3417–3432.

Whitehead 2001

Whitehead A, Omar RZ, Higgins JPT, Savaluny E, Turner RM, Thompson SG. Meta-analysis of ordinal outcomes using individual patient data. *Statistics in Medicine* 2001; 20: 2243–2260.