Published in final edited form as:

Psychol Methods. 2009 June ; 14(2): 81–100. doi:10.1037/a0015914.

Integrative Data Analysis: The Simultaneous Analysis of Multiple Data Sets

Patrick J. Curran and Andrea M. Hussong

University of North Carolina at Chapel Hill

Abstract

Both quantitative and methodological techniques exist that foster the development and maintenance of a cumulative knowledge base within the psychological sciences. Most noteworthy of these techniques is meta-analysis which allows for the synthesis of summary statistics drawn from multiple studies when the original data are not available. However, when the original data can be obtained from multiple studies, many advantages stem from the statistical analysis of the pooled data. The authors define integrative data analysis (IDA) as the analysis of multiple data sets that have been pooled into one. Although variants of IDA have been incorporated into other scientific disciplines, the use of these techniques are much less evident in psychology. In this paper the authors present an overview of IDA as it may be applied within the psychological sciences; a discussion of the relative advantages and disadvantages of IDA; a description of analytic strategies for analyzing pooled individual data; and offer recommendations for the use of IDA in practice.

The cornerstone of any field of scientific inquiry is the pursuit of a body of cumulative knowledge, yet the psychological sciences have often fallen short of this goal (e.g., Gans, 1992; Hunter & Schmidt, 1996; Meehl, 1978; Schmidt, 1996). This is not for want of trying. Both quantitative and methodological techniques have been developed to help build a cumulative knowledge base. Most noteworthy of these techniques is meta-analysis which allows for the synthesis of summary statistics drawn from multiple studies when the original data are not available (e.g., Cooper, in press; Glass, 1976; Rothstein, Sutton, & Borenstein, 2005; Smith & Glass, 1977). One of the original motivations for meta-analysis was that these techniques would further support the creation of a cumulative knowledge within the social sciences, particularly in psychology (e.g., Hunter & Schmidt, 1996; Schmidt, 1984). There is no doubt that meta-analysis has substantially advanced our science toward this goal.

Because the focus of meta-analysis is on the synthesis of summary statistics drawn from multiple studies, this approach is ideal when the original individual data used in prior analyses is inaccessible or no longer exists. However, as we discuss in greater detail below, there are many advantages to fitting models directly to the original raw data instead of synthesizing the relevant summary statistics when the original individual data are available for analysis (e.g., Berlin, Santanna, Schmid, Szczech, & Feldman, 2002; Lambert, Sutton, Abrams, & Jones, 2002). Recent developments within the scientific community, such as greater expectations for data sharing and better options for electronic data storage and retrieval, have increased the

Address for Correspondence: Patrick J. Curran, Department of Psychology, University of North Carolina, Chapel Hill, NC, 27599-3270, phone: 919.962.5235, fax: 919.962.2537, curran@unc.edu.

Publisher's Disclaimer: The following manuscript is the final accepted manuscript. It has not been subjected to the final copyediting, fact-checking, and proofreading required for formal publication. It is not the definitive, publisher-authenticated version. The American Psychological Association and its Council of Editors disclaim any responsibility or liabilities for errors or omissions of this manuscript version, any version derived from this manuscript by NIH, or other third parties. The published version is available at www.apa.org/journals/met.

potential for accessing original individual data for secondary analysis (i.e., the analysis of existing data). This in turn creates new opportunities for the development of alternative methods for integrating findings across studies by using original individual data to help overcome some of the unavoidable limitations of meta-analysis. (See Cooper and Patall, this issue, for a thoughtful comparison of the advantages and disadvantages of meta-analysis relative to the pooled analysis of raw data.)

Techniques for fitting models to pooled data go by a variety of names, none of which have been broadly adopted within the social sciences. Simply to offer a starting point, we will refer to this set of methodologies as *integrative data analysis* or IDA.¹ We chose the term *integrative* over options such as pooled, simultaneous, unified, or concomitant, to highlight our goal of creating "a whole by bringing all parts together" which is a common definition of *integrate* (e.g., American Heritage Dictionary of the English Language, 2009). Interestingly, IDA has been used in other areas of scientific inquiry for more than a decade. For example, IDA has been used in medicine to examine the efficacy of medications versus cognitive behavior therapy for severe depression (DeRubeis, Gelfand, Tang & Simons, 1999); to evaluate clinical trial outcomes for treatment of Alzheimer's Disease (Higgins, Whitehead, Turner, Omar, & Thompson, 2001); to examine the relation between fat intake and the risk of breast cancer (D. Hunter et al., 1996); to study the pharmacogenetics of tardive dyskinesia (Lerer et al., 2002); and to examine the relation between height, weight and breast cancer risk (van den Brandt et al, 2000).

Despite the broader use of IDA techniques in other disciplines, such applications are relatively novel within the behavioral sciences in general and within psychology in particular (but see Lorenz, Simons, Conger, Elder, Johnson & Chao, 1997; McArdle, Hamagami, Meredith & Bradway, 2000; and McArdle, Prescott, Hamagami, & Horn, 1998, for notable exceptions). One reason behind the slow adoption of these techniques may be the significant challenges that psychologists face in pooling across studies that are highly heterogeneous in their methodology, even when these studies examine the same topic. Differences between studies in sampling techniques and frame, historical timing, design characteristics and measurement create seeming barriers to study comparison and integration. However, by incorporating information about such between-study heterogeneity into our techniques for study integration, our conclusions may be more generalizable and our progress as a science more cumulative. Thus, IDA strives to capitalize upon such between-study heterogeneity to not only better understand findings across existing studies (i.e., study integration) but also to probe meaningful sources of between-study variability that may contribute to, and thus inform theories about, key psychological phenomenon (i.e., study comparison).

The topics that underlie IDA are both broad and complex and a comprehensive treatment is beyond the scope of any single manuscript. As such, our intent here is rather modest. Specifically we offer a general discussion of the core issues that typically arise in applications of IDA for study integration in the psychological sciences. These topics and our guiding perspective on IDA are largely culled from our experience in using these techniques on a project that we call Cross Study. Cross Study involves the integrated analysis of three independent longitudinal studies of children of alcoholic parents and matched controls. These data sets are unique in their excellent retention, breadth of measurement, and sampling of non-treatment samples. Nonetheless, the three studies differ in many respects, such as geographical location, developmental coverage, measurement, and assessment modality. Because applications of IDA are necessarily idiosyncratic to the theoretical questions and sample characteristics at hand, we wholly acknowledge that our experiences on Cross Study have shaped our views of IDA, and

¹We realize that psychology needs another acronym like a hole in the head, but we also believe that the set of techniques we explore here are in need of some shared terminology; thus IDA.

this in turn is reflected throughout our work here. However, it is this same sensitivity of IDA to the specific theoretical and methodological context that makes this both a broad topic eluding simple description and a flexible, informative set of techniques that is critically needed in our field.

In the current paper, we build on our work with Cross Study and aim to further establish IDA as a potential tool for pursuing and fostering a cumulative knowledge base in our field. We begin with a discussion of exactly what IDA is and what advantages IDA offers when appropriate data are available for analysis. Next we detail potential influences on between-sample heterogeneity that may serve either as nuisance factors when study integration is the goal or, in many instances, as sources of variance that offer novel insights about why a phenomenon may show study-to-study differences. We then explore general analytic strategies that address between-study heterogeneity. We conclude with future directions for research and recommendations for the use of these techniques in practice.

What Exactly is Integrative Data Analysis?

Because methods for pooling existing data can vary across discipline, we begin by offering a specific definition of integrative data analysis within the psychological sciences. IDA is the statistical analysis of a single data set that consists of two or more separate samples that have been pooled into one. What constitutes "separate" can sometimes be unambiguously determined and other times cannot. In some cases, minor design differences between samples may be present. For example, separate samples may be collected within a multi-site or rolling recruiting single-site design in which key design characteristics are held constant (e.g., recruitment, procedures, measurement) yet each study is conducted in a different setting (e.g., different hospitals or regions of country) or across different time periods (e.g., as recruitment rolls across different school years or birth cohorts). These separate samples are then pooled for analysis with some control for site or cohort differences (e.g., Kaplow, Curran, Dodge, and the Conduct Problems Prevention Research Group, 2002; Stark, Janicke, McGrath, Mackner, Hommel & Lovell, 2005). In other cases, many design differences between samples may be present. For example, multiple separate samples may be each collected as part of a different independent study that was conducted at different historical times using different sampling mechanisms, experimental procedures, and psychometric instruments. Thus what constitutes a "separate" sample ultimately resides on a continuum.

Although cognizant of the continuum of designs to which IDA may be applied, our focus here is explicitly on the latter situation; namely, where multiple samples are drawn from independent existing studies and pooled into a single data set for subsequent analyses. This was precisely our experience in Cross Study in which our focus was on data pooled from samples that were drawn from three independent studies in which participants differed from one another in both theoretically (e.g., status of family psychopathology) and methodologically (e.g., developmental level, measurement, recruitment strategies) meaningful ways. We believe that the broadest potential for future applications of IDA in psychological research relates to the pooling of data that are drawn from two or more existing studies. For this reason, we focus the remainder of our discussion on this topic. To highlight the potential applicability of these techniques in psychological research settings, we describe our work on the Cross Study and present exemplar findings that we believe could only be obtained using IDA.

A Motivating Example: Cross Study

Cross Study is an ongoing NIH-funded project in which data are pooled from three existing longitudinal studies of adolescent development with a particular focus on identifying developmental pathways leading to substance use and disorder. All three studies oversampled offspring who had at least one biological, alcoholic parent (i.e., children of alcoholics, or COA)

and included matched controls of offspring who had neither biological parent diagnosed as alcoholism (i.e., controls). The first study, the Michigan Longitudinal Study (MLS; Zucker, Fitzgerald, Refior, Puttler, Pallas, & Ellis, 2000), has amassed a broad data archive beginning with a sample of 2–5 year olds who were assessed over four waves (at the time) into early adulthood. The second study is the Adolescent and Family Developmental Project (AFDP; Chassin, Rogosch, & Barrera, 1991) which began interviewing families when adolescents were aged 11–15, with ongoing assessments continuing well into adulthood over five waves. The third study is the Alcohol, Health and Behavior Project (AHBP; Sher, Walitzer, Wood, & Brent, 1991) which began intensive assessments with college freshmen and has continued to survey participants over six waves into their thirties. Together, these three studies span the first four decades of life when early risk factors for later substance outcomes first emerge (childhood), when substance use initiation typically occurs (adolescence), when peak rates of substance use disorders are evident (young adulthood), and when deceleration in substance involvement is first apparent (adulthood). Table 1 presents a summary of pooled sample as a function of study membership and chronological age. Each cell in the table identifies the number of individuals assessed in a given wave of a given study at a given age. The column totals identify the total number of individuals assessed at a given age pooling across study and wave.

The Cross Study presented many methodological challenges within the context of studying early symptom trajectories associated with various forms of parent alcoholism. Notably, the three contributing studies varied substantially in design, including issues of participant recruitment, assessment strategies and instrumentation (see Table 2 for a summary of design characteristics). One of our goals in Cross Study was to use IDA to control for such between-study differences so as to examine our substantive questions of interest. The pursuit of this goal via IDA permitted us to study a longer developmental period than in any one study, larger subsamples of families with specific forms of alcoholism, and trajectories of symptomatology in analyses with greater statistical power.

Our approach to IDA in this work is best exemplified in our study combining data from all three studies to examine trajectories of internalizing symptomatology between 10 and 33 years of age (Curran, Hussong, Cai, Huang, Chassin, Sher & Zucker, 2008). The pooled sample consisted of a total of 1,827 individual participants (512 drawn from the MLS, 830 from the AFDP, and 485 from the AHBP). Each individual provided between one and five repeated measures resulting in a total of 7,377 person-by-time observations. We operationalized internalizing symptomatology using 27 dichotomous self-reported items; of these, 12 were drawn from the anxiety and depression subscales of the Brief Symptom Inventory (BSI; Derogatis & Spencer, 1982) and 15 from the anxiety and depression subscales of the Child Behavior Check List (CBCL; Achenbach & Edelbrock, 1981). Six of these items were unique just to the BSI, nine were unique just to the CBCL, and six items were shared between the BSI and CBCL. Given the six shared items, we used 21 unique items to define internalizing symptomatology (see Table 1 in Curran et al., 2008, for details). Of these 21 items, all were administered in the MLS, 10 of the 15 CBCL and none of the BSI items were administered in the AFDP, and all 12 BSI items but none of the CBCL items were administered in the AHBP. We applied a series of item response theory (IRT; e.g. Thissen & Wainer, 2001) models to calculate scale scores for each individual at each time point. Importantly, these scale scores were all anchored to a shared metric regardless of to which set of items the individual responded or to which study the individual belonged. Finally, we fitted a series of multilevel piecewise growth models to examine the fixed and random effects characterizing the developmental trajectories of internalizing symptomatology between ages 10 and 33.

We have used these same IDA methodologies to test many other theoretical questions using data pooled from two or three of the component studies. For example, in Hussong, Cai, Curran, Flora, Chassin and Zucker (2008) we used IDA disaggregate the distal, proximal and time-

varying effects of parental alcoholism on children's internalizing symptomatology between ages two and 17. In Hussong, Wirth, Edwards, Curran, Chassin and Zucker (2007) we used IDA to examine the relation between the number of alcoholic parents in the family, the specific subtype of parental alcoholism, and the gender of the child in the prediction of developmental trajectories of externalizing symptomatology between ages 2 and 17. And in Hussong, Flora, Curran, Chassin and Zucker (2008) we used IDA to examine the unique predictability of trajectories of child internalizing symptoms from parental alcoholism above and beyond the parental co-morbid diagnoses of depression and antisocial personality disorder. Over the course of this work, we believe that we have been able to use IDA to empirically test hypotheses in ways that would not otherwise be possible. However, this has not been without a cost. We have addressed a seemingly endless parade of challenges, some foreseen and others not, and some that were ultimately surmountable and others that were not. We draw upon these experiences to organize and navigate a more general discussion about the potential advantages and disadvantages IDA may hold for other applications in the social sciences.

Potential Advantages of Integrative Data Analysis

Of course IDA is not universally appropriate for pooling data from any two independent studies. Nor is IDA intended to replace meta-analysis or any other method of research synthesis. Rather, IDA is an additional tool that may be used for the purposes of study integration and comparison given conducive research contexts. Indeed, under the proper conditions, IDA offers a host of significant advantages compared to available alternatives. We believe that there are seven specific advantages of IDA that are particularly salient when used within many areas of psychological inquiry.

Replication

IDA provides a direct mechanism to test whether a set of findings replicate across independent studies by explicitly quantifying effects that represent tests of our hypotheses both within and between studies.² Unlike other approaches to research synthesis that are based on summary statistics, IDA can directly model potential influences on between-study heterogeneity at the level of the observed data. This permits explicit analysis of study equivalence at multiple levels of design that can often incorporate differences in sampling, geographic region, history, assessment protocol, psychometric measurement, and even hypothesis testing. This advantage also makes IDA well-suited to the task of testing novel hypotheses that may not have been considered in the original within-study analysis of the data. Thus, IDA may provide tests of replication of novel hypotheses within a single analysis of independent studies. Moreover, IDA permits an exploration of between-study differences that helps mitigate the need for creating new studies designed to resolve conflicting findings across studies posited to result from between-study design differences.

Increased statistical power

IDA has the potential to provide substantial increases in statistical power for testing research hypotheses through the combination of multiple individual data sets. It is well known that most research applications within psychology are often chronically under-powered such that there is an unacceptably low probability that a given effect will be found if that effect truly exists in the population (Cohen, 1992; Maxwell, 2004). However, when multiple independent samples are combined, there is often a marked increase in power when testing the same hypotheses based upon the aggregated versus independent sample.

²Because of the close relation between *study* and *sample* within our focus on IDA (in which a single study results in a single sample of data) we use these terms interchangeably throughout our paper.

Increased sample heterogeneity

A closely related advantage is that IDA often allows for a more heterogeneous pooled sample. For a variety of reasons, many studies in psychology use sampling methods that result in the under-representation of potentially important subgroups in the population of interest (e.g., groups based on gender, race, socioeconomic status, age). However, a pooled sample that is aggregated across multiple studies, each of which may have been conducted in a different geographic setting or using a different sampling mechanism, allows for more distinct groups or individual characteristics to be simultaneously considered. Moreover, given adequate sample representation within studies, group comparisons may be possible within IDA that are not possible due to small samples sizes within the individual studies. This in turn increases the external validity of the IDA findings fitted to the aggregated data.

Increased frequencies of low base-rate behaviors

This same logic is evident as an advantage of IDA in pooling studies of low base-rate behaviors. For example, each contributing study may have 5% of the sample reporting heavy drug use. Although such behaviors will retain an overall low base rate in the pooled IDA analyses (e.g., assuming equal sample sizes, the aggregate sample would still reflect 5% heavy drug use), the overall absolute number of individuals engaging in the behavior will necessarily be greater in the pooled sample relative to the individual contributing studies (e.g., there may be 20 of 400 individuals reporting heroin use in a single study, but 80 of 1600 individuals when pooling four studies). As a result, the stability of model estimation is improved, the influence of extreme observations is reduced, and more complicated models can be fitted than would otherwise be possible within the individual studies.

Broader psychometric assessment of constructs

IDA often results in a broader and more rigorous psychometric assessment of the key theoretical constructs under study. In any single-study design, theoretical constructs are typically assessed using a discrete set of items that are shared across all members of the sample (e.g., all subjects respond to the same 10-item scale assessing depression). A common challenge in many areas of psychological research is the need to reconcile the wide array of operationalizations of our constructs across studies. This is a seeming limitation for study-to-study comparison but is in turn a distinct advantage for increased construct validity in IDA. Typically researchers select the psychometric instruments for any given study based on the specific characteristics of their sample (e.g., age, gender, ethnicity), although this in turn limits the generalizability of the subsequent results to the characteristics of the sample under study. Yet when multiple samples are combined, the psychometric assessment of a given construct can often be substantially broadened by incorporating the multiple methods of assessment that were used in each individual study. This often results in much stronger psychometric properties of the assessment of the theoretical constructs in the aggregated sample compared to any given single sample.

Extended period of developmental study

Although researchers may use IDA to pool data drawn from either cross-sectional (i.e., single time-point) or longitudinal (i.e., multiple time-point) studies, the pooling of longitudinal studies presents several distinct advantages. Most importantly, a single sample is obviously limited to the age range observed within that study. However, when multiple longitudinal studies are combined, a much broader developmental period can be considered, given overlapping age ranges across the set of contributing studies. For example, in Cross Study the MLS assessed children between the ages of 2 and 24; the AFDP between ages 10 and 34; and the AHBP between ages 17 and 40. Under many situations IDA can allow for inferences across the entire range of ages two through 40 even though each individual study followed participants for a fraction of this time (see, e.g., Table 1). This advantage is then amplified by the combined

psychometric assessments of the theoretical constructs that were deemed optimal within the particular age range under study.

Support of data sharing and building a cumulative science

Finally, IDA is directly supportive of recent practical concerns about efficiency in psychological research, namely, increased calls for data sharing and decreasing resources available to support new research endeavors. First, the issue of data sharing has been addressed both at the level of federal funding mechanisms (e.g., both the National Science Foundation and the National Institutes of Health have formal policies regarding data sharing) and at the level of appropriate ethical practices in research (e.g., Section 8.14 of the Ethical Principles of Psychologists and Code of Conduct, American Psychological Association, 2003). Not only are individual researchers increasingly called upon to share data but technological advances further support these efforts through the accessible electronic storage and distribution of even the largest data sets. Second, in recent years, non-defense related federal funding for research and development has stagnated and sometimes decreased (American Association for the Advancement of Sciences, 2008).³ As a result, the analysis of existing data is an extremely cost efficient mechanism for conducting research (Kiecolt & Nathan, 1985). This efficiency is further realized by considering not just one but multiple existing samples of data. Thus IDA meets several practical needs in terms of data sharing and maximizing limited resources.

Summary

Despite the many potential advantages to adopting an IDA framework for data integration, this remains an uncommon practice within the psychological sciences. One potential reason may be that conducting such analyses can be an extremely complex and challenging task. Key practical issues associated with data acquisition and data management are often eclipsed by a multitude of difficulties that arise from sometimes substantial study-to-study differences. Whereas the initial temptation might be to embark on IDA with the desire to minimize between-study heterogeneity (i.e., to attempt to carefully select contributing studies that are as similar as possible), we believe that certain types of between-study differences can actually help us to simultaneously understand both within-study and between-study differences in our findings. We next turn to a closer examination of potential sources of between-study heterogeneity that are likely to arise in many areas of psychological research.

Potential Sources of Between-Study Heterogeneity

When two or more independent data sets are pooled within an IDA framework, it is important to closely consider the combination of study characteristics that uniquely define each individual study. For the purposes of study integration, we are typically most interested in controlling for these differences so that we may obtain findings that are maximally externally valid (e.g., Shadish, Cook & Campbell, 2002). For the purposes of study comparison, we may be directly interested in between-study heterogeneity as a means of testing the generalizability of our findings. Whether for purposes of control or exploration, identifying important sources of between-study heterogeneity is a critical aspect of IDA. This process is complicated because some sources of between-study heterogeneity are confounded and thus cannot be disentangled (i.e., geographic influences and ethnicity cannot be distinguished in pooled analyses of a study of Caucasian youth in Indiana and a separate study of Latino youth in Arizona). Another complicating factor is that there are multiple sources of heterogeneity that must be considered simultaneously, but many of these may interact with one another in potentially complex ways. Fortunately, it is not necessary in IDA to independently model and identify all sources of

³This does not include the effects of the 2009 *Recovery and Reinvestment Act* federal stimulus package that was passed in the same week that we completed this paper.

between-study heterogeneity for purposes of study integration. Rather, we can use techniques that control more globally for between-study differences to obtain findings across studies and to determine in which studies these findings hold.

However, for purposes of study comparison, we can use information about between-study heterogeneity to model potential moderating influences on the generalizability of our findings to the extent that we are aware of and able to identify and operationalize these sources of variance. Importantly, understanding the points of convergence and divergence in findings among a set of studies can often inform our understanding of findings within the individual studies themselves. As we will describe later, not only can between-study heterogeneity be directly factored into many IDA applications, but some of these study-to-study differences may be of substantive interest in their own right. Indeed, this latter point is what makes IDA such an intriguing endeavor. Although there are many sources of between-study heterogeneity that we might consider, here we focus on five: sampling, geographical region, history, design characteristics, and measurement.

Heterogeneity Due to Sampling

As we will explore throughout our paper, one of the key benefits of IDA is that it prompts us to think more carefully about important issues that typically receive limited or no attention in single-sample analysis. One prime example is sampling. By sampling we mean the implicit or explicit mechanism by which a sample of observations is drawn from a targeted population with the goal of making inferences back to a population (e.g., Cochran, 1977). Drawing on our own personal experiences spanning the past several decades, only on one occasion have either of us been asked to address sampling issues within any of our single-study papers or grant applications; in contrast, this issue has been raised in some form or another on every single manuscript that has come out of Cross Study. Sampling is clearly an important issue within all areas of psychological research, but this is particularly salient within the IDA framework.

Briefly, there are two general approaches to sampling. The first approach is *probability sampling* in which all members of a defined population have some known probability of being selected into the sample; examples of this type include simple, systematic, cluster, and stratified sampling. Because the probability of selection is known, one or more sampling weights are typically available for each individual included in the sample (e.g., Pfeffermann, 1993). The second approach is *non-probability sampling* in which some (or, more typically, all) members of the population have an unknown probability of being selected into the sample; examples of this type include convenience, snowball, and quota sampling. Because the probability of selection is unknown, no sampling weights are available for individuals included in samples obtained using non-probability procedures. The majority of research conducted in the psychological sciences is characterized by non-probability sampling (Sterba, 2009).

The statistical methods needed for making valid inferences from the sample back to the population depends upon what sampling mechanism was used to obtain the sample. There are two approaches that characterize nearly all research applications in the social sciences. The first is the *model-based* procedure proposed by Fisher (1922) that can be used to make population inferences based upon samples that were drawn using a non-probability framework.⁴ Fisher invoked three conditions to allow for this: a structural model must be hypothesized; a parametric distribution must be assumed; and any design characteristics impacting sample selection must be included in the model (e.g., over-sampling or cluster-sampling). The second approach is the *design-based* procedure developed by Neyman (1934) that can be used to make

⁴Model-based procedures can also be used with probability samples assuming that the selection mechanism is explicitly incorporated into the fitted model (Sterba, 2009)

population inferences based upon samples that were drawn using a probability sampling framework. Neyman developed these methods to overcome what he viewed as inappropriate subjectivity that was inherent in the model-based approach, particularly with respect to the selection of the hypothesized model and assumed distribution. In contrast to the model-based approach, the design-based approach allows for valid inferences to be made through the direct incorporation of the individual-specific sampling weights into the statistical analysis.⁵

Most important for our discussion here, combining data from two or more studies within the IDA framework provides an exciting opportunity to examine these same sampling issues much more carefully than was previously possible. More specifically, IDA offers the potential to conduct direct empirical evaluations of the effects of heterogeneity in sampling mechanism in the pooled sample that cannot be conducted within any single-sample analyses. From a practical standpoint, the first step is to determine what type of sampling mechanism was used within each study that is to be pooled within the IDA. In some situations this mechanism may be unambiguous, such as if data were drawn from probability design-based studies such as Add Health (Harris, Halpern, Entzel, Tabor, Bearman, & Udry, 2008) or the National Longitudinal Survey of Youth (Bureau of Labor Statistics, 2002); for data such as these, one or more sets of sampling weights, strata indicators, and cluster indicators will likely be available that can potentially be used in the pooled IDA. However, in other situations the sampling mechanism may be less clear, such as if data were drawn from a non-probability design in which subjects were sampled from an undergraduate psychology subject pool or a volunteer sample was identified via publically posted fliers. More often than not, the available data sets will likely fall between these two extremes. For example, all three contributing samples used in Cross Study were non-probability samples, yet all three oversampled on certain known demographics (e.g., parental alcoholism) and all three incorporated strict inclusion and exclusion criteria. This information can then be directly incorporated into a model-based framework that fulfills Fisher's necessary conditions for valid inference.

The critical point to appreciate about the role of sampling in IDA is that our goal is not necessarily to unequivocally establish that all of the component samples are precisely equivalent with respect to sampling prior to proceeding to tests of our substantive questions. (This is sometimes a misplaced goal in IDA; namely, to only include samples that are deemed functionally identical to one another). Instead, we must test for potential differences across the samples, to incorporate adjustments for these differences, and, if possible, to come to some understanding about why such differences exist. Indeed, meaningful between-study heterogeneity with respect to sampling mechanism may not only be incorporated within the IDA framework, but it may provide a unique opportunity to empirically evaluate the role of sampling in ways that would not be possible within typical single-study designs.

Heterogeneity Due to Geographical Region

Although in principle two or more studies might have been independently conducted within the same geographical region, it is much more likely that this is not the case. Indeed, even if two independent studies were conducted in New York City, the samples may not be comparable as a function of the specific borough in which they were conducted; and if they were conducted in the same borough, they may not be comparable as a function of neighborhood; and so on. To further complicate matters, in most IDA applications we typically find that sampling mechanism and geographical region are not only closely intertwined, but are often completely confounded. For example, if data are drawn from three independent studies, it is likely that not only did each individual study incorporate a unique sampling mechanism but each was also

⁵For more detailed discussions about the similarities and differences between model-based and design-based procedures as well as current hybrid designs that combine the two approaches see Guo & Hussy (2004), Little (2004), Lenhard (2006), Muthén and Satorra (1995) and Sterba (2009).

located in a distinctly different geographical region. In such situations it may not be possible to disentangle differences due to sampling from differences due to region. One exception is the use of multiple sites within a single study design. When data are pooled from a multi-site study, so called "site differences" offer unique insights into heterogeneity in characteristics associated with region while holding the sampling mechanism constant. However, in the majority of IDA applications we must jointly consider sampling and geographical location.

When evaluating potential between-study heterogeneity associated with geographical location, it is important to consider which characteristics of the specific location may be responsible for the observed differences across the studies. There may be nothing inherently meaningful in the direct comparison of a sample of individuals drawn from Denver with a comparable sample of individuals drawn from Indianapolis. Instead, the challenge is to identify the specific characteristics that serve to distinguish Denver from Indianapolis that are manifested within the data set. There are likely a large number of such characteristics including ethnic composition, median income, availability of social services, per capita rates of crime, seasonal weather patterns, and proximity to neighboring urban centers, to name just a few. Of course it may be difficult to even identify a discrete region within which a given study was conducted. For example, one study may have recruited subjects from all incoming college students at a single university whereas another study may have recruited subjects from incoming college students at all the universities within a single state. In this case there is not a comparable geographical location to contrast the two samples; all that might be deduced is that the samples were obtained from two different sources, but it is not possible to move beyond this gross level of assessment.

Heterogeneity Due to History

Whereas heterogeneity due to region is associated with potential differences in place, heterogeneity due to history is associated with potential differences in time. We use the term history here to capture the essence of this concept as it arises within quasi-experimental designs (e.g., Shadish et al., 2002). Broadly speaking, history as a threat to internal validity refers to any events that might have occurred during a study that could have accounted for an observed effect. Although this concept is implicitly longitudinal (i.e., an event that occurred during a study), history can also play an important role when conducting an IDA of pooled cross-sectional (i.e., single time-point) data sets. For example, if two cross-sectional studies were conducted 10-years apart, there could be historical influences that differentially impacted each of the two studies (e.g., two studies on general anxiety conducted pre- and post-9/11; two studies on minority aspirations conducted pre- and post-President Obama).

A key question that arises within the IDA framework is whether some effect that is observed in the pooled sample can be at least partially attributable to between-study heterogeneity associated with historical period. Although one can conceive of many different measures of time (e.g., chronological age, time since diagnosis, age of a particular event), here we are primarily focused on calendar time. In other words, are there influences associated with the particular year (or month or day) on which a given subject was assessed? Just as two or more independent studies are likely conducted using different sampling mechanisms within different geographical regions, it is also likely that the studies were conducted within different historical periods. It is thus important to consider the potential ways in which the effects of history might be manifested in both cross-sectional and longitudinal IDA applications. Although historical influences can be present when using IDA to evaluate pooled cross-sectional data, this tends to be a far greater challenge when pooling two or more independent longitudinal data sets.

When pooling cross-sectional data, one must evaluate the comparability of two samples of individuals who were assessed at different points in time. For example, in a cross-sectional IDA of adolescent drug use one might need to consider the impact of historical context as a

function of whether the subjects were assessed in 1970, 1980 or 1990. That is, did societal norms, legal sanctions, drug processing, and popular routes of administration differ across these three decades in a way that might influence the very meaning of drug use within each of the independent samples. Epidemiological data shows that the annual prevalence of illicit drug use among graduating high school seniors was 54% in 1979, 35% in 1989, and 42% in 1999 (Johnston, Bachman, & O'Malley, 2006); given this, care must be taken to not treat the absolute measure of drug use as necessarily equivalent when pooling data that was collected over this 20-year period. Drug use is just one example, and similar issues can arise across a whole host of psychologically relevant outcomes that might be of interest with a cross-sectional IDA application.

When pooling longitudinal data, outcomes are embedded in time which even further complicates potential history effects. In longitudinal IDA, heterogeneity across studies must then be considered in over-time trajectories rather than cross-sectional levels of behavior over time. Extending the cross-sectional example above, we might need to compare developmental trajectories of drug use for individuals who matured through adolescence during the 1970's in one study, the 1980's in a second study, and the 1990's in a third. Interestingly, this issue is closely related to the classic age-period-cohort distinction first raised nearly 50 years ago in developmental psychology (e.g., Schaie, 1965). Briefly, Schaie (1965) argued that to understand the development of an individual over time, one must simultaneously take into account the individual's chronological age, the historical period in which the individual was assessed, and the birth cohort to which the individual belonged. Although initial attempts were made to simultaneously disentangle these three influences, later work argued that knowledge of any two dimensions of time defined the third (e.g., knowing the age of the individual when they were assessed and the year in which they were assessed in turn defines the cohort in which they were born; Palmore, 1978; Schaie, 1994). Thus, we must often choose the historical unit of interest (e.g., age, period, cohort) given our methodological and theoretical application. As such, a common goal of an IDA of pooled longitudinal data sets is to evaluate and potentially control for between-study heterogeneity associated with the historical period, however, indexed, during which a set of repeated observations were obtained.

Heterogeneity Due to Other Design Characteristics

Another likely source of between-study heterogeneity is study-to-study differences in characteristics of the study design. Such design characteristics may include methods of data collection and sample retention, which in turn exert an important impact on the sampled data. As a sobering example,⁶ Harford (1994) found that discrepancies in the order and style of two items assessing alcohol use at two different periods of time resulted in a large and observable change in the reported levels of heavy alcohol use. This demonstrates that something as simple as altering the order of presentation of two items at two points in time can lead to substantial differences in how subjects respond to the items even within the very same study. Given this, care must clearly be taken in pooling across multiple studies that may differ in structural design characteristics.

Yet here lies another challenge. It is typically possible to generate an extraordinarily long list of potentially meaningful differences between all of the contributing studies and this list grows exponentially as the number of individual studies increases. Examples are abundant. Were subjects assessed face-to-face or via telephone or internet? If face-to-face, was the assessment conducted at the home of a subject or was the subject brought into a controlled facility? Was the assessment conducted as a personal interview or was it computer-based? Was the same assessment battery used at each time period or were items deleted and others added over time?

⁶Pardon the pun.

And on and on and on. This list-generation exercise is not dissimilar to "medical school hypochondriasis" in which first-year medical students diagnose themselves as having each new condition about which they learn in class.

Clearly there are an unmanageable number of differences across studies in design characteristics. In most typical IDA applications these differences are confounded with one another. Thus, it is not only unrealistic but also not useful to exhaustively identify, track and code the entire set of differences in design characteristics across the set of contributing samples. A more useful goal is to identify those specific characteristics that are thought to be most salient for the given application at hand. The selection of these characteristics can be guided by the same factors related to drawing valid inferences from between-group single-study designs (Shadish et al., 2002). For example, if studying a sensitive topic such as illegal or high risk behaviors, were individual assessments conducted using face-to-face interviews or using confidential computer-assisted personal interview methods. If studying diagnostic criteria, were all items presented to all subjects or could individual subjects "skip out" of a set of items that were deemed irrelevant as a function of the individual's response to some qualifying item. If studying in-group / out-group behavior, were subjects assessed individually or in the presence of other subjects? The goal here is not to develop an all-inclusive catalog of potential design-related differences. Instead, the goal is to identify those study-to-study characteristics that are most likely to be related to the specific constructs that are under study. These measures can then be included in the analytic strategies we describe later.

Heterogeneity Due to Measurement

We have saved what is often the most important and yet most challenging source of between-study heterogeneity for last: heterogeneity due to measurement. Here we use the classic definition offered by Stevens (1946) that broadly establishes measurement as "...the assignment of numerals to objects or events according to rules" (p.677). Of course measurement plays a critical role in nearly all areas within the psychological sciences, and the field of psychometrics is characterized by more than a century of work in this domain (e.g., Cudeck & MacCallum, 2007). In most IDA applications, a key goal is to optimally capture the measurement of specific theoretical constructs both within each sample individually and, more importantly, within the aggregated sample as a whole. Indeed, measurement might be considered the most fundamental source of between sample heterogeneity in IDA because the reliability and validity of the analytic results rely directly on the reliability and validity of the contributing measures drawn from each individual study.

As with other factors that arise in IDA, considering issues of measurement within the aggregated sample often prompts us to more closely consider these same issues within each contributing study. Our motivating goal here is to develop an analytic framework that allows us to create a valid and reliable aggregate measure that is sensitive to potential study differences on dimensions such as design characteristics, specific items administered, subject age, and calendar year. As we will see, in many IDA applications we will be able to implement such a framework, although naturally this comes at a cost. Simultaneously balancing the need to maximize the quality of measurement and to minimize the associated cost is one of the most salient challenges of IDA. Here we briefly explore two closely related measurement issues that are often of greatest importance in IDA: measurement invariance and measurement comparability.⁷ (Further discussions of these important issues are presented in Bauer & Hussong, this issue, and McArdle et al., this issue, and in other manuscripts from our project

⁷The term *measurement comparability* is not widely used in the traditional invariance literature, but this becomes an important issue when considering the psychometric equivalence of the assessment of a theoretical construct across multiple independent studies.

including Curran, Edwards, Wirth, Hussong & Chassin, 2007, Curran et al., 2008, and Flora, Curran, Hussong & Edwards, 2008.)

To think more specifically about these measurement issues, let us again consider the Cross Study project in which we estimated developmental trajectories of internalizing symptomatology using data pooled from three separate studies (Curran et al., 2008). For our measurement models we used 21 items drawn from the BSI and CBCL. In the first study subjects responded to all 21 items from the BSI and CBCL; in the second study subjects responded to 10 of the 15 CBCL items and none of the BSI items; and in the third study subjects responded only to the BSI items. Based on our theoretical perspective, we believe that there is some underlying individual-specific latent propensity for a child to experience depressive symptomatology and that this latent propensity is manifested in the child's response to a particular set of items. In our study we had a set of internalizing symptomatology items shared across all three studies as well as items that were unique to a given study or studies (e.g., three items were administered in just one study, 14 items in two studies, and four items in all three studies). Despite these differences, all three studies are attempting to assess precisely the same underlying latent construct. The concept of measurement invariance applies to the studies that share an item set; the concept of measurement comparability applies to the studies that use a unique item set.⁸

Measurement invariance—Generally speaking, measurement invariance addresses the extent to which a set of items reliably and validly assesses an underlying construct in a similar (if not identical) fashion across groups or over time. Classic examples of measurement invariance include examining gender differences or racial differences in the psychometric properties of a particular measure (e.g., Rusticus, Hubley & Zumbo, 2008) or examining developmental differences in the expression of a set of behaviors over time (e.g., Pentz & Chou, 1994). Nearly all prior research on measurement invariance has focused on differences across groups or over time within a single sample of observations (but see Nesselroade, Gerstorf, Hardy & Ram, 2007, for a recent discussion of idiographic approaches to measurement invariance). However, we can draw on this same literature to apply concepts of measurement invariance across groups and over time not only within a particular study, but more importantly between two or more independent studies. Indeed, given sufficient empirical data we can examine the interaction between measurement invariance across group or over time with study group membership. For example, we were able to explicitly compare differences in the magnitude of non-invariance in the assessment of depressive symptomatology as a function of age and gender across our three contributing studies (Curran et al., 2008; Flora et al., 2008).

Measurement comparability—The topic of measurement invariance itself is extremely broad and seminal contributions include Thurstone (1947), Horn and McArdle (1992), Millsap (1995; 1997), and Meredith (1993). There are actually many different types of measurement invariance that have been proposed over the years and different terminologies have been used to describe similar types (Vandenberg & Lance, 2000). To better understand these types of invariance, it is helpful to approach measurement using concepts drawn from factor analysis. Within a factor analytic framework, a set of items (or indicators) are used to define an underlying latent variable, the existence of which is believed to have given rise to the pattern of observed correlations among the items (e.g., Bollen, 2002). The items are linked to the underlying factor via the factor loadings that serve as partial regression coefficients that index the extent of change in the individual items resulting from a one-unit change in the latent factor. Each item is also defined by an item-specific intercept and residual variance and each factor

⁸This distinction between invariance and comparability is an oversimplification in that both concepts will ultimately apply to both situations. However, we believe this distinction offers a helpful starting point to understand the core issues at hand.

(or latent variable) is defined by a mean and variance. The different types of invariance conditions (e.g., strong, weak, configural, etc.) are then defined with respect to the equality (i.e., invariance) or inequality (i.e., non-invariance) of these model parameters.

At the extreme, strict invariance (Meredith, 1993; also called complete invariance, Millsap, 1995) reflects that all the parameters that define the measurement model are equal across group or over time. In our earlier example of depression, a measurement structure that is strictly invariant indicates that boys and girls express depression identically and thus a single set of parameters validly defines the assessment of this construct for both genders. However, if one or more parameters are found to differ across group or over time, then the items are characterized by weaker forms of invariance (e.g., configural, pattern, or scalar invariance). As such, given a partially non-invariant measurement structure, we would conclude that boys and girls express depression in a functionally different way and that, most importantly, different parameters are needed to validly define the assessment of depression within each gender. If non-invariance is ignored (whether within a single sample or within the pooled sample) we cannot unambiguously establish whether an observed relation between depression and some other construct is valid or is instead an artifact of imposing an improper measurement model; see Meredith (1993), Millsap (1995), and Vandenberg and Lance (2000) for further details.

The importance of this situation within the IDA framework is highlighted in an extreme hypothetical example in which there is strict measurement invariance across all observed groups and over all time points within each independent contributing study. Although there is unequivocal evidence for strict invariance within each sample, this is not sufficient to imply that there is measurement invariance across the pooled set of samples. This may occur, for example, if measurement invariance holds within a given range of age (e.g., invariance holds within each of three studies, one that assessed subjects from ages five to 10, one from ages 10 to 15, and one from ages 15 to 20) yet invariance does *not* hold across the full pooled range of ages (e.g., from ages five to 20). Thus, due to any number of reasons, children might express depression in a functionally different way across studies even though they responded to precisely the same set of items. However, this situation becomes even further complicated when subjects in one study respond to a *different* set of items than subjects in another study. This brings us to measurement comparability.

The topic of measurement comparability is less well studied relative to measurement invariance, yet it is no less important. The reason that it has received much less attention is that this issue does not often arise within a single study. In typical applications a single sample of subjects will all respond to a given set of items, and questions of invariance arise with respect to how the shared set of items might operate differently across group or over time. But when applying IDA to an aggregated sample comprised of two or more independent samples, it is common to encounter the use of partially or wholly different scales to assess a shared underlying construct. In this situation, the core issues that constitute classic measurement invariance simply do not apply. In the classic invariance scenario, the hypothesis being tested is that a given set of items relate to the underlying construct in equivalent ways across groups and over time. As such, it is illogical to raise the question of whether item #1 is related to the underlying construct in the same way as is item #2 across two separate studies; in many cases this is analogous to comparing apples with oranges and our classic methods of assessing invariance are not applicable.

Interestingly, much more is known about this issue in the field of education using Item Response Theory (IRT), particularly as applied to standardized testing of academic skills. For example, it is often theorized that there is some underlying individual-specific latent math ability, but this is assessed with fundamentally different items across grade (e.g., first by addition, then fractions, then algebra, then calculus, etc.). A vast literature exists that focuses

on methods in IRT that allow for things like test equating, scaling and linking to deal with these complex issues in practice (see, e.g., Kolen & Brennan, 2004, and Thissen & Wainer, 2001). These very same issues arise in many (if not most) IDA applications, yet several factors arise that make dealing with this particularly challenging. Issues such as small sample sizes, small numbers of shared items, and multidimensional factor structures combine to limit the use of standard IRT equating and linking procedures (Curran et al., 2007; 2008). Given these current limitations, we must think extremely carefully about issues related to both measurement invariance and measurement comparability when combining multiple data sets into one. We cannot simply compute the standardized mean of 10-items in Study A and compute the standardized mean of 15-items in Study B and assume that these are equivalent measures of the same underlying construct within the pooled sample. We must instead consider all theoretical and empirical evidence that can strengthen our confidence in whether we are assessing the same construct within each individual sample as well as within the pooled sample in a psychometrically equivalent way. Future research in this area is clearly needed.

General Analytic Strategies

We have reviewed five potential sources of between-study heterogeneity that can arise in applying IDA to pooled data: sampling, geography, history, other design characteristics, and measurement. There are of course many other sources that we have not addressed here. However, in our own research, we have sought to address each of these sources of between-study heterogeneity and, in doing so, come to see them as among the most likely to be encountered in IDA applications. To reiterate an important earlier point, it is not paramount that all of the contributing samples be precisely comparable on all possible dimensions to allow for a valid analysis of the aggregated data. Indeed, this was certainly not the case with the Cross Study in which our three studies differed along all five dimensions. As we will see, there are several analytic strategies that allow us to directly test and incorporate potential between-study differences into the analysis of pooled data; this will be done in a way that is consistent with Fisher's model-based procedures that we described earlier. Further, better understanding potential sources of between-study heterogeneity may offer unique insights into each sample individually while we endeavor to generalize our findings in the aggregated sample as a whole. Strategies for addressing these sources of heterogeneity in IDA are newly emerging, and we next summarize several general approaches to IDA that address one or more sources of between-study heterogeneity.

Evaluating Heterogeneity Due to Study-Specific Characteristics

We can think of our available set of independent samples that are to be combined within the IDA in two distinct ways. The first is to conceptualize the collection of data sets as randomly drawn realizations from a homogeneous population of data sets; we refer to this approach as random-effects IDA. The second is to treat the collection of data as fixed as known; we refer to this approach as fixed-effects IDA. There are advantages and disadvantages of each approach, the relative utility of which ultimately depends on the specifics of the IDA application at hand. We briefly explore each in turn.

Random-effects IDA—In any given single study, we typically assume that the sample of observations are randomly selected from some larger (if not infinitely large) population. Random sampling is one of the cornerstones of inferential statistics that permits probabilistic inferences to be drawn from samples to populations. We can extend this concept within the IDA framework where we consider our multiple samples to *themselves* be randomly drawn from some larger population. In this way we literally have a random sample of random samples. This concept is similar to the standard multilevel model in which a sample of independent sampling units (ISUs; e.g., schools) are randomly selected, and then individual observations

are randomly selected within ISU (e.g., students within schools). Here, independent studies are randomly sampled from a population of studies, and individual subjects are sampled within each study. This two-stage sampling introduces two potential sources of variability into our observed data: variability due to the sampling of studies, and variability due to the sampling of individual observations within study (Raudenbush & Bryk, 2002). As such, we can potentially approach IDA from a multilevel perspective.

There are two issues that we must consider here. First, we must somehow establish, even if in theory only, that the multiple data sets can be meaningfully considered as representing random draws from a homogeneous population of data sets. If the independent samples are deemed uniquely and distinctly different from one another in theory or design, then these should likely not be treated as random draws from a single population; the fixed-effects IDA approach should be considered instead. Second, from a more practical perspective, there must be a sufficient number of independent samples to allow for the reliable estimation of the random variability both within and between the samples. There are no infallible rules that dictate how many independent samples are sufficient to allow for proper estimation of the random effects, but in the general multilevel framework often 20 to 30 are viewed as a minimum (e.g., Kreft & de Leeuw, 1998). Although there are certainly potential applications of IDA that would have access to this many independent data sets, most applications within the social sciences would typically be based on substantially fewer samples than this (e.g., in Cross Study we combined just three data sets, yet this still provided nearly 2000 individual observations followed over nearly a 40-year period). As before, if an insufficient number of studies are available to allow for the estimation of a random-effects IDA, we would instead need to consider a fixed-effects approach.

There are, however, a variety of significant advantages to conducting a random-effects IDA if the necessary data are available. Most importantly, if the multiple samples can be treated as randomly drawn from a homogeneous population, we can then consider incorporating study-level predictors to model between-study variability on the outcome measure of interest. This is precisely analogous to having multiple students nested within one of multiple schools and being able to disaggregate student-level effects, school-level effects and the student by school cross-level interactions (e.g., Bauer & Curran, 2005). In IDA we have multiple subjects nested within one of multiple samples and we are able to disaggregate subject-level effects, sample-level effects, and (quite importantly) subject by sample cross-level interactions. Examples of sample-specific measures might include the type of sampling mechanism that was used, the geographical location of the study, and whether data was collected via personal interview or over the internet. A multilevel model could be estimated that simultaneously evaluates the main effects of the within-sample predictors on the outcome, the main effects of the between-sample predictors on the outcome, and the interaction between the within-sample and between-sample predictors on the outcome. Not only is sample-to-sample heterogeneity directly incorporated into the overall model, but this becomes an important empirical and theoretical question in its own right.

As an important final point, although an anonymous reviewer of our paper noted that they found our above arguments compelling, they also felt that the conceptualization of a truly random-effects IDA was "something of a stretch". Their primary concern was that because individual studies are so complexly determined on so many different dimensions that it may be inherently impossible to consider these as truly random realizations from a population of potential studies. We agree. However, we also agree with the reviewer's subsequent suggestion that, under the right circumstances, a random-effects IDA allows for the direct estimation of the between-study variability that could be quite useful if for nothing other than a description of the relative heterogeneity among a set of studies. That is, although we might ultimately not want to make direct inferences back to some hypothetical population of studies, the random-

effects IDA nonetheless provides intriguing insights into the similarities or dissimilarities among a set of studies that would otherwise not be possible. That said, most applications of IDA in the social sciences may simply lack the necessary number of individual data sets to support the estimation of the random-effects model and we thus must turn to the fixed-effects counterpart.

Fixed-Effects IDA—Within random-effects IDA, we treat each data set as an independent random draw from a homogeneous population of data sets. Within fixed-effects IDA, we instead treat study membership as a fixed and known characteristic of each individual observation nested within that study. To accomplish this, we simply incorporate one of several available coding schemes (e.g., dummy codes, effect codes, weighted effect codes) to denote study membership as a fixed characteristic of each individual observation. This is precisely the same strategy we might use to incorporate gender or ethnicity as a fixed characteristic of a given individual; however, here the fixed characteristic of the individual is the study to which they belong. For example, for any given observation, our design matrix might denote a particular individual as male, African American, and belonging to Study 4. These dummy- or effect-coded variables are then entered as predictors in our fitted models in a way consistent with Fisher's model-based inferential methods described earlier. A key advantage of this strategy is that we can also estimate multiplicative interactions between individual characteristics (e.g., gender, ethnicity) and study group membership. This in turn allows for the potential of differential impact of individual characteristics on the outcome across the set of studies. We have used precisely this strategy extensively on Cross Study (e.g. Hussong, Cai et al., 2008; Hussong, Flora et al., 2008; Hussong, Wirth et al., 2007).

There are several critical distinctions between the random-effects and fixed-effects approaches. Most importantly, in the random-effects framework we treat the set of independent samples as random draws from a population and can thus (in principle) make inferences back to an infinite population of samples. In contrast, in the fixed-effects framework we treat the set of independent samples as fixed and known and are thus only able to make inferences back to the specific samples under study. In our view, this is likely a more realistic goal in many IDA applications in psychology. Further, within the random-effects framework we can explicitly disaggregate within-sample effects, between-sample effects, and cross-level interactions. In contrast, in the fixed-effects approach we treat each sample as fixed and known such that the inclusion of the set of study membership variables removes all between-sample sources of variability from the model (although there are some situations in which more complex contrast coding schemes could be used to partially disaggregate these effects; see, e.g., Maxwell & Delaney, 2004).

Thus, whereas in the random-effects model we can estimate the effects of one or more sample-specific measures (e.g., sampling mechanism or geographical region), in the fixed-effects it is typically not possible to include sample-specific measures once the dummy- or effect-coded variables have been entered into the model. This last issue can be simultaneously viewed as a limitation and an advantage of the fixed-effects approach. The obvious limitation is that finer distinctions cannot be made among specific characteristics that define each unique sample because all between-sample variability has been removed from the model. But the associated advantage is that, because the entry of the effect-codes eliminates all between-sample sources of variability, any between-sample differences are controlled even if specific measures regarding these differences are not available. Given the plethora of potential sources of between-sample heterogeneity that exist, controlling for all of these simultaneously can be both beneficial and efficient.

Evaluating Heterogeneity Due to Historical Time

Our discussion thus far has focused on fixed- and random-effects models for evaluating between-study heterogeneity associated with study-specific design characteristics (e.g., differences due to sampling mechanism, geographical location, etc.). These techniques can equivalently be applied to pooled samples that consist of either cross-sectional or longitudinal data. However, several intriguing opportunities arise when considering the role of time when pooling longitudinal data, particularly when there is an interest in disaggregating within-person from between-person differences over time (e.g., Baltes, Reese & Nesselroade, 1988). Indeed, we can use a similar fixed-effects strategy for incorporating information about historical time as we did earlier for study-specific heterogeneity. To accomplish this we draw on methods currently available for testing and combining multiple cohorts within a single study (e.g., Miyazaki & Raudenbush, 2000). However, here we must deal with the added complexity that the multiple cohorts are drawn from multiple studies, and this in turn introduces the possibility of a cohort-by-study interaction.

To begin, consider a single longitudinal study that consists of three repeated measures on a sample of children who were 11 to 15 years of age at first assessment (e.g., Chassin, Rogosch & Barrera, 1991). In this situation a cohort-sequential design would typically be used so that developmental trajectories could be estimated between 11 and 17 years of age. This allows for the estimation of trajectories over a seven-year age span despite the fact that any given child was assessed only three times (e.g., Mehta & West, 2000). In this situation there are five cohorts (ages 11 to 15 at first assessment), three periods (the first, second and third assessments) and seven ages (11 to 17 years). Miyazaki and Raudenbush (2000) provide an excellent general discussion of this type of design and propose an analytic method for testing age, cohort, and age by cohort interactions.

Within the IDA framework we can extend these existing methods in two ways. First, we can adopt the approach of Miyazaki and Raudenbush (2000) in which each specific cohort is dummy-coded and entered as a predictor into the model. A series of nested models are then estimated to evaluate the extent to which a single trajectory underlies the set of repeated observations, or if cohort-specific trajectories are needed. The straightforward expansion we can introduce is the estimation of the interactions between cohort and study membership. This provides a direct test of the extent to which cohort may be differentially related to the outcome within each specific study. These differences could each be formally tested and retained to statistically control for cohort by study interactions prior to evaluating key theoretical predictors of interest.

A second extension we can consider capitalizes on the potential for large numbers of distinct cohorts that may be available as a result of pooling multiple longitudinal data sets. Because Miyazaki and Raudenbush (2000) considered a single data set with seven distinct cohorts, they logically chose to treat these as discrete groups and thus used a nominal coding scheme of dummy variables and orthogonal contrasts. However, if a sufficient number of cohorts are available, these could instead be included as a quantitative variable that would then allow for additional ways to evaluate age-related and cohort-related change.⁹ For example, within the final pooled sample on Cross Study there were 29 unique birth cohorts spanning from 1964 to 1992. Because of the large number of unique birth years it is possible to include this as a quantitatively-scaled between-subjects predictor in subsequent models. This in turn introduces the exciting possibility of simultaneously incorporating two measures of time in a single model: between-subject functions of historical time (as measured by birth year) and within-subject functions of developmental time (as measured by chronological age). Given sufficient data,

⁹We thank Dan Bauer for the original development of this idea.

these two dimensions of time could be allowed to interact with one another as well as with study group membership allowing for a variety of important hypotheses to be evaluated that were not previously possible.

Evaluating Heterogeneity Due to Measurement

There are two general situations in which issues of measurement arise in most IDA applications. The first situation is when studying the measurement properties of a set of items or tests such that measurement itself is the primary question of interest. For example, multiple independent studies might be pooled to allow for the examination of the factor structure underlying some set of items that assess constructs such as personality traits, intelligence, or psychopathology. The second situation is when measurement is used more as a means to some other end such that the goal is to produce a reliable and valid scale score that could then be used in other types of analyses. For example, a confirmatory factor analysis model might be fitted to a set of items to create a scale score that would then be used as a criterion measure in a separate multilevel model (e.g., Curran et al., 2008). Of course these two applications of measurement are highly related in that one must first examine the measurement properties prior to computing an associated scale score. Despite the obvious potential applications of the first situation, here we focus on the second given the likely broader use of this approach in many IDA applications in the social sciences.

One of the reasons that measurement is so critically important within the IDA framework relates to our goal of making valid inferences back to theory based upon the empirical characteristics of the pooled data set. We must establish that we are measuring the same theoretical construct in the same manner for all individuals across all data sets. How to best accomplish this in large part depends upon the specific characteristics of the data sets to be pooled. For example, say that the theoretical construct of interest is childhood depression. In an ideal situation all studies used precisely the same items, response scales (e.g., 1 to 10 Likert scale), and time frame (e.g., past 30 days) to measure depression. A less ideal situation is one in which all studies used precisely the same items but some studies used different response scales (1 to 5 Likert scale versus 1 to 10 Likert scale) and some studies used different time frames (e.g., past 30 days versus past 60 days). Finally, probably the most typical situation is one in which each study used some unique combination of items, response scales, and time frames. Different analytic strategies are more or less well suited for handling these different combinations of item characteristics across studies (see Bauer & Hussong, this issue, and McArdle, Grimm, Hamagami, Bowles & Meredith, this issue, for further discussions of these issues).

To establish the optimal measurement strategy, the first step is to identify the set of available items that are believed to assess the underlying construct pooling across the set of contributing studies. It is extremely helpful if at least some portion of this pooled item set contain items that are shared across all contributing studies. These shared items form a set of "anchor" items that can be of great use in later stages of measurement development (Kolen & Brennan, 2004). For example, there may be an item that is worded in precisely the same way across all studies (e.g., "My child cries easily"); this could unambiguously serve as a potential anchor item. However, there may be items that are not worded in precisely the same manner but might nonetheless be psychometrically operating in a similar fashion. For example, several studies might have asked the item in slightly different ways (e.g., "My child cries easily", "My child cries often", and "My child cries more frequently than other children"). Although the wording may appear to be closely related across the three items, these differences may or may not be sufficient to treat them as separate indicators. Finally, there will be items that will be distinctly unique to just one or a small number of contributing studies and were clearly not assessed in other studies. For example, only a single contributing study may have assessed the item "My child feels sad

even in the presence of others", yet from a theoretical standpoint this item is a valid indicator of childhood depression and should be retained if possible.

Selection of the proper statistical measurement model depends in large part on the form of the response scales that were used for each individual item. The standard linear confirmatory factor analysis (CFA) model is well suited for items that are (at least approximately) interval scaled, although there is some controversy regarding what constitutes "approximately" (e.g., Muthén & Kaplan, 1985; 1992). If individual items do not sufficiently approximate an interval scale, two or more individual items can be combined to create item parcels that better correspond to the assumed continuous distribution (e.g., Little, Cunningham, Shahar, & Widaman, 2002; MacCallum, Widaman, Zhang, & Hong, 1999). Alternatively, if the items are discretely scaled (e.g., binary or a small number of ordinal responses) and the creation of item parcels is not a viable option, then the assumption of linearity underlying the standard CFA model is not well met and a non-linear measurement model is necessary (Flora & Curran, 2004). There are two key options in this situation: non-linear factor analysis (NLFA; e.g., Skrondal & Rabe-Hesketh, 2004) and item response theory (IRT; e.g., Thissen & Wainer, 2001). There is a vast literature dedicated to both the NLFA and IRT models; a thorough review of more general issues in non-linear measurement models is given in Wirth and Edwards (2007), and we have discussed these issues as they relate to applications within IDA in Bauer and Hussong (this issue) and Curran et al. (2007; 2008).

Once the potential item pool has been identified and the optimal statistical model has been chosen, then actual model fitting can begin. As with many other points in our discussion, the specific steps to be taken will depend on the goals and unique characteristics of the given IDA. However, in most applications there are four general steps in the measurement portion of the analysis (see Curran et al., 2008, for a more detailed discussion of these steps). First, some assessment must be made of the dimensionality underlying the set of items. The standard IRT model assumes unidimensionality, although recently developed techniques allow for the estimation of multidimensional IRT models under certain circumstances (e.g., Fox & Glas, 2001; Rabe-Hesketh, Skrondal, & Pickles, 2004). In contrast, the CFA and NLFA models typically allow for either unidimensionality or multidimensionality. Second, measurement models are first fitted within each study separately and then across all studies simultaneously to establish an initial understanding of the psychometric properties of the scales. These properties take the form of factor loadings and item intercepts in the factor analysis models, and of discrimination and severity parameters in the IRT model. Third, some type of assessment is needed of measurement invariance either across study, demographic group, or over time. Multiple group analysis provides such tests in the factor analysis framework, and this is accomplished using differential item functioning (DIF) in the IRT model. Bauer and Hussong (this issue) and McArdle et al. (this issue) describe other options for evaluating invariance in IDA. Finally, once a comprehensive measurement model has been established, then scale scores are calculated that are jointly based on the observed pattern of responses to the items and the parameter estimates from the final measurement model. These scale scores can be calculated using one of several available factor score estimates in the factor model (e.g., Grice, 2001) and using MAP or EAP scoring in the IRT model (e.g., Thissen & Wainer, 2001). Regardless of approach, the motivating goal is to create a person-specific scale score that incorporates information about study group membership and, potentially, demographic group membership; these scale scores can then be used in subsequent statistical analyses.

Conclusion

We view IDA as a product of our times. These methods respond to an increased demand for collaborative efforts that make efficient use of limited resources in the pursuit of a cumulative science. Moreover, these methods call for the integration of cutting edge techniques concerning

longitudinal modeling and measurement evaluation that in tandem have the power to address many of the vexing problems of IDA surrounding study integration and study comparison. We have attempted to outline the advantages of IDA and the types of applications where IDA may be particularly useful. As we noted at the outset, we clearly recognize that IDA will not be possible in all applications. In our own work, we have succeeded in conducting IDA across our three contributing studies in some instances (e.g., Curran et al., 2008), needed to drop to using only two studies in others (e.g., Hussong, Cai et al., 2008; Hussong, Flora et al., 2008), and relied on parallel, single study analysis on another occasion (Hussong, Bauer, Huang, Chassin, Sher & Zucker, 2008). Thus, in our experience, the feasibility of IDA rests on the characteristics of the contributing studies as they bear on specific questions of theoretical interest. Understanding the boundaries of IDA within those applications where raw data are available for studies of homogenous populations is an important area of further development for these methods.

The focus of this paper on analysis of existing data is merely a starting point for IDA. Collaborative models for IDA using primary data collection efforts in psychology come from existing multi-site designs (e.g., Hofer & Picinnin, this issue). However, we believe that IDA has much more to offer in study design that is yet unrealized. Greater attention to assessing study characteristics, and not simply participant characteristics, would greatly aid in efforts to understand between-study sources of heterogeneity. Measurement approaches that are both tailored to the target sample of a given study but also include linkage items across studies create important opportunities for IDA. Discipline and field-level efforts to coordinate research may offer enormous potential for IDA applications, such as the suggested core items for assessing alcohol use put forth by the National Institute on Alcohol Abuse and Alcoholism (2003). Such efforts to provide a core but not necessarily exclusive set of items for assessment permit greater study comparison while retaining the need for study innovation. As such, overlapping but non-redundant item sets are often ideal for IDA. Regardless of specifics, efforts to reach beyond single study design planning are clearly needed to facilitate the future integration of findings pooling across multiple studies, and IDA can be an important part of those efforts.

Acknowledgments

This work was partially supported by grant DA013148 awarded to the first author and grant DA15398 awarded to the second author. We are deeply indebted to Dan Bauer, Dave Flora, and all of the past and present members of the Cross Study group whose creativity, commitment, and exceedingly high frustration tolerance have made this project possible.

References

- Achenbach T, Edelbrock C. The classification of child psychopathology: A review and analysis of empirical efforts. *Psychological Bulletin* 1981;85:1275–1301. [PubMed: 366649]
- American Association for the Advancement of Science. AAAS Report XXXIII: Research and Development FY2009. Intersociety Working Group, 2008. AAAS Publication Number 08-1A. 2008.
- American Heritage® Dictionary of the English Language. Fourth Edition Retrieved January 29, 2009, from Dictionary.com website <http://dictionary.reference.com/browse/integrate>.
- American Psychological Association. Ethical principles of psychologists and code of conduct. *American Psychologist* 2003;57:1060–1073.
- Baltes, PB.; Reese, HW.; Nesselroade, JR. *Life-Span Developmental Psychology: Introduction to Research Methods*. Hillsdale, NJ: Lawrence Erlbaum Publishers; 1988.
- Bauer DJ, Curran PJ. Probing interactions in fixed and multilevel regression: Inferential and graphical techniques. *Multivariate Behavioral Research* 2005;40:373–400.
- Bauer DJ, Hussong AM. Psychometric approaches for developing commensurate measures across independent studies: Traditional and new models. *Psychological Methods*, [this issue]. 2009

- Berlin JA, Santanna J, Schmid CH, Szczech LA, Feldman HI. Individual patient-versus group-level data meta-regressions for the investigation of treatment effect modifiers: Ecological bias rears its ugly head. *Statistics in Medicine* 2002;21:371–387. [PubMed: 11813224]
- Bollen KA. Latent variables in psychology and the social sciences. *Annual Review of Psychology* 2002;53:605–634.
- Bureau of Labor Statistics, U.S. Department of Labor. Produced and distributed by the Center for Human Resource Research. Columbus, OH: The Ohio State University; 2002. National Longitudinal Survey of Youth 1979 cohort, 1979–2002 (rounds 1–20) [computer file].
- Chassin L, Rogosch F, Barrera M. Substance use and symptomatology among adolescent children of alcoholics. *Journal of Abnormal Psychology* 1991;100(4):449–463. [PubMed: 1757658]
- Cochran, WG. Sampling Techniques. Third Edition. NY: John Wiley & Sons; 1977.
- Cohen J. A power primer. *Psychological Bulletin* 1992;112:155–159. [PubMed: 19565683]
- Cooper, HM. Research Synthesis and Meta-Analysis: A Step-by-Step Approach. 4th Ed. Thousand Oaks, CA: Sage; (in press).
- Cooper HM, Patall EA. The relative benefits of meta-analysis using individual participant data or aggregated data. *Psychological Methods*, [this issue]. 2009
- Cudeck, R.; MacCallum, RC. Factor Analysis at 100: Historical Developments and Future Directions. Psychology Press; 2007.
- Curran, PJ.; Edwards, MC.; Wirth, RJ.; Hussong, AM.; Chassin, L. The incorporation of categorical measurement models in the analysis of individual growth. In: Little, T.; Bovaird, J.; Card, N., editors. Modeling ecological and contextual effects in longitudinal studies of human development. Mahwah, NJ: LEA; 2007. p. 89-120.
- Curran PJ, Hussong AM, Cai L, Huang W, Chassin L, Sher KJ, Zucker RA. Pooling data from multiple prospective studies: The role of item response theory in integrative analysis. *Developmental Psychology* 2008;44:365–380. [PubMed: 18331129]
- Derogatis, LR.; Spencer, P. Brief Symptom Inventory (BSI). Baltimore: Clinical Psychometric Research; 1982.
- DeRubeis RJ, Gelfand LA, Tang TZ, Simons AD. Medications versus cognitive behavior therapy for severely depressed outpatients: Mega-analysis of four randomized comparisons. *American Journal of Psychiatry* 1999;156:1007–1013. [PubMed: 10401443]
- Doherty M. Probability versus non-probability sampling in sample surveys. *The New Zealand Statistics Review*, March 1994;21–28.
- Fisher RA. On the mathematical foundations of theoretical statistics. *Philosophical Transactions of the Royal Society of London Series A* 1922;222:309–368.
- Flora DB, Curran PJ. An Empirical Evaluation of Alternative Methods of Estimation for Confirmatory Factor Analysis With Ordinal Data. *Psychological Methods* 2004;9(4):466–491. [PubMed: 15598100]
- Flora DB, Curran PJ, Hussong AM, Edwards MC. Incorporating measurement nonequivalence in a cross-study latent growth curve analysis. *Structural Equation Modeling*. 2008;15:676–704.
- Fox J-P, Glas CAW. Bayesian estimation of a multilevel IRT model using Gibbs sampling. *Psychometrika* 2001;66(2):271–288.
- Gans HJ. Sociological amnesia: The noncumulation of normal social science. *Sociological Forum* 1992;7:701–710.
- Glass GV. Primary, secondary, and meta-analysis. *Educational Researcher* 1976;5:3–8.
- Grice JW. Computing and evaluating factor scores. *Psychological Methods* 2001;6:430–450. [PubMed: 11778682]
- Guo S, Hussey DL. Non-probability sampling in social work research: Dilemmas, consequences, strategies. *Journal of Social Service Research* 2004;30:1–18.
- Hannah MG. Sampling and the politics of representation in US Census 2000. *Environment and Planning D: Society and Space* 2001;19:515–534.
- Harford TC. The effects of order of questions on reported alcohol consumption. *Addiction* 1994;89:421–424. [PubMed: 8025495]

- Harris, KM.; Halpern, CT.; Entzel, P.; Tabor, J.; Bearman, PS.; Udry, JR. The National Longitudinal Study of Adolescent Health: Research Design [WWW document]. 2008. Retrieved February 12, 2009 <http://www.cpc.unc.edu/projects/addhealth/design>.
- Higgins JPT, Whitehead A, Turner R, Omar RZ, Thompson SG. Meta-analysis of continuous outcome data from individual patients. *Statistics in Medicine* 2001;20:2219–2241. [PubMed: 11468761]
- Hofer SM, Picinnin AM. Integrative data analysis through coordination of measurement and analysis protocol across independent longitudinal studies. *Psychological Methods*, [this issue]. 2009
- Horn JL, McArdle J. A practical guide to measurement invariance in research on aging. *Experimental Aging Research* 1992;18:117–144. [PubMed: 1459160]
- Hunter DJ, Spiegelman D, Adami HO, Beeson L, van den Brandt PA, Folsom AR, Fraser GE, Goldbohm A, Graham S, Howe GR, Kushi LH, Marshall JR, McDermott A, Miller AB, Speizer FE, Wolk A, Yaun SS, Willett W. Cohort studies of fat intake and the risk of breast cancer: A pooled analysis. *The New England Journal of Medicine* 1996;334:356–361. [PubMed: 8538706]
- Hunter JE, Schmidt FL. Cumulative research knowledge and social policy formulation: The critical role of meta-analysis. *Psychology, Public Policy, and Law* 1996;2:324–347.
- Hussong AM, Bauer DJ, Huang W, Chassin L, Sher KJ, Zucker RA. Characterizing the Life Stressors of Children of Alcoholic Parents. *Journal of Family Psychology* 2008;22:819–832. [PubMed: 19102603]
- Hussong AM, Cai L, Curran PJ, Flora DB, Chassin LA, Zucker RA. Disaggregating the Distal, Proximal, and Time-Varying Effects of Parent Alcoholism on Children's Internalizing Symptoms. *Journal of Abnormal Child Psychology* 2008;36:335–346. [PubMed: 17891557]
- Hussong AM, Flora DB, Curran PJ, Chassin LA, Zucker RA. Defining Risk Heterogeneity for Internalizing Symptoms among Children of Alcoholic Parents: A Prospective Cross-Study Analysis. *Development and Psychopathology* 2008;20:165–193. [PubMed: 18211733]
- Hussong AM, Wirth RJ, Edwards MC, Curran PJ, Chassin LA, Zucker RA. Externalizing symptoms among children of alcoholic parents: Entry points for an antisocial pathway to alcoholism. *Journal of Abnormal Psychology* 2007;116:529–542. [PubMed: 17696709]
- Johnston, LD.; Bachman, JG.; O'Malley, PM. Monitoring the Future: Questionnaire responses from the nation's high school seniors 2005. Ann Arbor, MI: Institute for Social Research; 2006.
- Kaplow JB, Curran PJ, Dodge K. the Conduct Problems Prevention Research Group. Child, parent, and peer predictors of early-onset substance use: A multi-site longitudinal study. *Journal of Abnormal Child Psychology* 2002;30:199–216. [PubMed: 12041707]
- Kiecolt, KJ.; Nathan, L. Secondary analysis of survey data. Thousand Oaks, CA: Sage Publications; 1985.
- Kolen, MJ.; Brennan, RL. Test Equating, Scaling, and Linking: Methods and Practices. 2nd Edition. New York: Springer; 2004.
- Kreft, I.; de Leeuw, J. Introducing multilevel modeling. London: Sage; 1998.
- Lambert PC, Sutton AJ, Abrams KR, Jones DR. A comparison of summary patient-level covariates in mega-regression with individual patient data meta-analysis. *Journal of Clinical Epidemiology* 2002;55:86–94. [PubMed: 11781126]
- Lenhard J. Models and statistical inference: The controversy between Fisher and Neyman-Pearson. *British Journal of Philosophy of Science* 2006;57:69–91.
- Lerer B, Segman RH, Fangerau H, Daly AK, Basile VS, Cavallaro R, Aschauer HN, McCreddie RG, Ohlraun S, Ferrier N, Masellis M, Verga M, Scharfetter J, Rietschel M, Lovlie R, Levy UH, Meltzer HY, Kennedy JL, Steen VM, Macciardi F. Pharmacogenetics of tardive dyskinesia: Combined analysis of 780 patients supports association with dopamine D3 receptor Gene Ser9Gly polymorphism. *Neuropsychopharmacology* 2002;27:105–119. [PubMed: 12062911]
- Little RJA. To model or not to model? Competing modes of inference for finite population sampling. *Journal of the American Statistical Association* 2004;88:546–556.
- Little TD, Cunningham WA, Shahar G, Widaman KF. To parcel or not to parcel: Exploring the question, weighing the merits. *Structural Equation Modeling* 2002;9:151–173.
- Lorenz FO, Simons RL, Conger RD, Elder GH, Johnson C, Chao W. Married and recently divorced mothers' stressful events and distress: Tracing change across time. *Journal of Marriage and the Family* 1997;59:219–232.

- MacCallum RC, Widaman KF, Zhang S, Hong S. Sample size in factor analysis. *Psychological Methods* 1999;4:84–99.
- Maxwell SE. The persistence of underpowered studies in psychological research: Causes, consequences, and remedies. *Psychological Methods* 2004;9:147–163. [PubMed: 15137886]
- Maxwell, SE.; Delaney, HD. *Designing experiments and analyzing data: A model comparison perspective*. 2nd Edition. Mahwah, NJ: Lawrence Erlbaum Associations; 2004.
- McArdle JJ, Grimm K, Hamagami F, Bowles R, Meredith W. Modeling life-span growth curves of cognition using longitudinal data with multiple samples and changing scales of measurement. *Psychological Methods*, [this issue]. 2009
- McArdle JJ, Hamagami F, Meredith W, Bradway KP. Modeling the dynamic hypotheses of Gf-Gc theory using longitudinal life-span data. *Learning and Individual Differences* 2000;12:53–79.
- McArdle JJ, Prescott CA, Hamagami F, Horn JL. A contemporary method for developmental-genetic analyses of age changes in intellectual abilities. *Developmental Neuropsychology* 1998;14:69–114.
- Meehl PE. Theoretical risks and tabular asterisks: Sir Karl, Sir Ronald, and the slow progress of soft psychology. *Journal of Consulting and Clinical Psychology* 1978;46:806–834.
- Mehta PD, West SG. Putting the individual back into individual growth curves. *Psychological Methods* 2000;5:23–43. [PubMed: 10937321]
- Meredith W. Measurement invariance, factor analysis and factorial invariance. *Psychometrika* 1993;58:525–543.
- Millsap RE. Measurement invariance, predictive invariance, and the duality paradox. *Multivariate Behavioral Research* 1995;30:577–605.
- Millsap RE. Invariance in measurement and prediction: Their relationship in the single factor case. *Psychological Methods* 1997;2:248–260.
- Miyazaki Y, Raudenbush SW. A test for linkage of multiple cohorts from an accelerated longitudinal design. *Psychological Methods* 2000;5:44–63. [PubMed: 10937322]
- Muthén BO, Kaplan D. A comparison of some methodologies for the factor-analysis of nonnormal Likert variables. *British Journal of Mathematical and Statistical Psychology* 1985;38:171–180.
- Muthén BO, Kaplan D. A comparison of some methodologies for the factor analysis of non-normal Likert variables: A note on the size of the model. *British Journal of Mathematical and Statistical Psychology* 1992;45:19–30.
- Muthén BO, Satorra A. Complex sample data in structural equation modeling. *Sociological Methodology* 1995;25:267–316.
- National Institute on Alcoholism and Abuse. Task Force on Recommended Alcohol Questions - National Council on Alcohol Abuse and Alcoholism Recommended Sets of Alcohol Consumption Questions - October 15–16, 2003. 2003. Retrieved February 9, 2009. <http://www.niaaa.nih.gov/Resources/ResearchResources/TaskForce.htm>.
- Nesselroade JR, Gerstorf D, Hardy SA, Ram N. Idiographic filters for psychological constructs. *Measurement* 2007;5:217–235.
- Neyman J. On the two different aspects of the representative method: The method of stratified sampling and the method of purposive selection. *Journal of the Royal Statistical Society* 1934;109:558–606.
- Palmore E. When can age, period, and cohort be separated? *Social Forces* 1978;57:282–295.
- Pentz MA, Chou CP. Measurement invariance in longitudinal clinical research assuming change from development and intervention. *Journal of Consulting and Clinical Psychology* 1994;62:450–462. [PubMed: 8063972]
- Pfeffermann D. The role of sampling weights when modeling survey data. *International Statistical Review* 1993;61:317–337.
- Rabe-Hesketh S, Skrondal A, Pickles A. Generalized multilevel structural equation modelling. *Psychometrika* 2004;69:167–190.
- Raudenbush, SW.; Bryk, AS. *Hierarchical Linear Models: Applications and Data Analysis Methods*. Second Edition. Thousand Oaks: Sage Publications; 2002.
- Rothstein, HR.; Sutton, AJ.; Borenstein, M. *Publication bias in meta-analysis: Prevention, assessment and adjustments*. West Sussex, England: Wiley & Sons; 2005.

- Rusticus SA, Hubley AM, Zumbo BD. Measurement invariance of the appearance schemas inventory-revised and the body image quality of life inventory across age and gender. *Assessment* 2008;15:60–71. [PubMed: 18258732]
- Shadish, WR.; Cook, TD.; Campbell, DT. Experimental and quasi-experimental designs for generalized causal inference. Boston, MA: Houghton Mifflin Company; 2002.
- Schaie K. A general model for the study of developmental problems. *Psychological Bulletin* 1965;64:92–107. [PubMed: 14320080]
- Schaie, K. Developmental designs revisited. In: Cohen, SH.; Reese, HW., editors. Life-span developmental psychology: Theoretical issues revisited. Hillsdale, NJ: Erlbaum; 1994.
- Schmidt, FL. Meta-Analysis: Implications for cumulative knowledge in the behavioral and social sciences. Paper presented at the Annual Convention of the American Psychological Association; Toronto, Canada. 1984 Aug.
- Schmidt FL. Statistical significance testing and cumulative knowledge in psychology: Implications for training of researchers. *Psychological Methods* 1996;1:115–129.
- Sher KJ, Walitzer KS, Wood PK, Brent EE. Characteristics of children of alcoholics: Putative risk factors, substance use and abuse, and psychopathology. *Journal of Abnormal Psychology* 1991;100(4):427–448. [PubMed: 1757657]
- Skrondal, A.; Rabe-Hesketh, S. Generalized latent variable modeling: multilevel, longitudinal and structural equation models. Boca Raton: Chapman & Hall/CRC; 2004.
- Stevens SS. On the theory of scales of measurement. *Science* 1946;103:677–680. [PubMed: 17750512]
- Smith ML, Glass GV. Meta-analysis of psychotherapy outcome studies. *American Psychologist* 1977;32:752–760. [PubMed: 921048]
- Stark LJ, Janicke DM, McGrath AM, Mackner LM, Hommel KA, Lovell D. Prevention of osteoporosis: A randomized clinical trial to increase calcium intake in children with juvenile rheumatoid arthritis. *Journal of Pediatric Psychology* 2005;30:377–386. [PubMed: 15944165]
- Sterba SK. Alternative inferential frameworks for non-probability sampling in psychology and probability sampling in allied fields: Polarization to reconciliation. Manuscript submitted for publication. 2009
- Thissen, D.; Wainer, H. Test scoring. Mahwah, NJ: Erlbaum; 2001.
- Thomas SL, Heck RH. Analysis of large-scale secondary data in higher education research: Potential perils associated with complex sampling designs. *Research in Higher Education* 2001;42:517–540.
- Thurstone, LL. Multiple-factor analysis. Chicago, IL: University of Chicago Press; 1947.
- van den Brandt PA, Spiegelman D, Yaun SS, Adami HO, Beeson L, Folsom AR, Fraser G, Goldbohm RA, Graham S, Kushi L, Marshall JR, Miller AB, Rohan T, Smith-Warner SA, Speizer FE, Willett WC, Wolk A, Hunter DJ. Pooled analysis of prospective cohort studies on height, weight, and breast cancer risk. *American Journal of Epidemiology* 2000;152:514–527. [PubMed: 10997541]
- Vandenberg RJ, Lance CE. A review and synthesis of the measurement in-variance literature: Suggestions, practices, and recommendations for organizational research. *Organizational Research Methods* 2000;3:4–69.
- Wainer, H. Drawing inferences from self-selected samples. New York: Springer; 1986.
- Wirth RJ, Edwards MC. Item factor analysis: Current approaches and future directions. *Psychological Methods* 2007;12:58–79. [PubMed: 17402812]
- Zucker, RA.; Fitzgerald, HE.; Refior, SK.; Puttler, LI.; Pallas, DM.; Ellis, DA. The clinical and social ecology of childhood for children of alcoholics: Description of a study and implications for a differentiated social policy. In: Fitzgerald, HE.; Lester, BM.; Zuckerman, BS., editors. Children of addiction: Research, health and policy issues. New York: Garland Press; 2000. p. 174-222.

Participant Age

[illegible]

USA=Michigan Longitudinal Study annual assessments (see Zucker et al., 2000); AFDP=Adolescent and Family Development Project; AHP=Alcohol Health to the study indicator indicates wave of assessment (e.g., AFDP3 is the third assessment wave in the AFDP project).

Table 2**Exemplar Study Design Differences in Cross Study**

	MLS	AFDP	AHBP
Design			
Recruitment	Rolling community-based recruitment with COA families identified through father's court-arrest records and community canvassing	A community-based sample with alcoholic parents identified through court records, HMO wellness questionnaires and telephone surveys	Recruited through a screening of 3,156 first-time freshmen at the University of Missouri who reported on paternal alcoholism using the father-SMAST.
Assessment Schedule	Mothers and fathers completed up to four assessments when the children were between ages 2–5, 6–8, 9–11, and 11–15 at 3 year intervals.	Mothers, fathers and one child completed the first three annual waves of data on children age 10–17 and two subsequent follow-up waves at 5-year intervals; age-appropriate siblings were also included as targets in the follow-up waves	Children completed four annual assessments (years 1–4) and two additional post-college follow-ups (at 3 and 4 year intervals, or years 7 and 11).
Variables			
Parent Alcoholism	Life time diagnosis was made by a trained clinician based on DSM-IV criteria with parent self-report at each wave using three instruments: Diagnostic Interview Schedule, the Short Michigan Alcohol Screening Test, and the Drinking and Drug History Questionnaire.	Life-time diagnosis was made by interviews based on DSM-III criteria with parent self-report at the first wave using the computerized version of the DIS. In cases where a biological parent was not directly interviewed, the reporting parent was used as the informant using the FH-RDC.	Life-time diagnosis was made by survey assessment based on DSM-III criteria with target (child) report at baseline using the parent-SMAST and FH-RDC.

Note: MLS=Michigan Longitudinal Study; AFDP= Adolescent and Family Development Project; AHBP= Alcohol and Health Behavior Project; COA=Child of Alcoholic; DSM=Diagnostic and Statistical Manual of Mental Disorders; DIS=Diagnostic Interview Schedule; FH-RDC=Family History Research Diagnostic Criteria; SMAST=Short-form of the Michigan Alcohol Screening Test.