

A Guide to Incorporating Multiple Methods in Randomized Controlled Trials to Assess Intervention Effects

David W. Grissmer

Center for Advanced Study of Teaching and Learning
University of Virginia

Rena F. Subotnik

Center for Psychology in Schools and Education
American Psychological Association

Martin Orland

WestEd



A Guide to Incorporating Multiple Methods in Randomized Controlled Trials to Assess Intervention Effects

David W. Grissmer

Center for Advanced Study of Teaching and Learning
University of Virginia

Rena F. Subotnik

Center for Psychology in Schools and Education
American Psychological Association

Martin Orland

WestEd



A Guide to Incorporating Multiple Methods in Randomized Controlled Trials to Assess Intervention Effects

Available online at: <http://www.apa.org/ed/schools/cpse/activities/mixed-methods.aspx>

Printed copies available from:

Center for Psychology in Schools and Education

Education Directorate

American Psychological Association

750 First Street, NE

Washington, DC 20002-4242

Phone: 202-336-5977 • TDD/TTY: 202-336-6123

E-mail: rsubotnik@apa.org

Suggested bibliographic reference:

Grissmer, D. W., Subotnik, R. F., & Orland, M. (2009). *A guide to incorporating multiple methods in randomized controlled trials to assess intervention effects*. Washington, DC: American Psychological Association.

Copyright © 2009 by the American Psychological Association. This material may be reproduced in whole or in part without fees or permission provided that acknowledgment is given to the American Psychological Association. This material may not be reprinted, translated, or distributed electronically without prior permission in writing from the publisher. For permission, contact APA, Rights and Permissions, 750 First Street, NE, Washington, DC 20002-4242.

APA reports synthesize current psychological knowledge in a given area and may offer recommendations for future action. They do not constitute APA policy or commit APA to the activities described therein.

CONTENTS

Introduction	1
Background on the Use of RCTs and Multiple Methods	3
Box 1: Motivating Policy/Research Questions	5
Box 2: Desirability/Feasibility of an RCT Study	9
Box 3: Employing Multiple Methods in Designing and Implementing RCTs	15
Box 4: Role of Multiple Methods in Providing a Deeper Understanding of Study Findings	25
Box 5: Exploring Implications for Research, Policy, and Next Steps	28
References	33
Appendix: Project Descriptions	41
Project Star	41
New Hope	42
Moving to Opportunity	43

INTRODUCTION

This narrative and the accompanying chart (see insert at the end of this report) take the reader through the lifecycle of a research study on the effects and effectiveness of a social intervention. The narrative and chart are intended for instructional use by new researchers, as well as to assist consumers of education research such as policymakers and school leaders contracting research and evaluation work.

The narrative and chart had their origins in a December 2004 national forum on incorporating multiple social science research methods in conjunction with randomized controlled trials (RCTs) in education. The National Research Council hosted this forum in collaboration with the American Psychological Association, the American Educational Research Association, and the National Science Foundation. Its underlying premise was that the contribution of RCTs to research, policy, and practice could be greatly enhanced when multiple methods are used to help understand the effects of context, population, resource constraints, and generalizability of research findings.

In response to the enthusiasm expressed by forum participants for continued work in this area, the forum organizing committee believed that it could make an important contribution both to educational research and to policy development by constructing a guide describing the conditions and circumstances favoring the use of RCTs, as well as demonstrating the value of other research methodologies and data collections to the overall research

investigation. The specific objectives of the chart and this narrative are to provide the following:

- A rationale for incorporating multiple methods into RCTs
- A guide for designing RCTs with multiple methods
- Examples from the literature that illustrate the various steps in the process
- Examples from the literature that illustrate how additional data collected with multiple methods may be used to address the following questions:
 - What are the causative mechanisms involved in producing a targeted effect in the RCT and how do these mechanisms work to produce the effect?
 - Why does the intervention work better for some participants and not others?
 - Can existing theory predict these results, or do the results suggest the need for new theories?
 - How can this intervention be redesigned to be more effective or efficient?
 - How would a scaled-up program likely change the predicted costs and effect size?
 - Do effects of scaled-up existing programs using nonexperimental data provide results similar to those of experimental data, and if not, why not? Which measurement provides a more accurate prediction?

Each section of the chart (Boxes 1–5; see insert at back) has a related section in this narrative. Each section of the narrative is designed to be relatively independent of the others so that readers may select the section from the chart that they wish to read and study in the guide. You may request hard copies of the chart in poster and foldout form by contacting Rena Subotnik (rsubotnik@apa.org).

Introducing multiple methods into RCTs is becoming a more frequent occurrence. However, a deep vein of literature that illustrates the process by which professionals design and use data for multiple methods does not exist. Typically, it takes almost 8–10 years from the inception of an RCT for a complete set of analyses to appear in journal publications. RCTs with multiple methods take even longer. Finding illustrative examples of RCTs using multiple methods therefore requires looking at RCTs started in the late 1980s and early to mid-1990s. Fortunately, there were a few RCTs from this period that incorporated multiple methods, as well as some publications that illustrated the power of multiple-methods RCTs.

In providing descriptive material associated with the chart, we drew primarily from three such RCTs:

- **Project STAR** was an experiment in class size reduction (kindergarten to third grade) conducted in Tennessee beginning in 1986.
- The **New Hope** intervention, begun in Milwaukee in 1994, aimed to lift the earning wages of individuals living below the poverty level through higher paying, more stable jobs that would improve the life circumstances of those individuals and their children. New Hope provided income supplements, health insurance coverage, and paid day care over a 3-year period.
- **Moving to Opportunity**, begun in 1994, provided a chance for families living in public housing in very poor and risky neighborhoods to relocate to better neighborhoods.

These three RCTs with multiple methods have long-term follow-ups and a rich literature trail, and the data they have used address the questions we cited earlier. We provide brief descriptions of these RCTs in the Appendix; in addition, the References section contains most of the publications spawned by these multiple-methods RCTs.

Each of these RCTs with multiple methods was started between 1986 and 1994. Each of the experiments ran for 3 or 4 years and included long-term follow-up studies after completion. Each also generated productive literature

starting in the year or two after the end of the experiment and continuing through 2008. We chose these three illustrative and multiple-methods RCTs because they represent some of the best examples from the 1980s and 1990s and were among the first to incorporate multiple methods into their design.

New Hope stands out in several respects as the best current example illustrating the utility of multiple-methods RCTs. The multiple-methods data collections were numerous, well designed, and methodologically diverse. It is also important to note that the utilization of this data in the analyses has been well documented and directed toward diverse audiences. Journal articles have addressed key research issues.

Books have been published entirely devoted to illustrating and using the multiple-methods data to address the questions cited above. In addition, a summary book directed primarily toward policymakers makes use of the multiple-methods data to communicate an enriched, in-depth, and indelible understanding of the research results and analyses of policy options. For those desiring an understanding of the utility and inclusion of multiple-methods data in RCTs, reading this literature provides the best and most complete current source.

However, each of these RCTs has flaws in design, implementation, and analysis. Some of these flaws are inevitable parts of doing social experimentation. Others might be attributed to compressed time during planning and implementation or limited theoretical development and budgets. Hindsight is 20/20, yet reflecting on these flaws offers not only the opportunity to illustrate the complexity of designing RCTs and using multiple methods in the real world but also to learn from and improve future multiple-methods RCTs.

Finally, we have likely omitted articles, research, and viewpoints that could have improved this description. We hope that this effort will be seen as a starting point for an expanded discussion of multiple-methods RCTs in the research literature and that a richer and more diverse set of perspectives will emerge in future discussions.

We hope that this effort will be seen as a starting point for an expanded discussion of multiple-methods RCTs in the research literature and that a richer and more diverse set of perspectives will emerge in future discussions.

BACKGROUND ON THE USE OF RCTs AND MULTIPLE METHODS

There has been a long and messy debate in the social science community, and most recently in the educational research community, over the appropriate role and priority that should be given to experimentation with random selection in funding research and development (R&D). This debate is intertwined with at least two other long-running discussions. The first is between the role and priority of “qualitative vs. quantitative” evidence. The second debate is about how best to arrive at reliable predictions for large-scale social or educational programs. The latter argument involves two important questions:

- Whether and under what conditions results from nonexperimental data can provide unbiased estimates of scaled-up effects of social and educational interventions/policies.
- Whether and how smaller scale experimentation, which cannot be generalized outside its specific context, can contribute to making reliable predictions for large-scale programs.

Turning small-scale interventions into large-scale interventions, or even implementing small-scale interventions in different contexts, can be problematic because results from small-scale RCTs cannot be

generalized beyond the specific experimental population and context. Using small-scale experimental programs as a way to identify and design efficient large-scale social and educational programs still remains an elusive goal. RCTs with multiple methods can help address these debates and critical issues.

These discussions have been intense in the educational research community because of recent strong funding support favoring experimentation. Intellectual leaders

Using small-scale experimental programs as a way to identify and design efficient large-scale social and educational programs still remains an elusive goal. RCTs with multiple methods can help address these debates and critical issues.

and researchers from several disciplines, the National Academy of Science, and federal policymakers in the Department of Education encouraged the use of RCTs as a high priority in R&D funding beginning around 2003 (Borman, 2002; Boruch, 1997; Chalmers, 2003; Cook, 2002, 2003; Duncan & Magnuson, 2003;

Feuer, Towne, & Shavelson, 2002; Mosteller & Boruch, 2002; Raudenbush, 2005; Shadish, Cook, & Campbell, 2002; Shavelson & Towne, 2002; Slavin, 2002; Towne, Shavelson, & Feuer, 2001). This movement generated heated deliberations centering on the role of experimental and nonexperimental research and qualitative and quantitative evidence in R&D funding and in improving long-term policies in education (see, e.g., Eisenhart, 2005, 2006; Eisenhart & Towne, 2003; Howe, 1998, 2004; Maxwell, 2004). Angrist (2004) provided an interesting history of

this line of argument and suggested a set of analytical techniques in evaluating RCTs that address some inevitable flaws in design and execution. However, the debate between the utility of experimental and nonexperimental evidence considerably predates this most recent flare-up (see, e.g., Chen & Rossi, 1983; Cook & Campbell, 1979; Cronbach & Shapiro, 1982; Heckman & Smith, 1995; Mosteller, 1995).

Multiple-methods RCTs address many issues relevant to these debates. Multiple-methods RCTs represent an evolution from “black box” experimentation—whose only purpose is to measure the impact of a particular intervention—to combining quantitative and qualitative research methods that allow researchers to address a broader set of questions (see the Introduction, p. 1). This can considerably enhance the scientific and policy value of RCTs. Using multiple-methods RCTs addresses the questions cited in the introduction by

- incorporating methods of data collection, commonly referred to as “qualitative,” that become indispensable and powerful tools within an RCT for understanding why and how the effects (or lack of effects) of an intervention occur and why effects differ among participants;
- using these qualitative data to explore and predict how results are sensitive to contextual effects, thereby improving predictions of effects in different and/or larger scale settings;
- providing opportunities to compare and contrast experimental and nonexperimental measurements and test hypotheses as to why such results differ, thereby potentially improving the methods and reliability of nonexperimental analyses; and
- focusing attention of the research not only on whether an intervention works but why it works, thereby contributing to building more general theories that can improve predictions in all settings and better prioritize what future experimentation to fund.

Multiple-methods RCTs are a partial response to a long-recognized need for theory-driven experimentation aimed not only at accurate measurements of interventions but also at accounting for why and how effects occur (Chen & Rossi, 1983; Cook, 2002; Cook & Campbell, 1979; Cronbach & Shapiro, 1982; Donaldson, 2007; Duncan & Magnuson, 2003; Heckman & Smith, 1995; Raudenbush, 2005; Romich, 2006; Walshe, 2007). Theory development is a critical complementary process to experimentation because successful theories can dramatically reduce the need for future experimentation and allow better priorities to be assigned to future experimentation. Without the parallel development of theories, the process of experimentation

will not converge but will instead lead to choosing from an infinite number of possible experiments.

Multiple-methods RCTs can also help address whether results of measurements using nonexperimental data are reliable, why results may differ between experimental and nonexperimental measurements, and under what conditions nonexperimental results are more reliable. For instance, contextual effects can explain differences between experimental and nonexperimental measurements, and multiple methods within RCTs can expand the range of contextual factors that can be tested for their influence. In this and other ways, multiple-methods RCTs can help sort and integrate the large body of nonexperimental research with experimental research. In the end, scientific consensus requires that researchers explain and reconcile both experimental and nonexperimental measurements. Multiple methods can not only help reconcile these measurements but also improve the reliability of nonexperimental measurements, thus helping to form scientific consensus.

Multiple-methods RCTs may represent a significant advance that uses complementary approaches in the pursuit of scientific knowledge, similar to those described by Salomon (1991), by (a) helping to address persisting questions and arguments in the research community, (b) developing stronger social and educational theories, (c) reconciling experimental and nonexperimental measurements, and (d) enabling improved external validity and better predictions of social and educational policies. However, multiple-methods RCTs will also have some significant limitations, mostly due to their increased costs and complexity. More experience is needed to test whether their potential contributions can be realized. Thus, RCTs with multiple methods do not ensure termination of the methodological and R&D policy debates (Howe, 2004).

Note. Each of the following chapters is related to one of the five boxes shown in the chart that accompanies this report (see insert at the back of this booklet).

BOX 1: MOTIVATING POLICY/ RESEARCH QUESTIONS

Construct an Intervention X to be tested by a study

- What does theory suggest would be effective?

This is the first question to undertake when thinking about conducting a multiple-methods RCT. It is critical from the start to think about the possible theories that might explain (a) why effects are expected, (b) why such effects might be different across participants, (c) whether results are sensitive to contextual factors, and (d) how to design an experiment that provides the information needed to refine and improve the intervention and why and how such effects might change when scaled up. Thinking through the alternate possible causal mechanisms and how and why such mechanisms might produce effects leads inevitably to the use of multiple methods.

In the present context, a theory can range from simple hypotheses to much more complex sets of interacting causal mechanisms that provide an explanation of why and how the measured effects occur. Sometimes there is no direct link to a theory that might apply to a potential intervention. Rather, unique hypotheses or theories might need to be developed for specific interventions. Burton, Goodlad, and Croft (2006) and Tilley (2004) provided clear examples of the range of simpler hypotheses that might account for the results of their

crime prevention experiments and why such hypotheses are critical in research. Romich (2006) suggested ways in which social policy experiments can advance theory-based knowledge in child development. Kling, Liebman, and Katz (2007) explicitly stated four hypotheses that might account for the mechanisms involved in how neighborhoods impact adult labor market outcomes. Cohen, Raudenbush, and Ball (2003) presented a theoretical approach to modeling the relationship between educational resources and achievement, focusing on a casual role for instruction.

Project STAR's first publication suggested three ways in which class size may affect achievement: enhancing teacher

morale, improving the number and quality of student-teacher interactions, or improving student engagement (Finn & Achilles, 1990). Evidence collected in the experiment included teacher surveys and logs and observational data suggesting greater fourth-

grade student engagement for those in smaller K-3 classes (Finn & Achilles, 1999). Multiple-methods data collected from teacher aides also helped to address why teacher aides in larger classes did not have statistically significant effects over large classes with no aides (Finn & Achilles, 1999; Gerber, Finn, Achilles, & Boyd-Zaharias, 2001).

Investment in theories can have low payoff if the measurements that theories are developed to predict are not accurate. Perhaps the most important role of RCTs

Perhaps the most important role of RCTs is to provide more accurate measurements that make development of increasingly refined theories productive.

is to provide more accurate measurements that make development of increasingly refined theories productive (Heckman & Smith, 1995). Project STAR's compelling experimental evidence spawned a rich theoretical literature directed toward understanding the causative mechanisms that created these effects. Research studies spanning several disciplines suggested hypotheses about classroom processes and parental effects that might account for achievement gains in small classes (Blatchford, 2003, 2005; Blatchford, Bassett, & Brown, 2005; Blatchford, Bassett, Goldstein, & Martin, 2003; Blatchford, Goldstein, & Mortimore, 1998; Blatchford & Martin, 1998; Bonesrønning, 2004; Boozer & Cacciola, 2001; Bosker, 1998; Bosworth & Caliendo, 2007; Datar & Mason, 2008; Finn, Pannozzo, & Achilles, 2003; Grissmer, 1999; Hattie, 2005; Lazear, 2001; Webbink, 2005). This literature is one of the best examples of theory development to explain an experimental effect. Such a literature can help specify what additional RCTs might be pivotal in deciding between theories and can help eliminate many areas of experimentation that would not make any contributions.

New Hope was partly built on a hypothesis that working at least 30 hours a week over a 3-year period (if supplemented by additional health, income, and child care benefits) could lift individuals who were not working or whose earnings were below the poverty line into lives of more stable employment and increased wages. These outcomes would then improve participants' lives and the lives of their children over a longer term (Duncan, Huston, & Weisner, 2007). The primary initial experimental measures focused on labor force behavior, and the results showed statistically significant effects for the treatment group. However, the researchers were initially puzzled by several issues. The control group participants made substantial employment and wage gains that did not depend on their having received New Hope benefits, and these gains were much larger than the incremental gains of New Hope recipients. Moreover, many eligible for New Hope benefits did not use them, or used them only sporadically. These results were inconsistent with the project's theoretical framework, and if not for multiple-methods data, New Hope would have left only unanswered questions and small contributions to theory and policy.

However, the multiple-methods data collections in New Hope enabled substantial contributions to understanding and designing new policies for welfare, child care, health, education, and employment to improve the lives of the working poor. A refocusing on wider outcome measures for working mothers with children enabled the development of theories that helped explain (a) the experimental results, (b) why this pattern of results emerged, (c) why benefit

utilization was much lower than expected, (d) why New Hope made the difference for some but not for others, (e) why control group participants made such large gains, (f) why effects for boys more than for girls were particularly large and sustained in achievement and behavior, (g) why more flexible menus of benefits might have enhanced effects, (h) why and what kinds of targeting would have improved efficiency, and (j) what key contextual factors and other issues present in Wisconsin would need to be addressed in any large-scale statewide or national interventions (Duncan et al., 2007; Huston et al., 2001; Yoshikawa, Weisner, & Lowe, 2006).

Besides its contribution to policy related to the working poor, New Hope serves as perhaps the best current model for designing, utilizing, and documenting multiple methods in RCTs and in illustrating that simple theories will be inadequate in predicting the complexity and often chaotic lives of this population. This contribution to research methodology and theory building may be its most important and longest lasting legacy. Multiple methods were used extensively in the following ways: (a) a comprehensive set of outcome measures using surveys and testing; (b) in-depth interviews with participants, their children, and the children's teachers; and (c) an ethnographic study of 44 families during and after the experiment. Moreover, the design, analyses, and documentation of New Hope multiple-methods data extended beyond academic journals, with separate documents designed for researchers and policy audiences. For example, Yoshikawa et al. (2006) provided a volume of studies using multiple-methods data to address research questions, while Duncan et al. (2007) directed their work to both policy and research audiences.

Researchers pursued the Moving to Opportunity (MTO) experiment because results from scores of nonexperimental studies suggested that living in poor neighborhoods may adversely affect a wide range of adult and child outcomes. The results, however, elicited concern about the high correlation of neighborhood characteristics with individual, family, and school characteristics and the strong possibility of selectivity bias in nonexperimental measurements. Different theories about neighborhood effects also predicted opposite directions for the outcomes. Kling et al. (2007) stated the theoretical hypotheses as follows:

It is hard to judge from theory alone whether the externalities from having neighbors of higher socioeconomic status are predominately beneficial (based on social connections, positive role models, reduced exposure to violence, and more community resources), inconsequential (only family influences, genetic endowments, individual human capital investments, and the broader non-neighborhood social environment

matter), or adverse (based on competition from advantaged peers and discrimination). (p. 84)

For instance, Wilson (1996) articulated a theory about why unemployment was high and wages were low for inner-city residents and the effect on neighborhoods of changes in job opportunities within and close to inner cities. The MTO experiment was designed to test, among other things, whether changing neighborhoods caused changing adult labor market outcomes and which theory better predicted the outcomes. Furthermore, it tested whether those outcomes improved, worsened, or did not change.

When the basic MTO experimental results showed no effect on adult labor market opportunities or children's achievement from moving to better neighborhoods, researchers used multiple-methods data to explore the reason for the null labor market effects (Kling et al., 2007; Turney, Clampet-Lundquist, Edin, Kling, & Duncan, 2006). In the process, they discovered that large effects were registered on adult mental health measures and on behavioral measures for children. Thus, the use of multiple-methods data not only helped to explain null results on the original variables of primary interest but also to validate some theories while dismissing others. In addition, multiple-methods data enabled the identification of important outcomes not included in the original objectives. Clampet-Lundquist, Edin, Kling, and Duncan (2006) provided another example of using multiple-methods data to test four explicit hypotheses that might account for the positive behavioral effects shown for boys, but not for girls, in the MTO experiment.

• What is known from previous research?

Without a theory available to explain why differences are present, reviewing previous research has always been a difficult and nuanced task because of the almost universal disparity in outcomes present in previous studies. The critical question is how to distinguish among the studies—often among a multitude of studies—that could be relevant and how to synthesize these studies in the most meaningful way. For example, a long-running debate in education (mainly from the 1980s to early 2000s) addressed the effect of additional resources on educational outcomes. Several research studies used differing techniques to select and weigh the value of studies (including meta-analysis) to arrive at contrasting conclusions (Greenwald, Hedges, & Laine, 1996; Hanushek, 1997, 2002; Krueger, 2002, 2003).

A major motivating factor for RCTs and, in particular, for multiple-methods RCTs, is to move future literature reviews toward a scientific and/or policy consensus on a given question. The absence of consensus in previous nonexperimental studies has been a major motivating factor in moving toward experimentation. However, black box experimentation alone may not create either research or policy consensus because of the lack of generalizability of such experiments to different and larger scale settings and the inevitable flaws present in most social and educational experimentation. Consensus will require being able to explain why the current set of both experimental and nonexperimental measurements differ. Black box experimentation alone usually fails to provide evidence for why experimental and nonexperimental measurements are different, but multiple-methods data can provide considerable help in addressing these differences.

Four important reasons why previous results from various studies exploring the same phenomena can differ are methodological bias, the presence of contextual effects, differences in the characteristics of the population studied, and structural or other changes in programs/interventions during scale-up such that predictions from smaller scale programs have little predictive accuracy. Addressing potential bias requires a thorough knowledge of the strengths and weaknesses inherent in the various methodologies used in previous work.

A relatively new strategy groups studies into the following categories: experimental, quasi-experimental, “natural” experiments, and nonexperimental methods. Webbink (2005) provided an example of this type of review. However, within each of these categories there is usually wide variation in quality. Thus, simple categorization can be misleading. Duncan and Gibson-Davis (2006) and Duncan, Magnuson, and Ludwig (2004) provided advice on how to critique and interpret nonexperimental results. Cronbach and Shapiro (1982) and Heckman and Smith (1995) provided critiques of experimental studies. Cook, Shadish, and Wong (2005) compared and contrasted experimental and quasi-experimental results. Rosenzweig and Wolpin (2000) and O'Connor (2003) both provided perspectives from developmental psychology and economics on the strengths and weaknesses inherent in natural experiments.

Multiple-methods RCTs should be designed to address and settle issues that prevent the establishment of consensus in a literature review. For instance, multiple-methods RCTs can provide evidence on contextual effects that help reconcile previous disparate results. It is also possible to design

A major motivating factor for RCTs and, in particular, for multiple-methods RCTs, is to move future literature reviews toward a scientific and/or policy consensus on a given question.

multiple-methods RCTs that incorporate a nonexperimental measurement. For example, although Project Star did not incorporate a nonexperimental measurement component, two later studies chose nonexperimental samples from Tennessee and compared and contrasted experimental and nonexperimental measurements of particular outcomes. Krueger (1999) compared STAR experimental findings with results estimated nonexperimentally from the variation in large class sizes that show similar experimental and nonexperimental results. Using propensity scoring, Wilde and Hollister (2007) showed significant differences comparing experimental and nonexperimental results with a sample of Tennessee students outside the STAR experiment.

- **Consider relevance for the population(s) of interest.**

Using multiple methods in RCTs can be viewed as road testing a prototype intervention and can be initiated in the planning and design stage in the form of a “mini-efficacy trial.” The purpose is partly to obtain feedback from the population of interest about their attitudes, reactions, or predictions, as well as to suggest changes to a particular intervention. Another rationale for an efficacy trial would be to determine the groups that should be targeted for inclusion in a study.

Techniques such as focus groups, interviews, and surveys of a sample of participants might be appropriate for the planning stage of a multiple-methods RCT. Focus groups allow for exploring the appropriateness of the intervention, an array of participant reactions, potential new design features, and more. Interviews can allow a two-way conversation focusing on these same topics. Surveys can be less expensive when larger samples are required, but they lack the flexibility for unstructured feedback. Brock, Doolittle, Fellerath, and Wiseman (1997) and Poglinco, Brash, and Granger (1998) provided an example of using multiple methods in an efficacy trial on a small population of potential participants in New Hope during the extensive preplanning for the major study.

BOX 2: DESIRABILITY/FEASIBILITY OF AN RCT STUDY

Desirability

- Is Intervention X well-enough developed/defined to warrant a controlled study, or are efficacy studies needed first to clarify constructs and establish the basic efficacy of the proposed interventions?

The development of a theory, the review of the literature, and the use of multiple methods in the planning and design stage provide information for making a decision whether to proceed first with a small-scale efficacy trial or a larger and more formally structured RCT. Since RCTs, especially multiple-methods RCTs, are substantially more costly and require much more planning than efficacy trials, conducting efficacy trials prior to multiple-methods RCTs is likely to become the rule rather than the exception. Efficacy trials not only establish viability for an intervention—and provide potential redesign and retargeting insights to make it more effective—but also allow field testing for multiple-methods data collections and eventual redesign. Multiple methods may be as important to efficacy trials as they are to structured RCTs (see the previous Consider Relevance for the Populations of Interest section, p. 8). For instance, Duncan et al. (2007, pp. 23–26) described their 50-participant pilot project and adjustments made in the later intervention as a result of the pilot. According

to the researchers, the pilot project seemed crucial to understanding the population of interest and matching the program benefits to that population.

- Are the results generated likely to be worth the expense?

This question can be viewed in two different contexts: the R&D or scientific context and the policy context. The R&D or scientific questions address whether a particular intervention is a sound use of R&D funds. In this context, the normal scientific criteria used in peer reviews are relevant to evaluate proposed RCTs or other research approaches. An increasingly important question will be how a multiple-methods RCT will contribute to testing a particular theory or set of theoretical hypotheses. Without multiple methods, it will often be difficult to provide a strong argument about why an RCT would contribute to theory. If the only outcome is the measurement of the intervention effect—even if done with a sound design—the contribution to theory will often be minimal. A sound design for a multiple-methods RCT that addresses the questions cited in the Introduction can significantly enhance the scientific value of a research proposal. In contrast, it may be increasingly difficult to make the scientific case for black box RCTs due to their limited contributions to theory and

Without multiple methods, it will often be difficult to provide a strong argument about why an RCT would contribute to theory.

inability to provide explanations for disparate results from previous studies.

With regard to the policy context, the researcher must determine the value of the intervention—if successful—to society. The questions that arise in this context are not only costs versus benefits to society that would result from a successfully scaled-up intervention but also the chances that a successful small-scale intervention could be widely scaled up without a significant deterioration of effects. These questions are more difficult to address with black box RCTs than with multiple-methods RCTs. The latter methods provide much more information about potential scale-up issues arising from contextual effects; how to target an intervention to the population, showing larger effects; and how to redesign the intervention to make it more effective. Again, black box RCTs have less policy value when compared with a feasible multiple-methods RCT.

Perhaps more than any other single publication, the study by Kling, Liebman, and Katz (2005) should be read by those researchers and policymakers who question the scientific and/or policy value of funding multiple methods in RCTs. We cited this article in the Introduction, but the article goes on to elaborate in more detail how in-depth interviews influenced their research:

Our qualitative fieldwork had a profound impact on our MTO research. First it caused us to re-focus our quantitative data collection strategy on a substantially different set of outcomes. In particular, our original research design concentrated on the outcomes most familiar to labor economists: the earnings and job training patterns of MTO adults and the school experiences of their children. Our qualitative interviews led us to believe that MTO was producing substantial utility gains for treatment families, but primarily in domains such as safety and health that were not included in our original data collection plan. *In our subsequent quantitative work, we found the largest program effects in the domains suggested by the qualitative interviews* [italics added].

Second, our qualitative strategy led us to develop an overall conceptual framework for thinking about the mechanisms through which changes in outcomes due to moves out of high poverty areas might occur. *Our conversations with MTO mothers were dominated by their powerful descriptions of their fear that their children would become victims of violence if they remained in high poverty housing projects* [italics added]. . . . This fear appeared to be having a significant impact on the overall sense of well-being of these mothers, and it was so deep-seated that their entire daily routine was focused on keeping their children safe. . . . We hypothesized that the need to live life on the watch may have broad implications for the future prospects of these families.

Third, our fieldwork has given us a deep understanding of the institutional details of the MTO

program. This understanding has helped us to make judgments regarding the external validity of our MTO findings, particularly regarding the relevance of our results to the regular Section 8 program. *In addition, this understanding has prevented us from making some significant errors in interpreting our quantitative results* [italics added].

Fourth, *by listening to MTO families talk about their lives, we learned a series of lessons that have important implications for housing policy. For many of the things we learned, it is hard to imagine any other data collection strategy that would have led to these insights* [italics added]. (pp. Kling, Liebman, & Katz, 2005, pp. 244–245)

Feasibility

• Are the factors of interest amenable to experimental manipulation and control in the real world?

Can a particular intervention be successfully tested in an experimental framework? Social and educational experiments inevitably depart from ideal experimental conditions. These departures can sometimes seriously mitigate the scientific advantages that are inherent in ideal experiments. Gueron (2002), who had 30 years of experience at the Manpower Demonstration Corporation (which pioneered large-scale social welfare experimentation), provided the best resource for understanding the complexity of actually doing an RCT. Her article covers most of the real-world constraints and limitations that can compromise the internal validity of such efforts. Gueron (2003, 2007) also provided unique perspectives on both the difficulty and the utility of doing social welfare experiments.

Heckman and Smith (1995) examined a group of social experiments that measured the impacts of job-training programs conducted in the 1980s and early 1990s. One of their conclusions is that significant deviations from experimental conditions destroyed much of the scientific value of the results. These deviations were often peculiar to particular experiments, but their article identified and characterized many of the vulnerabilities inherent in social experiments. As a result, it is important to assess the susceptibility of a proposed intervention to the potential vulnerabilities that have plagued social experiments such as those described by Heckman and Smith.

Another potential vulnerability in experimentation is that the “signal-to-noise ratio” will turn out to be too small for successful measurement. In any intervention, there is an unbiased effect size (i.e., the signal), which in the absence of any noise (i.e., bias, errors, and uncertainties created by a finite sample) would emerge from an RCT. Yet, there are always conditions that will introduce noise. Some of

this noise can be predicted and limited by the selection of sampling parameters. However, such an analysis can only take account of sampling uncertainty and cannot fully take into account the inevitable other sources of noise from random sources and flaws present in social experiments.

In order to have successful RCTs, it is necessary for the signal to emerge clearly from the noise (a favorable signal-to-noise ratio). Since the amount of noise is always uncertain, a researcher would optimally like to create an intervention with a large signal. Two factors often control and limit the strength of the signal in an experiment. The first is that the costs of the experiment will usually increase with larger variations in the treatment. For example, the cost to the state of Tennessee for Project STAR was about \$13 million to sustain class size differences of 24 versus 16 pupils per class over 4 years. The costs were nearly proportional to the level of class size reduction. Although smaller reductions would cost much less, it is unlikely that a class size reduction of 2–4 students (a small signal) per class would have produced such definitive results.

The second factor limiting signal strength is that large variations can generate political problems arising from inequitable treatment of test and control participants (Gueron, 2002). As a result of such large differences in class sizes maintained over 4 years, parents of pupils assigned to large classes started lobbying for their children to participate in the experimental classes. Some parental opposition was mitigated by randomly dividing the large classes into two groups, with teacher aides in one group, leaving a smaller group of children without benefit. However, about 15% of the children assigned to large classes appeared in small classes over the course of the experiment—likely due to parental pressure (Finn & Achilles, 1999). Thus, large signals may also cause some compromise in the integrity of the experiment.

It is thus important to consider potential reactions of participants in the control group. For instance, Duncan et al. (2007, pp. 42–43) described the reactions of some New Hope participants to being assigned to the control group. Randomization was clearly described to participants from the beginning of their potential involvement. The researchers explained that although some participants would not receive New Hope’s additional benefits, no participants would lose any benefits as a result of the experiment. Nevertheless, some participants who “lost” the lottery were disgruntled and painted a negative picture of the program to the researchers.

- **Can an RCT be conducted without encountering ethical constraints?**

Any social or educational experiment will have to balance the potential benefits of carrying out the experiment with the possible costs to participants. This balancing is ultimately evaluated by human subjects review boards who have the independence and authority to protect research participants. However, because this balancing is often difficult and not straightforward, studies involving any significant costs to participants or possible ethical issues need early review by such boards.

Generally, the limitations on experiments imposed by ethical considerations and review boards would be predicted to cause some compromises from ideal experimental conditions. Gueron (2002) provided a real-world perspective on the ethical issues that arose in carrying out over 2 decades of social welfare experimentation, as well as some necessary compromises that sometimes limit internal and external validity. Often, these compromises can be addressed analytically in a way that still maintains the advantage inherent in random assignment. Angrist (2004) illustrated analytical methodologies that address noncompliance in treatment groups and crossovers from control groups. These techniques allow some flexibility in service denial or compelling participation without undue compromise in measuring and interpreting effects from random assignment experiments.

- **Is it likely that the study would gain the necessary cooperation and enough recruits to be assigned randomly to treatment conditions?**

Project STAR was mandated by the Tennessee legislature. Schools with at least three kindergarten classes were invited to participate (this eliminated smaller schools from consideration). About 100 schools volunteered to partake in the randomization of children to classes, and 79 schools were selected to participate. Project STAR was conducted prior to the need for parental permission for children’s involvement in research, so during the experiment virtually all students entering kindergarten in these schools took part, as did all students entering these schools in Grades 1–3. Given the compressed time available for planning Project STAR, and had the project been conducted more recently, it is possible that parental permission would have been a significant obstacle to carrying out the experiment. This likely would have introduced limitations on external validity and potential selectivity bias.

Unlike Project STAR, many experiments do not have a state mandate for participation. In assessing whether

sufficient numbers of participants can be recruited, there are three main considerations:

- Can a sufficient number of volunteers be obtained for the participation “lottery”?
- Will a sufficient number of lottery “winners”—those assigned to the experimental condition—actually use the benefits offered?
- How different (selective) will the volunteers and those using benefits be from the actual population of interest?

Both MTO and New Hope encountered some difficulty not only in recruiting “volunteers” for the lottery but also in having participants (compliers) in the treatment group actually take advantage of the benefits offered.

Duncan et al. (2007, pp. 36–41) described a year-long effort to recruit 1,357 participants to New Hope and the subsequent effort to understand why some participants did not take full advantage of the project benefits. Although New Hope participants ended up approximately matching the racial/ethnic characteristics of a national sample of working poor individuals, participants’ choice to take part likely made them somewhat different from those who did not volunteer based on other characteristics. Those who became eligible for New Hope benefits ended up using those benefits less than expected. Fortunately, the original sample was large enough (1,357) to allow a focus on working mothers (745)—the group that used benefits most often and accounted for much of the program effects. Because the subgroups of primary interest often emerge only after initial analysis, larger samples offer some insurance against this problem of lower than expected utilization.

Moving to Opportunity was carried out in five cities and recruited individuals from public housing in high-poverty areas. Individuals who volunteered may have been more motivated to move out of public housing. Over a 4-year period, over 4,000 families volunteered for the program. Of those who were randomly assigned to the treatment, about 47% actually moved (Kling et al., 2007). Clearly, self-selection of volunteers into the lottery group and further selectivity in the treatment group has the potential to bias experimental results and limit generalizability. Heckman and Smith (1995) provided further examples of this kind of selection bias in job training experiments.

Multiple-methods RCTs will be more likely than black box experimentation to provide “unexpected” results and/or discover unanticipated opportunities that would require added funding to exploit.

- Will funding be sufficient to support an RCT design with adequate statistical power?

Multiple-methods RCTs can be significantly more costly than black box experimentation. For instance, larger sample sizes are often required to measure whether effects differ across participants with different characteristics. Further, the additional data collection required in multiple-methods RCTs can add significant costs. In the longer term, the benefits derived from multiple-methods RCTs can be substantial because the derived theories can more efficiently guide future experimentation. In the short run, however, they will increase R&D costs. These additional costs may represent a significant barrier to proposing multiple methods—especially between black box RCTs and multiple methods RCTs—in the absence of firm guidance from funding agencies about their priorities.

A power analysis is the usual method used to determine the number of participants required to measure different effect sizes with various degrees of certainty. It is always difficult to incorporate the wide range of

possible factors that can introduce additional uncertainty and bias into any social experiment. Yet failure to incorporate these factors can make experiments too weak to measure desired effects accurately. Often it is the case that experimental effects are widely and unpredictably different across participants, and the major contribution of the study is to examine

what is causing the effects in a subpopulation of interest. Sample sizes larger than those dictated by power analysis provide some assurance that such effects can be studied.

New Hope had a board of directors whose responsibility was partly to garner necessary funding to carry out the project. Although funding was obtained to initiate the program, additional funding for various research components was added during and after the experiment. In the end, over 50 different foundations, government agencies, and businesses provided financial support for New Hope. The original funding to support a target population of about 1,200 allowed the experiment to start. However, additional funding supported a “family” study that incorporated multiple methods into the data collection. About two thirds of the participants of the family study were individuals with children—making the sample adequate for studying this population (Duncan et al., 2007). The later funding of this sample proved crucial in making the entire experiment so valuable. Multiple-methods RCTs will be more likely than black box experimentation to provide “unexpected” results and/or discover unanticipated

opportunities that would require added funding to exploit. In Project STAR, MTO, and New Hope, later funding to refocus study objectives, explore hypotheses in more detail, or do longer term follow-ups was crucial to their scientific and policy utility.

Occasionally, funding can be generous. Project STAR was funded by the state of Tennessee as part of a compromise that delayed the institution of smaller class sizes statewide until the completion of the study (Ritter & Boruch, 1999). The \$13 million needed to fund the study was a small proportion of the costs of implementing a policy of smaller classes statewide. This experiment represented a compromise between implementing an expensive statewide program and funding an experiment. In such circumstances, the sums needed for experimentation looked small to legislators but large to researchers. This sum supported very large sample sizes (over 6,000 in the kindergarten cohort) and extensive multiple-methods data collection. Project STAR encountered much more difficulty later in raising the smaller amounts of funding required for long-term follow-ups.

- **Will I know afterward what conditions are necessary for the intervention to be effective?**

One of the major vulnerabilities of black box experimentation is that it often provides little power for predicting effects in different contexts, in generalizing to different populations, or in scaling up programs. One of the major advantages of multiple-methods experimentation is an increased ability to estimate how impacts might change in different contexts, populations, and scaled-up versions. Perhaps the major reason for doing multiple-methods RCTs is the ubiquity of contextual effects in social and educational interventions and the need to develop theories that can make better predictions that apply to different contexts and populations.

The California Class Size Reduction initiative, which was partly motivated by Project STAR (and a huge one-time budget surplus in California), is being used as the poster child for the lack of predictability from contextual and scale-up effects. Project STAR was not a small-scale experiment but rather a fully scaled-up experiment involving 79 schools and over 12,000 children (Finn & Achilles, 1999). Much was learned from Project STAR that could guide policy and implementation in other states. Three important lessons were the following:

- Three to four years of small classes, starting with kindergarten, were needed for long-term effects.
- Effects were much larger for minority and disadvantaged children.
- The teachers in Project STAR were not newly recruited but were drawn from the pool of existing experienced teachers.

California mandated class size reductions statewide from around 30 to 20 pupils per class in Grades K–3 beginning in 1996 (Bohrnstedt & Stecher, 2002). The legislation passed one month before the start of school, along with

strong monetary incentives for immediate implementation and a set of rules governing implementation. These rules stated that before second-grade classes could be reduced, implementation should begin in first grade and be completed for all students. Likewise, second-grade implementation had to be

completed before either kindergarten or third-grade classes were reduced. These rules meant that kindergarten classes were not reduced until much later in the process and that most children in the first few years had less than 3–4 years of consecutively smaller classes. The program was not targeted to minority or disadvantaged children but rather to all children in Grades K–3.

The reduction placed a huge immediate demand on a teacher labor market unable to enhance the supply of teachers in the short run. Newly recruited teachers were often inexperienced and lacked certification, and classroom space was less than ideal. An evaluation concluded that the reductions had small effects on achievement in the short term (Bohrnstedt & Stecher, 2002) and cited these contextual effects as a possible explanation for the results. Jepsen and Rivkin (2002) did a longer term evaluation and found somewhat larger effects.

Unlike Tennessee, California did not prudently phase in the program beginning in kindergarten so that all children would experience 3–4 years of smaller classes, nor did they target the intervention to minority and disadvantaged children. This failure to phase in the program slowly led to shortages of teachers and classroom space and smaller short-term effects from children receiving only 1–2 years of smaller classes. Another unintended side effect of the California initiative was that the number of combination classes (more than one grade taught in a classroom) increased, and analysis of effects for these children showed negative results (Sims, 2004).

One of the major advantages of multiple-methods experimentation is an increased ability to estimate how impacts might change in different contexts, populations, and scaled-up versions.

Haste in implementation also failed to ensure that sufficient data were collected and available to provide unbiased measurements of short-term effects. For instance, comparable test scores were not available for the years prior to the experiment, and thus the evaluation lacked a critical source of comparative evidence. An unintended consequence of the rush to small classes and not targeting treatment to specific populations was that many high-quality teachers in central city schools left for the suddenly available jobs in suburban schools. Inner-city schools not only had to recruit teachers to reduce class size but also had to fill additional vacancies caused by those moving to suburban schools. These changes meant that it was impossible to predict California effects from Project STAR effects, or to provide unbiased measurements of the short-term effects of small classes in California.

In the case of MTO, there were no significant effects on the primary measures of employment, wages, and children's achievement. Turney et al. (2006) explored with multiple-methods data what might explain the null labor force effects and under what conditions positive effects might have been expected.

Duncan et al. (2007, chap. 5) used multiple-methods data to investigate why the impact on some New Hope participants was high while for others it was low. Part of this difference was predicted by the context of participants' lives. For some, their lives were dominated by serious obstacles like domestic abuse or addiction that prevented them from taking advantage of New Hope benefits. For these participants, the conditions necessary for effective intervention would have involved addressing those issues. For other participants, mainly men and women without families, utilization of benefits was low, and many were able to make significant labor market gains without New Hope. New Hope women with children had conditions in their lives that allowed them to make the most of New Hope benefits, including improved and reliable child care and health care, which enabled them to make gains for themselves and their children.

BOX 3: EMPLOYING MULTIPLE METHODS IN DESIGNING AND IMPLEMENTING RCTs

Considerations for Internal Validity

- What factors led to Intervention X working? Failing?
- What factors led to Intervention X working for some groups and not others?

Project STAR's first publication hypothesized that reduced class size would affect achievement in at least one of three ways: (a) by enhancing teacher morale, (b) by increasing student-teacher interactions, and (3) by increasing student engagement (Finn & Achilles, 1990). Later analysis of their teacher and classroom data at the fourth-grade level supported the student engagement hypotheses over the other two (Finn et al., 2003). In its aftermath, Project STAR appears to have spawned a rich theoretical literature and set of research studies spanning several disciplines that suggest hypotheses and theories about classroom processes and parental effects that might account for achievement gains in small classes (Blatchford, 2003, 2005; Blatchford et al., 1998, 2003, 2005; Blatchford & Martin, 1998; Bonesrønning, 2004; Boozer & Cacciola, 2001; Bosker, 1998; Bosworth & Caliendo, 2007; Datar & Mason, 2008; Finn et al., 2003; Grissmer, 1999; Hattie, 2005; Webbink, 2005). This literature serves as an example of the development of theories to explain an experimental effect and what such theories can look like. For instance, this work lent some support to the hypothesis that increased time spent by teachers with individual students in small classes might explain part of the intervention effect.

Project STAR also found larger short-term effects for minority and low-income children (Finn & Achilles, 1990, 1999; Krueger, 1999). At eighth grade, the reported long-term effects were somewhat mixed on whether there were differential achievement effects for minority and low-income students (Finn & Achilles, 1999; Krueger & Whitmore, 2001; Nye, Hedges, & Konstantopoulos, 2000a, 2002, 2004). But minority and low-income students were significantly more likely than similar students in large classes to take college admission tests, graduate from high school, and enroll in advanced courses (Finn, Gerber, Achilles, & Boyd-Zaharias, 2001; Finn, Gerber, & Boyd-Zaharias, 2005; Krueger & Whitmore, 2001). Evidence does suggest that teachers spend more time involved in one-on-one interactions with students in small classes (Blatchford et al., 2003).

Grissmer (1999) suggested that short-term differential effects may be due to increased individual teacher time devoted to minority and low-income students in smaller classrooms that compensate for lack of parental time with children on school-related topics. The lack of class size effect for more advantaged students may be due to shifts in parental time and resources in response to class size. For instance, parents may devote more time when class sizes are larger. Datar and Mason (2008) and Bonesrønning (2004) explored whether increased class size influenced types of parental behaviors. Ideally, Project STAR would have collected mixed-methods data from parents to assess whether parental time spent with children on school topics

is different across racial and socioeconomic status (SES) groups and whether parental time and resources change in response to changes in class size.

Duncan et al. (2007), Yoshikawa et al. (2006), and Weisner (2005) provided the richest examples in the literature of how data collected with multiple methods can be used to provide explanations and refine theoretical hypotheses about results. Duncan et al. (chap. 5) explained through the use of multiple-methods data why some participants took more advantage of benefits and made larger gains in income or employment than others. For instance, among women with children, the data suggest that about 20% of families had problems (drugs, alcohol, domestic abuse, etc.) that could not be addressed by the New Hope benefits offered. Another significant portion of participants eligible for benefits was not constrained in their economic life by factors that the benefits could address. For instance, no employment or income effects were measured for women without children partly because two of the key benefits of child care and health insurance for children were not barriers to employment for these women.

Duncan et al. (2007, pp. 77–79), drawing from Yoshikawa et al. (2006) and Huston et al. (2001), also used multiple-methods data from New Hope to summarize why behavior changes resulting from New Hope interventions were different for boys and girls. They suggested that boys' existing higher levels of risk, especially in poor neighborhoods, may have led parents to favor providing more and better day care and after-school activities for boys than for girls (as indicated by higher enrollment for boys than for girls).

The Moving to Opportunity experiment produced no significant effects on participants' earning and labor force behavior or achievement scores of their children (Kling et al., 2007; Sanbonmatsu, Kling, Duncan, & Brooks-Gunn, 2006). Turney et al. (2000) used multiple-methods data to develop hypotheses as to why earnings and employment did not change much in response to the intervention. Unanticipated effects that were large and significant occurred for mental health measures of adults' and children's behavior (Kling et al., 2007). For instance, Clampet-Lundquist et al. (2006) used multiple-methods data in the MTO experiment to try to explain differences in behavioral effects for boys and girls.

In some of these instances, the researchers collected data well into and after the experiment that were not part of the original design to further explore causal mechanisms and

differential effect sizes by group. Although some multiple methods can be built into the design of RCTs, it may also be efficient to institute a flexible response capability that allows for introducing new multiple-methods data collections in response to early RCT findings, especially if such findings are unexpected. In some cases, it is even possible to follow up with participants long after the end of the experiment to explore theoretical hypotheses. For instance, in Project STAR, while follow-up measurements included effects on high school graduation, taking college entrance exams, and advanced courses, no follow-up has thus far tried to collect data that would attempt to explain what differences between individuals in small and large class sizes might explain these long-term effects.

Include collection of baseline demographic and other measures to confirm that randomization was accomplished.

Randomization will still leave differences in average characteristics of treatment and control groups. Establishing baseline characteristics of the treatment and control groups can identify those characteristics when average differences do exist. If there are such differences, it would be important to include variables for those characteristics in equations estimating treatment effects. Although Project STAR shows that the average demographic characteristics of treatment and control groups were similar,

it would have been desirable to collect baseline test score data at the beginning of kindergarten. In New Hope, researchers collected tracking data on work, benefit usage, and supplementary income and state benefit utilization from the beginning of the experiment and showed balance between test and control groups. However, the first extended comprehensive survey was not conducted until 2 years after random assignment, when more detailed analysis of the results of randomization could be checked.

Use structured interviews and/or surveys to (1) assess fidelity of implementation, (2) document the existence of local policies and practices that might affect the outcomes of interest, and (3) document changes that occurred before and during the study (i.e., "history").

Use interviews or surveys to learn how subjects experienced the intervention.

Although some multiple methods can be built into the design of RCTs, it may also be efficient to institute a flexible response capability that allows for introducing new multiple-methods data collections in response to early RCT findings, especially if such findings are unexpected.

In Project STAR, researchers annually captured the experiences of K–3 teachers and teacher aides with time logs and a year-end survey that asked about their experiences in small and large classes. Researchers also asked fourth-grade teachers to assess the learning behaviors of each child in the experiment. Gerber et al. (2001) offered an analysis of time logs for teacher aides in the experiment, and Finn and Achilles (1999) supplied an analysis of teachers’ assessments of learning behaviors at fourth grade for each child. The latter data revealed that teachers had more positive perceptions of their students’ learning behaviors if their students had been in smaller classes in previous grades. Data collections and their analyses were invaluable in developing theories and hypotheses about why regular-sized classes with teacher aides did not have significant achievement effects and why small classes did have effects. It would also have been valuable to have captured data from children and their parents about their experiences in small and large classes. In addition, it would have been beneficial to have collected data from control and experimental groups in high school and beyond that could help explain why long-term effects persisted and why particularly large positive differences in high school graduation occurred for minority groups.

New Hope collected interview data with participants in the second year and at long-term follow-up 5 and 8 years after initiation. More important, researchers initiated an intensive sub-study during the experiment that involved the 745 participants with children to look at family effects and effects on children’s school performance (standardized assessments in mathematics and reading) and behavior. This opportunistic sub-study provided in-depth information about how parents and children experienced and changed as a result of the intervention. It incorporated the design and fielding of new surveys that included detailed information from parents about their children and also included interviews with children. In addition, children’s teachers completed surveys to report on performance and behavior.

Finally, researchers incorporated a unique ethnographic study that targeted 44 families. This study involved repeated visits and open-ended home interviews from 1998 to 2001 and again in 2004. Yoshikawa et al. (2006) provided 13 analyses of the ethnographic data by addressing a number of questions that illustrate the power of such data in understanding the lives and experiences of poor working mothers and their children. Duncan et

al. (2007) presented an interpretation of results using ethnographic data as follows:

The New Hope offer made a big difference for some people, but it was not a good fit for others. Some parents refused to entrust their children to the care of someone other than a family member. Many parents worked evenings and weekends, when few child-care centers or licensed homes were available. The child-care subsidy was therefore of little use to them. (p. 13)

Two books and articles offer rich perspectives on the issues and analyses of specific RCTs with mixed-methods data. Yoshikawa et al. (2006) provided a detailed set of analyses and interpretations of mixed-methods data from New Hope that incorporated the ethnographic data. This

volume is probably the single best resource for illustrating the value of mixed-methods data collection within a specific RCT. Weisner (2005) supplied a wider set of examples of RCTs that used mixed methods and the issues and analyses linked to inclusion of such data. These volumes consider the issues involved in the entire process, from idea origination, to

design of data collection instruments, to their analyses, and to interpretation and integration with the other data from the RCT. Yoshikawa, Weisner, Kalil, and Way (2008) provided a recent perspective on using multiple methods in developmental science and the range of methodological choices available in implementing mixed methods. Gardenhire and Nelson (2003) provided an assessment of the challenges and benefits of qualitative data in four RCTs, including New Hope.

A unique use of the ethnographic data allowed Duncan et al. (2007) to gather information about the lives and experiences of 3 participants in New Hope in a way that illustrates indelibly the complexity of the lives of poor families and their children. Since the life experiences of poor families differ dramatically from the lives of those who set policies and do research, understanding these lives remains a significant barrier to better policy outcomes and research questions. Mixed-methods data of the type collected and analyzed in New Hope can help bridge this “cultural” gap by increasing appreciation for the lives of those targeted by interventions. Such knowledge leads to improved theories, better design of future interventions, and better choice among candidate interventions.

Data collections and their analyses were invaluable in developing theories and hypotheses about why regular-sized classes with teacher aides did not have significant achievement effects and why small classes did have effects.

Check measured outcomes for indications that Intervention X worked better for some groups than others.

Each individual has a unique genetic endowment and follows a unique environmental trajectory. Environmental effects are also largely expressed through gene–environment interactions (Rutter, 2002). Both individual uniqueness and interactional dynamics make it unlikely that interventions will have identical effects across participants in any social or educational intervention. Exploring whether effects are different across groups is critical because the cost-effectiveness or benefit–cost ratios of interventions can be made more favorable by targeting interventions to those groups with larger effects (Grissmer, 2002).

Studies by Finn and Achilles (1999), Krueger (1999, 2002), Krueger and Whitmore (2002), and Nye et al. (2000a, 2002, 2004a) contained analyses of class size effects for Project STAR by income, race, and achievement level. These analyses of short-term effects from Project STAR have always found larger effects for minority and low-income students. In the longer term, reported differential effects were more mixed for eighth-grade achievement but were strongly significant for high school graduation and levels of college entrance test taking (Finn et al., 2005; Krueger & Whitmore, 2001).

Duncan et al. (2007, chap. 5) explained through multiple-methods data why some participants used benefits more and made larger gains in income or employment than others. They differentiated income and labor force effects for men without children, women without children, and women with children. For instance, among women with children, the data suggest that about 20% of families had problems (drugs, alcohol, domestic abuse, etc.) that could not be addressed by the New Hope benefits offered. Another significant portion of participants eligible for benefits did not become engaged with the program for a variety of reasons. No employment or income effects were measured for women without children partly because one of the key New Hope benefits was day care, which has been shown to be a critical barrier to employment for women with children. The barriers for women without children were different and were often not addressed by New Hope benefits.

Use more intensive interviews, case studies, and ethnographic research to investigate reasons for variability of effects within and between groups.

Several chapters/articles shed light on how to design and conduct multiple-methods data collections in RCTs (Brock, 2005; Brock et al., 1997; Cooper, 2005; Cooper, Brown, Azmitia, & Chavira, 2005; Datta, 2005; Duncan & Raudenbush, 1999, 2001; Fricke, 2005; Gibson-Davis & Duncan, 2005; Goldenberg, Gallimore, & Reese, 2005; Greene, 2005; Harkness, Hughes, Muller, & Super, 2005; Huston, 2005; Weisner, 2002; Weiss, Kreider, Mayer, Hencke, & Vaughan, 2005).

In the New Hope study, researchers used extensive interviews and ethnographic data to develop theories and hypotheses about the reasons for differential effects (Duncan et al., 2007; Huston et al., 2001; Weisner, 2005; Yoshikawa et al., 2006). For instance, Duncan et al. (chap. 5) created a new categorization that distinguishes participants by “potential obstacles” of utilizing New Hope benefits—partly based on interview and ethnographic data. For families with substantial barriers (drug and alcohol abuse, arrest records, presence of developmentally impaired children, domestic abuse, etc.), New Hope did not offer much to target such barriers, and intervention effects on these families were small or nonexistent. Some participants in both test and control groups, however, proved their abilities to accomplish New Hope objectives without New Hope benefits, leading to small or nonexistent overall effects. The largest effects were for families “poised to profit” from the specific benefits offered by New Hope. For instance, the child care benefit allowed some parents to upgrade the quality of their day care substantially. This illustrates that mixed-methods data allow analysis of categories that go far beyond the usual gender, race, and income categories.

Huston et al. (2001) explained why the achievement and behavior of boys improved more than for girls for New Hope participants, and why girls’ school behavior actually deteriorated for New Hope participants. In both cases viable hypotheses emerged from the mixed-methods data that placed boys at existing greater risk in poor neighborhoods, leading parents to favor using extra resources to protect boys against such risk. For instance, in New Hope, it was found that boys more often than girls participated in after-school programs with academic and recreational activities.

Bos, Duncan, Gennetian, and Hill (2007) provided an example of employing in-depth interview data to highlight the fear associated with threats to safety in the lives of poor families, especially single-parent families:

Exploring whether effects are different across groups is critical because the cost-effectiveness or benefit–cost ratios of interventions can be made more favorable by targeting interventions to those groups with larger effects.

In the qualitative sub-study, parents appeared to worry more about their boys than about their girls, especially when they reached early adolescence. There was experimental evidence that New Hope's child care supports were more likely to be used for boys than for girls. Mothers often said that their boys were vulnerable, and they used any resources they had to counteract negative influences. As one mother said, "It's different for girls. For boys, it's dangerous. [Gangs are] full of older men who want these young ones to do their dirty work. And they'll buy them things and give them money." (p. 12)

New Hope boys were more likely than girls to be in organized after-school programs where they received help with homework and had opportunities for recreation (Duncan et al., 2007). The larger impact on boys may be explained by the fact that from the parents' perspectives, boys had much more to gain from intervention than girls.

In addition, there are several other examples from the literature on using ethnographic, interview, and other mixed-methods data to investigate why effects occur and can change across participants (Bernheimer, Weisner, & Lowe, 2003; Datta, 2005; Lowe & Weisner, 2004). Clampet-Lundquist et al. (2006) used MTO data to assess why behavior and mental health measures improved for girls but not for boys who relocated into higher income neighborhoods. This article provided an excellent example of formulating four competing theoretical hypotheses that might explain these results and used multiple-methods data—including a new in-depth interview of teens—to test these hypotheses. Kling, Liebman, and Katz (2005) supplied a testimonial to the value of the in-depth interviews collected in MTO and illustrated how such data were used in the article (see p. 10 for a quotation from the introduction to the Kling et al., 2005, article).

- **Does Intervention X remain effective when different outcome measures are used?**

Include multiple quantitative outcome measures to assess different aspects of the desired outcomes (e.g., specialized outcome measures aligned with the purposes of Intervention X as well as more general measures such as standardized test scores).

Use case studies, interviews, and observations to detect unanticipated/unmeasured outcomes.

The experiences of Project STAR, New Hope, and particularly Moving to Opportunity, as well as of other RCTs with long-term follow-up studies, suggest that (a) the effect of social and educational interventions are unlikely to be confined to a single outcome measure or to a single generation, especially in the long term; (b) some outcomes are likely to be unpredictable and/or unanticipated,

especially in the long term; and (c) some of the primary measures often chosen by researchers can have small and/or null effects (e.g., achievement, labor force measures), while unanticipated measures, usually health and behavioral measures, can have large and significant effects.

It is also noteworthy that the importance of effects depends not only on their effect size but also on their contribution to long-term benefits. For instance, the Abecedarian and Perry Preschool experiments originally used IQ and achievement test scores to measure academic performance. Although the participants showed improvement on these measures, most of the benefits flowed either in the form of lower levels of grade retention and special education placement or changes in behavior that resulted in less involvement with the criminal justice system (Károly et al., 1998). Other important unanticipated effects included examples of generational effects like lower levels of addictive behavior and teen pregnancy (Károly et al., 1998; Károly, Kilburn, & Cannon, 2005; Masse & Barnett, 2002; Ramey et al., 2000; Reynolds, Temple, Robertson, & Mann, 2002; Reynolds et al., 2007; Schweinhart, 2004).

Failure to measure the full range of effects can result in significant underestimation of the benefits of an intervention. Although using multiple methods does not guarantee that all effects will be measured, these data provide the best opportunity to capture unanticipated outcomes and develop stronger theories that can better predict the full range of outcomes.

In Project STAR, the outcome measures used in the short term were standardized test scores, grade retention, and special education placement. These effects were certainly significant, but the effect sizes ranging from 0.2 to 0.3 generally would not be expected to have such large impacts on high school graduation or signing up for college entrance examinations. The achievement gains in elementary school translated into significantly higher secondary school graduation rates and increased levels of taking college entrance tests (Finn et al., 2001, 2005; Krueger & Whitmore, 2001).

New Hope's original objective was to move families out of poverty through more stable and higher paying jobs and better health care. However, a new set of outcome measures were introduced when the supplemental parent-child study started during the second year of the experiment. This study assessed, among other measures, changes in parenting practices, children's school performance (as rated by teachers and through standardized testing), and children's behavior (Duncan et al., 2007; Huston et al., 2001).

The original objectives of MTO involved improvements in income and labor force behavior and children's performance in school. In general, the experiment showed no significant

effects for any of these measures (Goering & Feins, 2003; Katz, Kling, & Liebman, 2001; Ladd & Ludwig, 1997; Rosenbaum & Harris, 2001; Sanbonmatsu et al., 2006). However, in-depth interviews alerted researchers to refocus their analyses on mental health, criminal behavior, and children's conduct measures, which showed large effects (Browning & Cagney, 2003; Kling, Ludwig, & Katz, 2005; Kling et al., 2007; Leventhal & Brooks-Gunn, 2003b; Ludwig, Duncan, & Hirschfield, 2001).

- **Are all of the components of Intervention X necessary for it to work, or are some unnecessary? Are some needed components missing?**

Plan to measure the various intervention components; build in case studies to learn which components mattered to different subjects and to generate hypotheses about other components that might have made Intervention X more effective.

Determining whether all components are necessary and whether some components mattered more to some participants than to others is critical, since simplifying and targeting an intervention can significantly reduce costs. For instance, Project STAR analyses asked whether 4 years of smaller classes were required to affect achievement or whether similar effects would occur with fewer years of intervention. Smaller class sizes are costly, and if each year did not make contributions to the effect, significant cost savings would be possible.

Hanushek (1999) suggested that most of the achievement effects occurred in the first year. However, Krueger (1999), Finn et al. (2001), and Hedges, Konstantopoulos, and Nye (2001) suggested that 3 or 4 years of small classes are required for sustained, long-term achievement effects. It would have been desirable in Project STAR to have systematically varied the class size rather than to have aimed for reductions of 8 pupils per class across all schools. Perhaps most of the achievement gains were due to reductions of 6 rather than 8 pupils. If so, then in future class size reduction initiatives, significant cost savings would be possible.

Project STAR analyses showed higher effects on achievement from smaller classes through Grade 3 for minority and disadvantaged students (Finn & Achilles, 1999; Krueger, 1999). However, Finn et al. (2001), Nye, Hedges, and Konstantopoulos (2000a, 2002, 2004) and Krueger and Whitmore (2002) suggested that the effect sizes declined somewhat by eighth grade and that the larger effects for minority and disadvantaged students were mixed at eighth grade. The studies by Krueger and Whitmore (2001) and Finn et al. (2001) showed significantly larger

effects on high school graduation and college entrance test taking for minority and disadvantaged students who participated in the experimental small classes.

Finishing high school requires more than direct cognitive gains. Other developmental skills such as social skills and behavioral and emotional skills play important roles in completing education and in labor force success. This suggests that Project STAR may have affected children's social, behavioral, or emotional trajectories as well as their cognitive trajectories. Finn and Achilles (1999) argued that improved classroom behavior may partially account for achievement gains, but ideally a wider range of developmental measures would have been included in kindergarten through third grade and in the longer term follow-ups.

This pattern of larger long-term effects for measures other than direct achievement measures seems to be emerging as a consistent finding from several early interventions of long duration. For instance, the Perry Preschool and Abecedarian projects had significant effects on many behavioral measures such as reduced involvement with the criminal justice system, even though achievement gains leveled off or declined in the longer term (Karoly et al., 1998, 2005; Masse & Barnett, 2002; Ramey et al., 2000; Reynolds et al., 2002, 2007; Schweinhart, 2004).

New Hope found that the utilization of three key benefits diverged widely across participants and that the largest effects were for those participants whose particular life circumstances "fit" the particular menu of offered benefits. For instance, the cost of day care was a prime benefit for mothers with children—especially those with multiple children—so men and women without children could not take advantage of this lucrative benefit. Also, many children had significant health or disability problems, and the health insurance benefit provided coverage for these kinds of issues. Perhaps one of the major lessons arising from New Hope is the need to characterize the diverse needs of a population before designing the benefit package and offering a wider and more flexible menu that would address a broader range of issues (Duncan et al., 2007). As an example, about 20% of participants had more severe problems linked to drugs, alcohol, or domestic abuse. For these families, other interventions were needed before they could take advantage of the New Hope benefits (Duncan et al., 2007, chaps. 3–4).

- **Are the treatment effects sustained over time?**

Plan extended follow-ups, particularly of treatment group members, using both quantitative and qualitative data (e.g., achievement data, case studies, interviews).

Project STAR has followed participants through high school. Achievement data were collected at eighth grade. Then, at the end of high school, data were collected on how many college entrance tests were taken as well as on high school graduation rates. The analyses suggest that the size of achievement effects declined somewhat at eighth grade, and earlier differential effects for minority and disadvantaged effects were mixed (Krueger & Whitmore, 2002; Nye et al., 2000a, 2002, 2004a). However, the studies by Finn et al. (2001) and Krueger and Whitmore (2001) showed large effects on high school graduation and levels of taking college entrance tests, with much larger effects for minority and disadvantaged students.

In general, the long-term effects of New Hope and Moving to Opportunity tended to be small or nonexistent for direct labor force measures such as income and employment. However, there were somewhat larger effects for selected behavioral and school performance of participants' children by gender and school subject (although MTO did not directly measure the effects on children's achievement) (Kling et al., 2007; Sanbonmatsu et al., 2006). Also, adult and children's mental health measures showed positive long-term effects (Kling et al., 2007; Leventhal & Brooks-Gunn, 2003b).

New Hope followed up with interviews 2 years and 5 years after the experiment ended to determine long-term effects on participants and their children. Duncan et al. (2007, chap. 11) reported that the larger and most persistent effects were on the children—particularly the achievement and behavior of the boys. This is another example of the importance of measuring generational effects. It is possible that smaller but persisting improvements in the lives of parents, particularly single mothers, can generate larger and longer lasting effects in the next generation. Thus far, the long-term generational effects on the children of the individuals who participated in interventions like Perry Preschool or Abecedarian have not been measured.

Clampet-Lundquist et al. (2006) used data from follow-up interviews with MTO participants 4–7 years after project initiation to explore differential effects on behavior changes in boys and girls. They also carried out an additional data collection with a subsample of teens focusing on a theory-based set of hypotheses directed at explaining gender differences in outcomes. Clampet-Lundquist et al.'s article is an excellent example of adding a multiple-methods data collection 5 years into the

experiment to further test hypotheses generated by the original data collection. The original analysis suggested no differences in risk behavior for boys in the treatment and control groups, but in fact, girls exhibited better mental health and lower risk behavior. The added in-depth interviews with teens, together with the original follow-up data, suggested specific viable explanations for the differences in outcomes between genders.

Considerations for External Validity

• How do contextual effect factors affect the impact of Intervention X?

Use case studies, administrative data, interviews, and observations to document contextual factors (e.g., local policy environment, resources, cultural concerns, history) and how they might interact with Intervention X.

Researchers face substantial obstacles in translating successful small-scale experiments into successful large-scale programs (see Schneider & McDonald, 2007, Vols. 1–2, for an excellent review of research on scaling up). Experiments only provide results with predictive validity if the conditions and context of the experiment can be duplicated

in other settings. However, experimental conditions can never be perfectly duplicated. In fact, the conditions in experimentation (a high degree of control of conditions, personnel selected by researchers, etc.) that are necessary to make experiments successful from a scientific perspective often guarantee smaller effects in scaled-up real-world settings. In addition, contextual effects seem to be ubiquitous in social and educational interventions. One of the key advantages of multiple methods in RCTs is to provide information that can better predict how results might change in different contexts, conditions, and scales.

Although Project STAR was carried out on a large scale in real-world conditions, results from the experiment cannot be assumed to transfer to different populations in different schools under different conditions. The relatively smooth implementation of STAR in 79 Tennessee schools stands in stark contrast to the statewide class size reductions in California that were plagued by teacher shortages and limited space (Bohrnstedt & Stecher, 2002).

In Project STAR, each of 79 schools represented a separate experiment because each school included at least one randomly assigned small class, a large class, and a large class with teacher aides. The context, however,

Perhaps one of the major lessons arising from New Hope is the need to characterize the diverse needs of a population before designing the benefit package and offering a wider and more flexible menu that would address a broader range of issues.

was different across schools, enabling the researchers to explore contextual effects. Teachers were also randomly assigned to classrooms to enable research on teacher effects in classrooms. Two simple and important examples of contextual effects in STAR are that (a) minority and disadvantaged students experienced higher achievement effects and (b) all students in small classes for 1–2 years, rather than for 3–4 years, experienced no sustained effects (Finn & Achilles, 1999). Thus, STAR effects would be predicted to vary by student characteristics and by the number of years of small classes between kindergarten and third grade. However, STAR data have also been used to explore more complex contextual effects.

For instance, Nye, Konstantopoulos, and Hedges (2004) and Peevely, Hedges, and Nye (2005) explored the effect on achievement gains of teacher experience and salary and classroom composition. They suggested that teacher experience effects are larger in math than in reading and that lower SES classrooms have larger variance in score gains due to teachers than do higher SES classrooms. Dee (2004) suggested that students who have same-race teachers have higher score gains. Nye, Hedges, Konstantopoulos (2000b) analyzed contextual effects of class composition and school location and concluded that class composition and location do not change effect size significantly.

New Hope was a small-scale experiment conducted in an economic and welfare policy environment (Wisconsin) that was not typical of other states. State economic conditions produced a labor market that was quite favorable to finding and maintaining employment. For instance, Duncan et al. (2007) pointed out that the control group also generated substantial gains in employment and income, making the explanation of treatment effects challenging. However, the New Hope participants made even larger gains than did the controls in employment and income. Wisconsin was also at the forefront of welfare reform, making such generalizations of New Hope to other states problematic. However, the rich multiple-methods data enabled researchers to do more than speculate on how a program like New Hope might be redesigned and scaled up nationally (Duncan et al., chaps. 6–7).

Interesting neighborhood contextual factors were identified in MTO. Turney et al. (2006), using multiple-methods data from in-depth interviews with participants, identified barriers to employment. These barriers may explain how, in spite of relocation to better neighborhoods

in Baltimore, participants did not experience employment and earnings effects. Identifying such barriers helps to delineate conditions in other cities that might be needed for earning and employment effects to occur.

- **How close are the measured outcomes to outcomes of interest?**

When designing the study, interview key stakeholders to determine the relevance/appropriateness of the outcome measures proposed for the study.

The ultimate stakeholder in social and educational experimentation is the American taxpayer. For these stakeholders, a commonly used criterion is that the long-term benefit to society (measured in monetary terms) must at least exceed societal costs and, it is hoped, have a rate of return that justifies government borrowing. However, Karoly and Bigelow (2005) suggested that many government programs would not meet this criterion, and they offered an alternative—that

a particular early childhood program only needs to have a benefit–cost ratio higher than other government programs.

Near-term stakeholders in Project STAR included the Tennessee legislature, Tennessee teachers, and parents of K–3 students. The Tennessee legislature authorized Project STAR, which clearly indicated that raising achievement was a prime objective. But the impetus for smaller class sizes at the policy level was due to pressure from parent and teacher groups. Stakeholders were directly involved in specifying the intervention, as well as setting objectives (see Ritter & Boruch, 1999, for a history of STAR). However, while the initial focus was on immediate achievement, the most important outcomes for society showed up in the long-term follow-ups, where researchers were able to register substantial gains in high school graduation and college entrance behavior.

The evolution of New Hope had a long history, dating from over 15 years before the experiment was initiated (Duncan et al., 2007, chaps 1–2). A clear objective was to provide evidence for how to design an improved welfare system. Besides federal and state policymakers, stakeholders included the business community and welfare participants themselves. The broadened objectives in New Hope transformed from a primary emphasis on adult labor force outcomes to a strong emphasis on behavioral outcomes for both parents and children, as well as schooling outcomes for boys. For the children’s outcomes, the dual emphasis on school achievement and behavior both in and out of school

One of the key advantages of multiple methods in RCTs is to provide information that can better predict how results might change in different contexts, conditions, and scales.

proved to be an important and persisting generational result of the New Hope experiment.

The federal Department of Housing and Urban Development was the sponsor of MTO. Its clear purpose was to determine how important neighborhoods were to adult outcomes, so that better housing policies could be promoted. However, MTO expanded from a primary emphasis on improvements in labor market measures, which showed no statistically significant effects, to measures of mental health, parenting, and children's behavior (Kling et al., 2007; Sanbonmatsu et al., 2006). No achievement effects were found, but some effects on children's behavior and on the mental health of adults and children were significant (Kling et al., 2007; Leventhal & Brooks-Gunn, 2003b). One of the key findings from MTO was that nonexperimental research had overestimated neighborhood effects. In fact, neighborhood effects were smaller, involved a broader range of outcomes, and were more complex than previously thought (see, e.g., Booth & Crouter, 2001; Duncan & Raudenbush, 1999, 2001; Kling et al., 2007; Leventhal & Brooks-Gunn, 2003a; Turney et al., 2006).

- **How would resource constraints affect the institutionalization of Intervention X if it were found to be effective?**

Build collection of cost data into the study and conduct cost, cost-effectiveness, and cost-benefit analyses.

Levin and McEwan (2000, 2002) provided a good introduction to conducting either cost-effectiveness or cost-benefit analyses and distinguishing between them. A cost-effectiveness analysis compares the proposed intervention to alternate interventions that focus on a single common outcome (e.g., higher achievement). Cost-benefit analyses use a single intervention to see if long-term monetary benefits exceed costs from all outcomes.

There are several good examples of conducting analyses incorporating costs and benefits. Karoly et al. (1998) compared the costs and benefits of the Perry Preschool programs and a nurse visiting program. Karoly and Bigelow (2005) analyzed the costs and benefits of universal preschool in California. Masse and Barnett (2002) estimated the costs and benefits from the Abecedarian Project. Reynolds et al. (2002) performed a cost-benefit analysis of a Chicago early childhood intervention. Lynch (2004) summarized several cost-benefit analyses of early childhood programs. Grissmer's (2002) study contained a cost-effectiveness analysis of four options for improving achievement.

Such analyses cannot usually yield reliable predictions for scaled-up programs in different contexts. Although it is usually assumed that effect sizes can change with context,

costs as well as effects are sensitive to context, so costs measured in experimental settings may change dramatically in large-scale settings. Brewer, Krop, Gill, and Reinhardt (1999) illustrated the variance in the cost of class size reductions depending on location (cost-of-living differences), the specific rules used to implement such reductions, the availability of space, the hiring practices of teachers, the pay scales of teachers, and the characteristics of the students targeted for smaller class size. For instance, implementing class size reductions in inner cities carries higher space and teacher costs but also leads to larger effect sizes.

Moreover, scaling up small-scale interventions to large-scale public sector programs can carry several additional cost considerations. For instance, because such programs depend on forming a successful political coalition for passage, powerful stakeholders can lobby for wider eligibility in experimental groups, which can lead to higher average costs and lower effects. Gordon (2004) suggested that federal allocations of Title I funding (for low-income students) to local governments get partially diverted by local governments for alternative noneducational uses. Finally, programs are rarely fully funded. Duncan et al. (2007, chap. 8) provided analyses of cost-benefits for New Hope and a discussion of the implications for expanding New Hope to a larger state or national program.

- **How do the details of the intervention and the controls imposed by the study design differ from the real-world conditions under which Intervention X might be implemented?**

Collect and report descriptive data that will allow policymakers to assess the similarity of the sample population and setting to those in other situations to which they might want to generalize results.

Project STAR was a large-scale intervention involving over 12,000 students in Grades K–3 in mostly large suburban and urban schools in Tennessee. Tennessee children in STAR included disproportional numbers of minority and disadvantaged participants compared to all Tennessee students, and Tennessee students are disproportionately more disadvantaged than U.S. students (Grissmer, 1999). The larger effects for minority and disadvantaged children mean that average effects can change markedly as the composition of students changes across states. Tennessee implemented Project STAR almost entirely in suburban and inner-city schools. The costs and effects may change for rural schools (where recruiting teachers may be more difficult) or in states with higher or lower costs of living than in Tennessee. Tennessee also utilized experienced teachers rather than hiring new teachers, although the experiment

did not provide for specific preparation or instruction directed at teachers working with smaller classes.

Researchers learned many lessons from Project STAR that could guide policy and implementation in other states. Three important lessons were that (a) 3–4 years of small classes were needed for long-term effects, (b) effects were much larger for minority and disadvantaged children, and (c) the teachers in Project STAR were not newly recruited but were drawn from the pool of existing experienced teachers.

Project STAR did spur class size reductions in many states beginning in the 1990s, which extended into the next decade. In general, these reductions were more often directed to schools and districts with larger proportions of minority and disadvantaged students. Grissmer, Flanagan, Kawata, and Williamson (2000) and Grissmer and Flanagan (2006) used state National Assessment of Educational Progress scores to assess the effects of class size reductions and other initiatives across states from 1990 to 2005. They concluded that the average effect of such class size reductions is consistent with the results from Project STAR.

California was the notable exception in successfully building off of Project STAR. California mandated sizable class size reductions statewide for all students in Grades K–3 beginning in 1996 (Bohrnstedt & Stecher, 2002).

The short time between mandate and implementation and the fact that all children in Grades K–3 were eventually involved left school districts throughout the state unprepared for hiring the necessary additional teachers and finding the needed classroom space. Unlike Tennessee, California did not prudently phase in the program beginning in kindergarten so that all children would experience 3–4 years of smaller classes, nor did they target the intervention to minority and disadvantaged children. This failure to phase in the program more slowly generated shortages in teachers and classroom space. The initiative led to smaller short-term effects due to the inclusion of more advantaged children and children receiving only 1–2 years of smaller classes.

Such haste also failed to ensure that sufficient data were collected and available to provide unbiased measurements of short-term effects. For instance, comparable test scores were not available for the years prior to the experiment, so the evaluation lacked a critical source of comparative

evidence. An unintended consequence of the rush to small classes and lack of targeting was that many better quality teachers in central city schools left for the suddenly available jobs in suburban schools. Inner-city schools not only had to recruit teachers needed to reduce classes but also had to fill additional vacancies caused by those moving to suburban schools. These changes meant that it was impossible to predict California effects from Project STAR effects, or to provide unbiased measurements of the short-term effects of small classes in California.

New Hope was a small-scale program with volunteer participants. Thus, expansion to a large-scale program would mean incorporating those populations that did not volunteer, as well as all the cost and effectiveness issues associated with scaling up from small experimental programs (Quint, Bloom, Black, & Stephens, 2005; Schneider & McDonald, 2007). Duncan et al. (2007, chap. 8) provided analyses of cost–benefits for New Hope and a discussion of the implications and uncertainty involved in scaling New Hope to a state or national program.

An unintended consequence of the rush to small classes and lack of targeting was that many better quality teachers in central city schools left for the suddenly available jobs in suburban schools. Inner-city schools not only had to recruit teachers needed to reduce classes but also had to fill additional vacancies caused by those moving to suburban schools. These changes meant that it was impossible to predict California effects from Project STAR effects, or to provide unbiased measurements of the short-term effects of small classes in California.

BOX 4: ROLE OF MULTIPLE METHODS IN PROVIDING A DEEPER UNDERSTANDING OF STUDY FINDINGS

- In addition to using quantitative measures to assess outcomes, use data from case studies, interviews, surveys, and/or observations to interpret the observed outcomes (e.g., how intervention was experienced and responded to by subjects in differing circumstances).

Perhaps the main reason why collecting multiple-methods data in RCTs is necessary is that each participant in any social science RCT has a unique genetic and developmental history, and many of the forces shaping development involve gene–environment interactions (Rutter, 2002). The a priori expectation should therefore be of differential effect size across participants. If theories are to be ultimately successful in predicting behavior, they must in some way take account of and incorporate this wide diversity inherent in study subjects. Data from multiple methods can be seen as a start to understanding this uniqueness and diversity and exploring ways of identifying groups with similar enough paths and responses that enable more efficient targeting and more accurate predictions. These individual paths and responses to interventions can only be captured by multiple-methods data.

New Hope collected an extremely rich set of multiple-methods data, and researchers have used these data to try to understand several emerging issues. These issues include why control participants who did not receive New

Hope benefits made large labor market gains; why the incrementally larger gains made by participants receiving New Hope benefits were statistically significant but of modest size; and why many participants eligible for New Hope benefits did not utilize their benefits or utilized them only sporadically. In addition, the differential effects for boys on behavior and achievement were puzzling.

The 12 analyses contained in Yoshikawa et al.'s (2006) study, using New Hope data, provide outstanding examples of (a) how multiple-methods data can address unexpected utilization and results; (b) how to make these types of interventions and similar policies more effective; and (c) in general, how results are dependent on context. Huston et al. (2001) and Duncan et al. (2007) provided examples of incorporating the analyses of multiple-methods data into academic and policy publications. A volume by Weisner (2005) contains 12 chapters that illustrate the value of multiple-methods data to address research questions not necessarily embedded in RCTs. As such, it provides material that is helpful in learning how such methods have been used across different research areas, what kinds of methods have been employed, and how these data have contributed to testing hypotheses and theories about why and how behavioral effects occur.

Perhaps more important, these analyses illustrate the complexity and uniqueness of the lives of poor, working mothers and why it is difficult to design interventions and

These individual paths and responses to interventions can only be captured by multiple-methods data.

policies that could have a great impact on large numbers of such women. For instance, Lowe, Weisner, and Geis (2003) provided a picture of the challenge of finding day care for the children of poor, working mothers and the problem with “one size fits all” benefit packages. An important lesson drawn from New Hope is the need for more extensive and flexible benefit options to address the complexity and diversity in the lives of poor working mothers and their children (Duncan et al., 2007).

In Project STAR, Finn et al. (2003) used observational data, teacher surveys, and interviews to address the question of why small classes work. Gerber et al. (2001) also employed teacher and teacher aide logs, surveys, and interviews to address why adding teacher aides to classrooms did not have large effects.

Clampet-Lundquist et al. (2006) used follow-up interview data in Moving to Opportunity to develop and test hypotheses as to why the experience of changing neighborhoods was different for girls than for boys and why girls fared better than boys in new neighborhoods. Turney et al. (2006) used interview data from MTO participants to explain why moving to higher income neighborhoods had no effects on employment, income, or welfare utilization.

- Use data from case studies and/or interviews to illustrate findings in a compelling manner.

There are two primary audiences for the findings of RCTs: researchers and policymakers. Researchers tend to be concerned about whether effects occur and how big the effects are when compared with alternatives. However, to develop theories and improve research designs, researchers will increasingly have to address questions such as why effects occur. The constraints on normal academic publications often preclude the longer page length required to address such questions. In complex RCTs with extensive multiple methods, results need to be communicated through edited books or longer summary publications. Yoshikawa et al. (2006) provided an indispensable resource for researchers designing mixed-methods RCTs and communicating their results to other researchers. This volume is entirely directed toward using multiple-methods data to address key issues in explaining the pattern of New Hope results, particularly the differential effects across adults and children. Weisner (2005) provided examples for researchers from a wider range of studies.

Policymakers also need publications that illustrate findings in a compelling manner and address concerns

specific to their roles. Policymakers must develop political support for any new program; thus, both legislators and the public need to be convinced of the merits of a program. Policymakers will face questions from legislators and the public about why a particular program will work, how it can be targeted to achieve larger effects, and what it will cost. In such contexts, relating stories about how individual participants responded, why it worked for some participants and not others, and how lives were changed can be effective methods of communication for policymakers.

Policymakers also need research translated into “readable” and compelling prose. Duncan et al. (2007) provided an outstanding example of communicating the results of mixed-methods analyses to a more general audience, including policymakers. This volume wraps the basic results of the intervention around a compelling narrative that illustrates how and why the intervention worked in individual cases, how one could change the design to obtain a more effective and efficient intervention, how to set eligibility rules, and in what contexts this intervention might or might not work, as well as the potential risks of moving to a large-scale program.

In complex RCTs with extensive multiple methods, results need to be communicated through edited books or longer summary publications.

- Examine quantitative and qualitative results to determine whether additional hypotheses (e.g., about additional outcomes, modifications to the intervention) might be pursued in subsequent studies or different stages of the current RCT.

RCTs are usually envisioned as having a fairly unchangeable design that includes well-defined planning, implementation, and analysis stages. This format is dictated largely by the federal proposal process. Such a research design is generated in accordance with preexisting theories and a fixed set of outcome measures, yet they are problematic when RCTs show unexpected results or have unexpected outcomes. Although follow-up RCTs might be designed to address these issues, it may be more efficient and timely to use resources to expand data collection either during the RCT or in longer term follow-up. In fact, multiple-methods RCTs are directed toward answering a more complex set of questions than a black box RCT, and the chances of unexpected results are correspondingly higher. Puzzling differential outcomes and results require new theories. This kind of research likely requires a more flexible and opportunistic research funding process that is able to respond to unexpected findings.

Both New Hope and MTO might be described as having an evolving and opportunistic research strategy that responded to emerging research findings with additional and expanded multiple-methods data collections. These data collections were targeted toward identifying a wider set of outcomes, framing and testing new emerging hypotheses, and providing explanations of unexpected effects.

MTO was the first large-scale RCT designed to explore the effect of neighborhoods on adults' economic outcomes and children's schooling outcomes by randomly assigning differential access to higher income neighborhoods. However, no significant effects were found on adult employment and income across the five locations, and no effects were found on children's schooling outcomes (Kling et al., 2007; Sanbonmatsu et al., 2006). The null effects from these primary measures resulted in a redirection of the study to determine if there were effects on the mental and physical health of adults and children and on the incidence of risky behavior among youth.

These data came from a follow-up survey, conducted about 7 years after the initiation of MTO, of each participating adult and up to two children per household that included a much wider set of outcome measures and also explored possible explanations for the null effects on economic outcomes. Kling et al. (2007) provided a summary of these results that suggests large and significant positive effects on adult mental health measures but no effects for physical health measures. Young females experienced positive effects on physical and mental health and lower incidence of risky behavior. However, male youth showed no effects or offsetting negative effects on each of these measures.

Kling et al. (2007) also provided some hypotheses to explain the null effects on adult economic measures. Turney et al. (2006) used the long-term follow-up and an additional in-depth interview with 67 participants to explain why moving to higher income neighborhoods had no effects on employment, income, or welfare utilization. Clampet-Lundquist et al. (2006) also used this follow-up interview and an additional in-depth interview of 86 teens to develop and test hypotheses as to why girls fared better than boys when parents moved to a higher income neighborhood. Multiple-methods data collections were critical in instigating the redirection of data gathering and interpretation.

An article by Kling, Liebman, and Katz (2005) provided unusual and compelling testimony that highlighted the value of qualitative, in-depth interviews and eventually changed the research strategy for the MTO project (see p. 10 for commentary from the researchers on how these interviews helped to identify mechanisms driving the outcomes and offered insights into interpreting results). Kling, Liebman, and Katz (2005) should be read by those

researchers and policymakers who question the value of embedding multiple methods in RCTs

The original focus of New Hope was also on adult economic measures: increased employment and wages and less welfare dependency. However, partly because of funding opportunities and lower utilization of benefits by adults without children, the experiment increasingly focused on outcomes for mothers—about 71% of the total sample. Researchers introduced new data collection measures that focused on a wider set of adult health, parenting, and other behavioral measures, as well as measures to link performance in schools to health and behavioral measures of the participants' children (Duncan et al., 2007; Huston et al., 2001).

Bos et al. (2007) provided an example of employing qualitative data that, like the MTO in-depth interviews, highlights the fear associated with threats to safety in the lives of poor families—especially single-parent families (see quote from this study on p. 19). New Hope boys were more likely than girls to be in organized after-school programs where they received help with homework and had opportunities for recreation (Duncan et al., 2007). The larger impact on boys may be explained by the fact that from parents' perspectives, boys had much more to gain from the intervention than did girls.

BOX 5: EXPLORING IMPLICATIONS FOR RESEARCH, POLICY, AND NEXT STEPS

Relate outcomes of the current study to findings from prior research

- Do the results of this RCT confirm or contradict results from other studies of similar interventions? Consider results from RCTs and other types of studies (e.g., quasi-experimental, correlational, and ethnographic).
- What factors may account for differences in results between this RCT and previous studies? Take account of variations in study design, characteristics of participants, outcome measures, settings, times, and fidelity of implementation.

Ideally, a literature review (as outlined in the What Is Known From Previous Research section; see pp. 7–8) is available and can serve as the basis for integrating the new results with outcomes from previous studies. A major motivation for conducting a multiple-methods RCT is to make it more likely that future literature reviews will generate a scientific and/or policy consensus. The integration of the present results with the previous literature review should thus focus on if and how the present results settle disparities existing in the previous literature and/or raise new questions and issues that must be the subject of future research.

Multiple-methods RCTs are likely to emerge as the primary method to explain both direct and indirect disparities in past measurements. The direct contribution arises from being able to measure contextual effects, measure differential effects across participants, and eliminate many nonexperimental sources of bias. Each of these contributions will help reconcile existing disparities in the literature. The indirect contribution will come from building stronger theories. Theories are successful only to the extent that they can accurately predict the results of many measurements.

In this context, it is important to realize that consensus usually emerges only when the disparate results from previous research can be reasonably reconciled or explained by viable theories. Consensus is generally not achieved by any single gold standard experiment alone. Project STAR probably comes the closest to a gold standard intervention. But Project STAR also provided many explanations for the disparity in previous measurements by showing differential effects (larger effects for minority and disadvantaged students), necessary components (3–4 years required for sustained effects), the presence of strong pressure for interference and selectivity (pupils assigned to large classes often made their way to smaller classes), and absence of strong teacher and school contextual effects. All of these helped to explain some of the disparities in

Multiple-methods RCTs are likely to emerge as the primary method to explain both direct and indirect disparities in past measurements.

previous measurements. Ehrenberg, Brewer, Gamoran, and Willms (2001) and Grissmer (1999) provided examples of integrating Project STAR results with previous class size measurements.

This integration of new multiple-methods RCT results with previous literature requires a thorough knowledge of the strengths and weaknesses of analysis using nonexperimental data, data from natural experiments, quasi-experiments, and experimental measurements. Three important explanations for why previous results may differ include (a) measurement bias, (b) the presence of contextual effects, or (c) differences in the characteristics of the population studied. Since the potential for bias usually differs by research method, one strategy is to group them into experimental, quasi-experimental, natural experiments, and nonexperimental methods.

Webbink (2005) provided an example of this type of review. However, within each of these categories, there is usually wide variation in quality, so that simple categorization often can be misleading. Because of this, the review must also assess the quality of studies within each category.

A substantial literature helps with conducting such a critique. Duncan and Gibson-Davis (2006), Duncan and Magnusson (2003), and Duncan et al. (2004) discussed the way in which experimental methods address measurement bias issues in nonexperimental data. They also argued for the potential of natural experiments. Cronbach and Shapiro (1982) and Heckman and Smith (1995) provided critiques of experimental studies. Cook et al. (2005) compared and contrasted experimental and quasi-experimental results. O'Connor (2003) and Rosenzweig and Wolpin (2000) provided a developmental psychology and economics perspective on both advantages and disadvantages inherent in natural experiments.

Other useful resources that summarize and interpret schooling effects for adolescents from eight experiments in welfare reform policy include Gennetian et al. (2004). Leventhal and Brooks-Gunn (2003a), Oakes (2004), and Kling et al. (2007) discussed the difficult issues involved in measuring neighborhood effects and contrasted findings from experimental and nonexperimental studies. Krueger (1999) and Wilde and Hollister (2007) provided some more direct comparison of experimental and nonexperimental results using Project STAR data.

Respond to finding positive effects

- Consider policy implications.
- Decide whether further scale-up is needed. If so, decide whether replicate studies are needed before going to scale and what the cost, cost-effectiveness, and cost-benefit of going to scale would be.

The value of positive results from RCTs that do not collect multiple-methods data can be degraded significantly if (a) results cannot be generalized to different populations, (b) contextual effects cannot be identified, (c) weaknesses in the design of the intervention cannot be identified and improvements suggested, and (e) the issues in scaling up to larger programs cannot be addressed. Only multiple-

methods data can be used to address these issues, and the next steps after garnering positive results is to carefully assess these issues using the multiple-methods data. Each of these four issues should be addressed in analyses and publications before proceeding to make decisions on next steps.

Perhaps more important, a publication needs to address the implications of the results on an understanding of why and how researchers achieved the effects.

New Hope provides an outstanding example of the additional work and documentation required after obtaining positive effects for both adults and children. Duncan et al. (2007) provided a summary of the many analyses undertaken and targeted policymakers who might be interested in undertaking a statewide or national program. They were able to address these policy issues only because of New Hope's comprehensive multiple-methods data collections.

Probably the most difficult challenge is making predictions of costs and effects for a scaled-up program from data/results originally collected in small-scale programs. Schneider and McDonald (2007, Vols. 1–2) provided a comprehensive assessment of the issues involved in scaling-up education programs. In general, small-scale experiments should not be used as the basis for a major program implementation unless compelling cases can be made that effects and costs will not change in different contexts or at different scales.

In education, the Success for All intervention comes with an interesting history of moving from smaller to larger scale and evaluating the results experimentally (Borman & Hewes, 2002; Borman et al., 2005, 2007). Efforts to gradually increase implementation of Success for All across

In general, small-scale experiments should not be used as the basis for a major program implementation unless compelling cases can be made that effects and costs will not change in different contexts or at different scales.

different types of schools allowed many of the contextual hypotheses and scaling issues to be tested. Slavin (2002, 2008), Chatterji (2005, 2008), and Briggs (2008) also brought useful perspectives to the question of synthesizing research evidence and deciding when programs should be recommended for wider implementation.

The results from Project STAR had a major impact on class size policies throughout the nation from 1995 to 2007. This influence was partly due to its experimental design, large sample, and the transparency of its findings to policymakers. It also benefited from widespread public support and belief in smaller class sizes, especially when coupled with expanding state revenues. With the exception of California, states implemented smaller classes in a way that took account of Project STAR's findings. Smaller classes were often targeted to minority and disadvantaged children, and reductions were usually made for 3–4 years in early grades.

Project STAR benefited from not facing many of the scale-up issues inherent in other educational interventions. Project STAR was already operating at a large scale—in 79 schools. Implementation only required finding additional teachers and more classroom space. In general, outside California, implementation was gradual and targeted enough to allow for careful identification of teachers and space. Scale-up was also easier because class size effects in Tennessee were achieved with no additional teacher training. Providing quality training is often a key issue in scale-up. However, it is also possible that teacher training could have enhanced the STAR effects and could be the focus of additional research.

Respond to finding marginal or no effects

- Examine the implementation data.
- Examine the design; check the analyses.
- *If marginal effects are found:* Design new or additional studies to clarify results or abandon effort.
- *If no effects are found:* Rethink theoretical model; plan for a new study or abandon effort.

Null or marginal findings in experiments are often more important in making scientific progress than are findings of large effect sizes. This can be particularly true when null effects are found where current theories and understanding

would have predicted large effects in a given RCT. Such null effects directly undermine current theories and understanding and present an opportunity to develop new theories and discard old ones. One of the most important experiments in physics was the Michelson and Morley (1887) experiment that measured whether light propagating at right angles held the same speed. The null result paved the transition from Newtonian mechanics to special relativity. Multiple-methods data are crucial in helping to reject a current theory and moving to a new theory by

Multiple-methods data are crucial in helping to reject a current theory and moving to a new theory by generating an understanding of why assumptions in the old theory are untenable and what alternative theory might better explain the new results.

generating an understanding of why assumptions in the old theory are untenable and what alternative theory might better explain the new results. It is also important to publish null results in the literature because they provide crucial information for individuals involved in theory building. The current bias toward publication of studies with significant effects can

meaningfully impair the work of theory development.

Moving from public housing to a higher income neighborhood in MTO was hypothesized to improve adult job opportunities, employability and income, and children's schooling outcomes due to the accessibility of better schools and achieving better parent outcomes. However, researchers found no significant effects on adult employment and income of children's schooling outcomes across the five locations in the experiment (Kling et al., 2007; Sanbonmatsu et al., 2006). These null effects focused research and additional data collections on teasing out the flaws in the theories that predicted positive effects (Kling, Liebman, & Katz, 2005, 2007; Sanbonmatsu et al., 2006; Turney et al., 2006). For instance, Turney et al. conducted in-depth interviews with 67 participants in Baltimore to explore why the economic outcomes were insignificant:

The voucher group did not experience employment or earnings gains in part because of human capital barriers that existed prior to moving to a low-poverty neighborhood. In addition, employed respondents in all groups were heavily concentrated in retail and health care jobs. To secure or maintain employment, they relied heavily on a particular job search strategy—informal referrals from similarly skilled and credentialed acquaintances who already held jobs in these sectors. Though experimentals were more likely to have employed neighbors, few of their neighbors held jobs in these sectors and could not provide such referrals. Thus controls had an easier time garnering such referrals. (p. 137)

Project STAR found experimentally that having teacher aides in Grades K–3 had no consistent and significant

effect on achievement. Multiple-methods data indicated that aides spent only about 25–30% of their time on direct instructional tasks, with the remaining time spent on administrative or noninstructional interactions with students. However, even when aides did spend more time on instruction, this did not lead to effects on achievement. It was assumed that administrative and noninstructional work provided by teacher aides would allow teachers to be more effective, leading to achievement gains (Gerber et al., 2001). However, multiple-methods data suggested that teachers' perceptions of their ability to manage time, cope with student misbehavior, or engage students in the learning process was no different for teachers with or without aides (Gerber et al., 2001).

Managing aides demands additional teacher time that may reduce teacher productivity. The productivity of aides must therefore exceed the possible lost teacher productivity from managing teacher aides to register net gains. These data pointed to an emerging hypothesis that teacher aides had no specific training or educational background that would prepare them for the job. It is also possible that training might be needed to help teachers utilize aides effectively. Both of these hypotheses could be pursued through future research.

REFERENCES

- Angrist, J. D. (2004). American education research changes tack. *Oxford Review of Economic Policy*, 20, 198–212.
- Bernheimer, L. P., Weisner, T. S., & Lowe, E. D. (2003). Impacts of children with troubles on working poor families: Mixed method and experimental evidence. *Mental Retardation*, 41, 403–419. doi:10.1352/0047-6765(2003)41<403:IOCWTO>2.0.CO;2
- Blatchford, P. (2003). A systematic observational study of teachers' and pupils' behaviour in large and small classes. *Learning and Instruction*, 13, 569–595. doi:10.1016/S0959-4752(02)00043-9
- Blatchford, P. (2005). A multi-method approach to the study of school class size differences. *International Journal of Social Research Methodology*, 8, 195–205 doi:10.1080/13645570500154675
- Blatchford, P., Bassett, P., & Brown, P. (2005). Teachers' and pupils' behavior in large and small classes: A systematic observation study of pupils aged 10 and 11 years. *Journal of Educational Psychology*, 97, 454–467. doi:10.1037/0022-0663.97.3.454
- Blatchford, P., Bassett, P., Goldstein, H., & Martin, C. (2003). Are class size differences related to pupils' educational progress and classroom processes? Findings from the Institute of Education class size study of children aged 5–7 years. *British Educational Research Journal*, 29, 709–730. doi:10.1080/0141192032000133668
- Blatchford, P., Goldstein, H., & Mortimore, P. (1998). Research on class size effects: A critique of methods and a way forward. *International Journal of Educational Research*, 29, 691–710. doi:10.1016/S0883-0355(98)00058-5
- Blatchford, P., & Martin, C. (1998). The effects of class size on classroom processes: "It's a bit like a treadmill—Working hard and getting nowhere fast!" *British Journal of Educational Studies*, 46, 118–137. doi:10.1111/1467-8527.00074
- Bohrnstedt, G. W., & Stecher, B. M. (Eds.). (2002). *What we have learned about class size reduction in California?* (Capstone Report, Class Size Reduction [CSR] Research Consortium). Palo Alto, CA: California Department of Education.
- Bonesrønning, H. (2004). The determinants of parental effort in education production: Do parents respond to changes in class size? *Economics of Education Review*, 23, 1–9. doi:10.1016/S0272-7757(03)00046-3
- Booth, A., & Crouter, A. C. (Eds.). (2001). *Does it take a village? Community effects on children, adolescents and families*. Mahwah, NJ: Erlbaum. doi:10.1016/S0140-1971(02)00138-0
- Boozer, M. A., & Cacciola, S. E. (2001). *Inside the "black box" of Project STAR: Estimation of peer effects using experimental data* (Discussion Paper No. 832). New Haven, CT: Economic Growth Center. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=277009
- Borman, G. D. (2002). Experiments for educational evaluation and improvement. *Peabody Journal of Education*, 77, 7–27. doi:10.1207/S15327930PJE7704_2
- Borman, G. D., & Hewes, G. M. (2002). The long-term effects and cost-effectiveness of success for all. *Educational Evaluation and Policy Analysis*, 24, 243–266. doi:10.3102/01623737024004243

- Borman, G. D., Slavin, R. E., & Cheung, A. C. K., Chamberlain, A. M., Madden, N. A., & Chambers, B. (2005). The national randomized field trial of success for all: Second-year outcomes. *American Educational Research Journal*, 42, 673–696. doi:10.3102/00028312042004673
- Borman, G. D., Slavin, R. E., Cheung, A. C. K., Chamberlain, A. M., Madden, N. A., & Chambers, B. (2007). Final reading outcomes of the national randomized field trial of Success for All. *American Educational Research Journal*, 44, 701–731. doi:10.3102/0002831207306743
- Boruch, R. F. (1997). *Randomized experiments for planning and evaluation: A practical guide*. Thousand Oaks, CA: Sage. doi:10.1016/S0149-7189(97)00050-5
- Bos, H., Duncan, G. J., Gennetian, L. A., & Hill, H. D. (2007). *New Hope: Fulfilling America's promise to "make work pay"* (Discussion Paper No. 16). Washington, DC: Brookings Institution Press. Retrieved from http://www.brookings.edu/papers/2007/12_work_gennetian.aspx
- Bosker, R. J. (1998). The class size question in primary schools: Policy issues, theory, and empirical findings from the Netherlands. *International Journal of Educational Research*, 29, 763–778. doi:10.1016/S0883-0355(98)00062-7
- Bosworth, R., & Caliendo, F. (2007). Educational production and teacher preferences. *Economics of Education Review*, 26, 487–500. doi:10.1016/j.econedurev.2005.04.004
- Brewer, D. J., Krop, C., Gill, B. P., & Reichardt, R. (1999). Estimating the cost of national class size reductions under different policy alternatives. *Educational Evaluation and Policy Analysis*, 21, 179–192. doi:10.3102/01623737021002179
- Briggs, D. C. (2008). Comments on Slavin: Synthesizing causal inferences. *Educational Researcher*, 37, 15–22. doi:10.3102/0013189X08314286
- Brock, T. (2005). Viewing mixed methods through an implementation research lens: A response to the New Hope and Moving to Opportunity evaluations. In T. S. Weisner (Ed.), *Discovering successful pathways in children's development: Mixed methods in the study of childhood and family life* (pp. 317–325). Chicago, IL: University of Chicago Press.
- Brock, T., Doolittle, F., Fellerath, V., & Wiseman, M. (1997). *Creating New Hope: Implementation of a program to reduce poverty and reform welfare*. New York, NY: Manpower Demonstration Research Corporation. (ERIC Document Reproduction Service No. ED414443)
- Browning, C. R., & Cagney, K. A. (2003). Moving beyond poverty: Neighborhood structure, social processes, and health. *Journal of Health and Social Behavior*, 44, 552–571. Retrieved from <http://www.jstor.org/stable/1519799>
- Burton, P., Goodlad, R., & Croft, J. (2006). How would we know what works? Context and complexity in the evaluation of community involvement. *Evaluation*, 12, 294–312. doi:10.1177/1356389006069136
- Chalmers, I. (2003). Trying to do more good than harm in policy and practice: The role of rigorous, transparent, up-to-date evaluations. *The ANNALS of the American Academy of Political and Social Science*, 589, 22–40. doi:10.1177/0002716203254762
- Chatterji, M. (2005). Evidence on “what works”: An argument for extended-term mixed-method (ETMM) evaluation designs. *Educational Researcher*, 34, 14–24. doi:10.3102/0013189X034005014
- Chatterji, M. (2008). Comments on Slavin: Synthesizing evidence from impact evaluations in education to inform action. *Educational Researcher*, 37, 23–26. doi:10.3102/0013189X08314287
- Chen, H. T., & Rossi, P. H. (1983). Evaluating with sense: The theory-driven approach. *Evaluation Review*, 7, 283–302. doi:10.1177/0193841X8300700301
- Clampet-Lundquist, S., Edin, K., Kling, J. R., & Duncan, J. G. (2006). *Moving at-risk youth out of high-risk neighborhoods: Why do girls fare better than boys?* (Working Paper No. 509). Princeton, NJ: Princeton University, Industrial Relations Section. Retrieved from <http://www.irs.princeton.edu/pubs/pdfs/509.pdf>
- Cohen, D. K., Raudenbush, S. W., & Ball, D. L. (2003). Resources, instruction, and research. *Educational Evaluation and Policy Analysis*, 25, 119–142. doi:10.3102/01623737025002119
- Cook, T. D. (2002). Randomized experiments in educational policy research: A critical examination of the reasons the educational evaluation community has offered for not doing them. *Educational Evaluation and Policy Analysis*, 24, 175–199. doi:10.3102/01623737024003175
- Cook, T. D. (2003). Why have educational evaluators chosen not to do randomized experiments? *The ANNALS of the American Academy of Political and Social Science*, 589, 114–149. doi:10.1177/0002716203254764
- Cook, T. D., & Campbell, D. T. (1979). *Quasi-experimentation: Design & analysis issues for field settings*. Boston, MA: Houghton Mifflin.
- Cook, T. D., Shadish, W. R., & Wong, V. C. (2005). *Within-study comparisons of experiments and non-experiments: Can they help decide on evaluation policy?* Unpublished manuscript.
- Cooper, C. R. (2005). *Developmental pathways through middle childhood: Rethinking contexts and diversity as resources*. Mahwah, NJ: Erlbaum.
- Cooper, C. R., Brown, J., Azmitia, M., & Chavira, G. (2005). Including Latino immigrant families, schools, and community programs as research partners on the good path of life. In T. S. Weisner (Ed.), *Discovering successful pathways in children's development: Mixed methods in the study of childhood and family life* (pp. 359–385). Chicago, IL: University of Chicago Press.
- Cronbach, L. J., & Shapiro, K. (1982). *Designing evaluations of educational and social programs* (1st ed.). San Francisco, CA: Jossey-Bass.

- Datar, A., & Mason, B. (2008). Do reductions in class size “crowd out” parental investment in education? *Economics of Education Review*, 27, 712–723. doi:10.1016/j.econedurev.2007.10.006
- Datta, L. (2005). Mixed methods, more justified conclusions: The case of the Abt Evaluation of the Comer Program in Detroit. In T. S. Weisner (Ed.), *Discovering successful pathways in children’s development: Mixed methods in the study of childhood and family life* (pp. 65–83). Chicago, IL: University of Chicago Press.
- Dee, T. S. (2004). Teachers, race, and student achievement in a randomized experiment. *Review of Economics and Statistics*, 86, 195–210. doi:10.1162/003465304323023750
- Donaldson, S. I. (2007). *Program theory-driven evaluation science: Strategies and applications*. New York, NY: Erlbaum.
- Duncan, G. J., & Gibson-Davis, C. M. (2006). Connecting child care quality to child outcomes: Drawing policy lessons from non-experimental data. *Evaluation Review*, 30, 611–630. doi:10.1177/0193841X06291530
- Duncan, G. J., Huston, A. C., & Weisner, T. S. (2007). *Higher ground: New Hope for the working poor and their children*. New York, NY: Russell Sage Foundation.
- Duncan, G. J., & Magnuson, K. A. (2003). The promise of random-assignment social experiments for understanding well-being and behavior. *Current Sociology*, 51, 529–541. doi:10.1177/00113921030515005
- Duncan, G. J., Magnuson, K. A., & Ludwig, J. (2004). The endogeneity problem in developmental studies. *Research in Human Development*, 1, 59–80. doi:10.1207/s15427617rhd0101&2_5
- Duncan, G. J., & Raudenbush, S. W. (1999). Assessing the effects of context in studies of child and youth development. *Educational Psychologist*, 34, 29–41. doi:10.1207/s15326985ep3401_3
- Duncan, G. J., & Raudenbush, S. W. (2001). Neighborhoods and adolescent development: How can we determine the links? In A. Booth & A. C. Crouter (Eds.), *Does it take a village? Community effects on children, adolescents, and families* (pp. 105–136). Mahwah, NJ: Erlbaum.
- Ehrenberg, R. G., Brewer, D. J., Gamoran, A., & Willms, J. D. (2001). Class size and student achievement. *Psychological Science in the Public Interest*, 2, 1–30. Retrieved from https://www.psychologicalscience.org/journals/pspi/pdf/pspi2_1.pdf
- Eisenhart, M. (2005). Hammers and saws for the improvement of educational research. *Educational Theory*, 55, 245–261. doi:10.1111/j.1741-5446.2005.00002.x
- Eisenhart, M. (2006). Qualitative science in experimental time. *International Journal of Qualitative Studies in Education*, 19, 697–707. doi:10.1080/09518390600975826
- Eisenhart, M., & Towne, L. (2003). Contestation and change in national policy on “scientifically based” education research. *Educational Researcher*, 32, 31–38. doi:10.3102/0013189X032007031
- Feuer, M. J., Towne, L., & Shavelson, R. J. (2002). Scientific culture and educational research. *Educational Researcher*, 31, 4–14. (ERIC No. EJ662137)
- Finn, J. D., & Achilles, C. M. (1990). Answers and questions about class size: A statewide experiment. *American Educational Research Journal*, 27, 557–577. doi:10.3102/00028312027003557
- Finn, J. D., & Achilles, C. M. (1999). Tennessee’s class size study: Findings, implications, misconceptions. *Educational Evaluation and Policy Analysis*, 21, 97–109. doi:10.3102/01623737021002097
- Finn, J. D., Gerber, S. B., Achilles, C. M., & Boyd-Zaharias, J. (2001). The enduring effects of small classes. *Teachers College Record*, 103, 145–183. doi:10.1111/0161-4681.00112
- Finn, J. D., Gerber, S. B., & Boyd-Zaharias, J. (2005). Small classes in early grades, academic achievement, and graduating from high school. *Journal of Educational Psychology*, 97, 214–223. doi:10.1037/0022-0663.97.2.214
- Finn, J. D., Pannozzo, G. M., & Achilles, C. M. (2003). The “why’s” of class size: Student behavior in small classes. *Review of Educational Research*, 73, 321–368. doi:10.3102/00346543073003321
- Fricke, T. (2005). Taking culture seriously: Making the social survey ethnographic. In T. S. Weisner (Ed.), *Discovering successful pathways in children’s development: Mixed methods in the study of childhood and family life* (pp. 185–221). Chicago, IL: University of Chicago Press.
- Gardenhire, A., & Nelson, L. (2003). *Intensive qualitative research: Challenges, best uses, and opportunities* (MDRC Working Paper). New York, NY: MDRC. Retrieved from <http://www.mdrc.org/publications/339/full.pdf>
- Gennetian, L. A., Duncan, G., Knox, V., Vargas, W., Clark-Kauffman, E., & London, A. S. (2004). How welfare policies affect adolescents’ school outcomes: A synthesis of evidence from experimental studies. *Journal of Research on Adolescence*, 14, 399–423. doi:10.1111/j.1532-7795.2004.00080.x
- Gerber, S. B., Finn, J. D., Achilles, C. M., & Boyd-Zaharias, J. (2001). Teacher aides and students’ academic achievement. *Educational Evaluation and Policy Analysis*, 23, 123–143. doi:10.3102/01623737023002123
- Gibson-Davis, C. M., & Duncan, G. (2005). Qualitative/quantitative synergies in a random-assignment program evaluation. In T. S. Weisner (Ed.), *Discovering successful pathways in children’s development: Mixed methods in the study of childhood and family life* (pp. 283–303). Chicago, IL: University of Chicago Press.
- Goering, J., & Feins, J. D. (Eds.). (2003). *Choosing a better life: Evaluating the Moving to Opportunity social experiment*. Washington, DC: Urban Institute Press.

- Goldenberg, C., Gallimore, R., & Reese, L. (2005). Using mixed methods to explore Latino children's literacy development. In T. S. Weisner (Ed.), *Discovering successful pathways in children's development: Mixed methods in the study of childhood and family life* (pp. 21–46). Chicago, IL: University of Chicago Press.
- Gordon, N. (2004). Do federal grants boost school spending? Evidence from Title I. *Journal of Public Economics*, 88(9–10), 1771–1792. doi:10.1016/j.jpubeco.2003.09.002
- Greene, J. (2005). A reprise on mixing methods. In T. S. Weisner (Ed.), *Discovering successful pathways in children's development: Mixed methods in the study of childhood and family life* (pp. 405–419). Chicago, IL: University of Chicago Press.
- Greenwald, R., Hedges, L. V., & Laine, R. D. (1996). The effect of school resources on student achievement. *Review of Educational Research*, 66, 361–396. doi:10.3102/00346543066003361
- Grissmer, D. W. (1999). Conclusion: Class size effects: Assessing the evidence, its policy implications, and future research agenda. *Educational Evaluation and Policy Analysis*, 21, 231–248. doi:10.3102/01623737021002231
- Grissmer, D. W. (2002). Cost-effectiveness and cost-benefit analysis: The effect of targeting interventions. In H. M. Levin & P. J. McEwan (Eds.), *Cost-effectiveness and educational policy* (pp. 97–108). Larchmont, NY: Eye on Education.
- Grissmer, D. W., & Flanagan, A. (2006). *Improving the achievement of Tennessee students: An analysis of the National Assessment of Educational Progress*. Santa Monica, CA: Rand. Retrieved from http://www.rand.org/pubs/technical_reports/2006/RAND_TR381.sum.pdf
- Grissmer, D. W., Flanagan, A., Kawata, J. H., & Williamson, S. (2000). *Improving student achievement: What state NAEP test scores tell us*. Santa Monica, CA: Rand. (ERIC No. ED440154)
- Gueron, J. M. (2002). The politics of random assignment. In F. Mosteller & R. F. Boruch (Eds.), *Evidence matters* (pp. 15–49). Washington, DC: Brookings Institution Press.
- Gueron, J. M. (2003). Fostering research excellence and impacting policy and practice: The welfare reform story. *Journal of Policy Analysis and Management*, 22, 163–174. doi:10.1002/pam.10110
- Gueron, J. M., (2007). Building evidence: What it takes and what it yields. *Research on Social Work Practice*, 17, 134–142. doi:10.1177/1049731506293095
- Hanushek, E. A. (1997). Assessing the effects of school resources on student performance: An update. *Educational Evaluation and Policy Analysis*, 19, 141–164. doi:10.3102/01623737019002141
- Hanushek, E. A. (1999). Some findings from an independent investigation of the Tennessee STAR experiment and from other investigations of class size effects. *Educational Evaluation and Policy Analysis*, 21, 143–163. doi:10.3102/01623737021002143
- Hanushek, E. A. (2002). Evidence, politics and the class size debate. In L. Mishel & R. Rothstein (Eds.), *The class size debate* (pp. 37–65). Washington, DC: Economic Policy Institute.
- Harkness, S., Hughes, M., Muller, B., & Super, C. M. (2005). Entering the developmental niche: Mixed methods in an intervention program for inner-city children. In T. S. Weisner (Ed.), *Discovering successful pathways in children's development: Mixed methods in the study of childhood and family life* (pp. 329–358). Chicago, IL: University of Chicago Press.
- Hattie, J. (2005). The paradox of reducing class size and improving learning outcomes. *International Journal of Educational Research*, 43, 387–425. doi:10.1016/j.ijer.2006.07.002
- Heckman, J. J., & Smith, J. A. (1995). Assessing the case for social experiments. *Journal of Economic Perspectives*, 9, 85–110. Retrieved from <http://www.jstor.org/stable/2138168>
- Hedges, L. V., Konstantopoulos, S., & Nye, B. A. (2001). The long-term effects of small classes in early grades: Lasting benefits in mathematics achievement at grade 9. *Journal of Experimental Education*, 69, 245–257.
- Howe, K. R. (1998). The interpretive turn and the new debate in education. *Educational Researcher*, 27, 13–20. doi:10.3102/0013189X027008013
- Howe, K. R. (2004). A critique of experimentalism. *Qualitative Inquiry*, 10, 42–61. doi:10.1177/1077800403259491
- Huston, A. C. (2005). Mixed methods in studies of social experiments for parents in poverty. In T. S. Weisner (Ed.), *Discovering successful pathways in children's development: Mixed methods in the study of childhood and family life* (pp. 305–315). Chicago, IL: University of Chicago Press.
- Huston, A. C., Duncan, G. J., Granger, R., Bos, J., McLoyd, V., Mistry, R., . . . Ventura, A. (2001). Work-based antipoverty programs for parents can enhance the school performance and social behavior of children. *Child Development*, 72, 318–336. Retrieved from <http://www.jstor.org/stable/1132487>
- Jepsen, C., & Rivkin, S. G. (2002). *Class size reduction, teacher quality, and academic achievement in California public elementary schools*. San Francisco, CA: Public Policy Institute of California. Retrieved from http://www.ppic.org/content/pubs/rb/RB_602CJRB.pdf
- Karoly, L. A., & Bigelow, J. H. (2005). *The economics of investing in universal preschool education in California* (Monograph No. 349). Santa Monica, CA: Rand. Retrieved from http://www.rand.org/pubs/monographs/2005/RAND_MG349.1.pdf
- Karoly, L. A., Greenwood, P. W., Everingham, S. S., Hoube, J., Kilburn, M. R., Rydell, C. P., . . . Chiesa, J. (1998). *Investing in our children: What we know and don't know about the costs and benefits of early childhood interventions* (Monograph Report No. 898). Santa Monica, CA: Rand.

- Karoly, L. A., Kilburn, M. R., & Cannon, J. S. (2005). *Early childhood interventions: Proven results, future promise* (Monograph No. 341). Santa Monica, CA: Rand. Retrieved from http://www.rand.org/pubs/monographs/2005/RAND_MG341.pdf
- Katz, L. F., Kling, J. R., & Liebman, J. B. (2001). Moving to opportunity in Boston: Early results of a randomized mobility experiment. *Quarterly Journal of Economics*, *116*, 607–654. doi:10.1162/00335530151144113
- Kling, J. R., Liebman, J. B., & Katz, L. F. (2005). Bullets don't got no name: Consequences of fear in the ghetto. In T. S. Weisner (Ed.), *Discovering successful pathways in children's development: Mixed methods in the study of childhood and family life* (pp. 243–281). Chicago, IL: University of Chicago Press.
- Kling, J. R., Liebman, J. B., & Katz, L. F. (2007). Experimental analysis of neighborhood effects. *Econometrica*, *75*, 83–119.
- Kling, J. R., Ludwig, J., & Katz, L. F. (2005). Neighborhood effects on crime for female and male youth: Evidence from a randomized housing voucher experiment. *Quarterly Journal of Economics*, *120*, 87–130. doi:10.1162/qjec.2005.120.1.87
- Krueger, A. B. (1999). Experimental estimates of education production functions. *Quarterly Journal of Economics*, *114*, 497–532. doi:10.1162/003355399556052
- Krueger, A. B. (2002). Understanding the magnitude and effect of class size on student achievement. In L. Mishel & R. Rothstein (Eds.), *The class size debate* (pp. 7–35). Washington, DC: Economic Policy Institute.
- Krueger, A. B. (2003). Economic considerations and class size. *Economic Journal*, *113*(485), F34–F63. doi:10.1111/1468-0297.00098
- Krueger, A. B., & Whitmore, D. M. (2001). The effect of attending a small class in the early grades on college-test taking and middle school test results: Evidence from Project STAR. *Economic Journal*, *111*, 1–28. doi:10.1111/1468-0297.00586
- Krueger, A. B., & Whitmore, D. M. (2002). Would smaller class sizes help close the black–white achievement gap? In J. E. Chubb & T. Loveless (Eds.), *Bridging the achievement gap* (pp. 11–46). Washington, DC: Brookings Institution Press.
- Ladd, H. F., & Ludwig, J. (1997). Federal housing assistance, residential relocation, and educational opportunities: Evidence from Baltimore. In Papers and Proceedings of the Hundred and Fourth Annual Meeting of the American Economic Association. *American Economic Review*, *87*, 272–277. <http://www.aeaweb.org/aer/index.php>
- Lazear, E. P. (2001). Educational production. *Quarterly Journal of Economics*, *116*, 777–803. doi:10.1162/00335530152466232
- Leventhal, T., & Brooks-Gunn, J. (2003a). Children and youth in neighborhood contexts. *Current Directions in Psychological Science*, *12*, 27–31. doi:10.1111/1467-8721.01216
- Leventhal, T., & Brooks-Gunn, J. (2003b). Moving to opportunity: An experimental study of neighborhood effects on mental health. *American Journal of Public Health*, *93*, 1576–1582. doi:10.2105/AJPH.93.9.1576
- Levin, H. M., & McEwan, P. J. (2000). *Cost-effectiveness analysis: Methods and applications* (2nd ed.). Thousand Oaks, CA: Sage.
- Levin, H. M., & McEwan, P. J. (Eds.). (2002). *Cost-effectiveness and educational policy*. Larchmont, NY: Eye on Education.
- Lowe, E. D., & Weisner, T. S. (2004). “You have to push it—Who’s gonna raise your kids?”: Situating child care and child care subsidy use in the daily routines of lower income families. *Children and Youth Services Review*, *26*, 143–171. doi:10.1016/j.childyouth.2004.01.011
- Lowe, E. D., Weisner, T. S., & Geis, S. (2003). *Instability in child care: Ethnographic evidence from working poor families in the New Hope Intervention* (Next Generation Working Paper No. 15). New York: MDRC.
- Ludwig, J., Duncan, G. J., & Hirschfield, P. (2001). Urban poverty and juvenile crime: Evidence from a randomized housing-mobility experiment. *Quarterly Journal of Economics*, *116*, 655–679. doi:10.1162/00335530151144122
- Lynch, R. G. (2004). *Exceptional returns: Economic, fiscal, and social benefits of investment in early childhood development*. Washington, DC: Economic Policy Institute.
- Masse, L. N., & Barnett, W. S. (2002). A benefit-cost analysis of the Abecedarian early childhood intervention. In H. M. Levin & P. J. McEwan (Eds.), *Cost effectiveness and educational policy* (pp. 157–173). Larchmont, NY: Eye on Education.
- Maxwell, J. A. (2004). Causal explanation, qualitative research, and scientific inquiry in education. *Educational Researcher*, *33*, 3–11. doi:10.3102/0013189X033002003
- Michelson, A. A., & Morley, E. W. (1887). On the relative motion of the earth and the luminiferous ether. *American Journal of Science*, *34*, 334–345.
- Mosteller, F. (1995). The Tennessee study of class size in the early school grades. *The Future of Children*, *5*, 113–127. Retrieved from http://futureofchildren.org/futureofchildren/publications/docs/05_02_08.pdf
- Mosteller, F., & Boruch, R. F. (Eds.). (2002). *Evidence matters: Randomized trials in education research*. Washington, DC: Brookings Institution Press.
- Nye, B. A., Hedges, L. V., & Konstantopoulos, S. (2000a). Do the disadvantaged benefit more from small classes? Evidence from the Tennessee class size experiment. *American Journal of Education*, *109*, 1–26. doi:10.1086/444257
- Nye, B. A., Hedges, L. V., & Konstantopoulos, S. (2000b). The effects of small classes on academic achievement: The results of the Tennessee class size experiment. *American Educational Research Journal*, *37*, 123–151. doi:10.3102/00028312037001123

- Nye, B. A., Hedges, L. V., & Konstantopoulos, S. (2002). Do low achieving students benefit more from small classes? Evidence from the Tennessee class size experiment. *Educational Evaluation and Policy Analysis, 24*, 210–217. doi:10.3102/01623737024003201
- Nye, B. A., Hedges, L. V., & Konstantopoulos, S. (2004). Do minorities experience larger lasting benefits from small classes? Evidence from a five-year follow-up of the Tennessee class size experiment. *Journal of Educational Research, 98*, 94–100. doi:10.3200/JOER.98.2.94-114
- Nye, B. A., Konstantopoulos, S., & Hedges, L. V. (2004). How large are teacher effects? *Educational Evaluation and Policy Analysis, 26*, 237–257. doi:10.3102/01623737026003237
- Oakes, J. M. (2004). The (mis)estimation of neighborhood effects: Causal inference for a practicable social epidemiology. *Social Science & Medicine, 58*, 1929–1952. doi:10.1016/j.socscimed.2003.08.004
- O'Connor, T. G. (2003). Natural experiments to study the effects of early experience: Progress and limitations. *Development and Psychopathology, 15*, 837–852. doi:10.1017/S0954579403000403
- Peevely, G., Hedges, L. V., & Nye, B. A. (2005). The relationship of class size effects and teacher salary. *Journal of Education Finance, 31*, 101–109.
- Poglinco, S. M., Brash J., & Granger, R. C. (1998). *An early look at community service jobs in the New Hope demonstration*. New York, NY: MDRC.
- Quint, J., Bloom, H. S., Black, A. R., & Stephens, L. (with Akey, T. M.). (2005). *The challenge of scaling up educational reform: Findings and lessons from First Things First (Final Report)*. New York, NY: MDRC. Retrieved from <http://www.mdrc.org/publications/412/full.pdf>
- Ramey, C. T., Campbell, F. A., Burchinal, M., Skinner, M. L., Gardner, D. M., & Ramey, S. L. (2000). Persistent effects of early childhood education on high-risk children and their mothers. *Applied Developmental Science, 4*, 2–14. doi:10.1002/9780470755778.ch3
- Raudenbush, S. W. (2005). Learning from attempts to improve schooling: The contribution of methodological diversity. *Educational Researcher, 34*, 25–31. doi:10.3102/0013189X034005025
- Reynolds, A. J., Temple, J. A., Ou, S. R., Robertson, D. L., Mersky, J., Topitzes, J. W., et al. (2007, March–April). *Effects of a preschool and school-aged intervention on adult health and well-being: Evidence from the Chicago Longitudinal Study*. Paper presented at the biannual meeting of the Society for Research on Child Development, Boston, MA.
- Reynolds, A. J., Temple, J. A., Robertson, D. L., & Mann, E. A. (2002). Age 21 cost-benefit analysis of the Title I Chicago child-parent centers. *Educational Evaluation and Policy Analysis, 24*, 267–303. doi:10.3102/01623737024004267
- Ritter, G. W., & Boruch, R. F. (1999). The political and institutional origins of a randomized controlled trial on elementary school class size: Tennessee's Project STAR. *Educational Evaluation and Policy Analysis, 21*, 111–125. doi:10.3102/01623737021002111
- Romich, J. L. (2006). Randomized social policy experiments and research on child development. *Journal of Applied Developmental Psychology, 27*, 136–150. doi:10.1016/j.appdev.2006.01.001
- Rosenbaum, E., & Harris, L. E. (2001). Residential mobility and opportunities: Early impacts of the Moving to Opportunity Demonstration Program in Chicago. *Housing Policy Debate, 12*, 321–346.
- Rosenzweig, M. R., & Wolpin, K. I. (2000). Natural “natural experiments” in economics. *Journal of Economic Literature, 38*, 827–874.
- Rutter, M. (2002). Nature, nurture, and development: From evangelism through science toward policy and practice. *Child Development, 73*, 1–21. doi:10.1111/1467-8624.00388
- Salomon, G. (1991). Transcending the qualitative-quantitative debate: The analytic and systemic approaches to educational research. *Educational Researcher, 20*, 10–18. doi:10.3102/0013189X020006010
- Sanbonmatsu, L., Kling, J. R., Duncan G. J., & Brooks-Gunn, J. (2006). Neighborhoods and academic achievement: Results from the Moving to Opportunity Experiment. *Journal of Human Resources, 41*, 649–691.
- Schneider, B. L., & McDonald, S. K. (Eds.). (2007). *Scale-up in education*. Lanham, MD: Rowman & Littlefield.
- Schweinhart, L. J. (2004). *The High/Scope Perry Preschool Study through age 40*. Ypsilanti, MI: High/Scope Educational Research Foundation. Retrieved from <http://www.highscope.org/Content.asp?ContentId=219>
- Shadish, W. R., Cook, T. D., & Campbell, D. T. (2002). *Experimental and quasi-experimental designs for generalized causal inference*. Boston, MA: Houghton Mifflin.
- Shavelson, R. J., & Towne, L. (Eds.). (2002). *Scientific research in education*. Washington, DC: National Academy Press.
- Sims, D. P. (2004). *Unintended consequences of education and housing reform incentives* (Master's thesis). Available from Proquest Dissertation Express. (AAT 0807684)
- Slavin, R. E. (2002). Evidence-based education policies: Transforming educational practice and research. *Educational Researcher, 31*, 15–21. doi:10.3102/0013189X031007015
- Slavin, R. E. (2008). Perspectives on evidence-based research in education—What works? Issues in synthesizing educational program evaluations. *Educational Researcher, 37*, 5–14. doi:10.3102/0013189X08314117

- Tilley, N. (2004). Applying theory-driven evaluation to the British crime reduction programme: The theories of the programme and of its evaluations. *Criminal Justice*, 4, 255–276. doi:10.1177/1466802504048465
- Towne, L., Shavelson, R. J., & Feuer, M. J. (Eds.). (2001). *Science, evidence, and inference in education: Report of a workshop*. Washington, DC: National Academy Press.
- Turney, K., Clampet-Lundquist, S., Edin, K., Kling, J. R., & Duncan, G. J. (2006). Neighborhood effects on barriers to employment: Results from a randomized housing mobility experiment in Baltimore. In G. Burtless & J. R. Pack (Eds.), *Brookings-Wharton papers on urban affairs: 2006* (pp. 137–187). Washington, DC: Brookings Institution Press.
- Walshe, K. (2007). Understanding what works—and why—in quality improvement: The need for theory-driven evaluation. *International Journal for Quality in Health Care*, 19, 57–59.
- Webbink, D. (2005). Causal effects in education. *Journal of Economic Surveys*, 19, 535–560.
- Weisner, T. S. (2002). Ecocultural understanding of children's developmental pathways. *Human Development*, 45, 275–281. doi:10.1159/000064989
- Weisner, T. S. (Ed.). (2005). *Discovering successful pathways in children's development: Mixed methods in the study of childhood and family life*. Chicago, IL: University of Chicago Press.
- Weiss, H. B., Kreider, H., Mayer, E., Hencke, R., & Vaughan, M. (2005). Working it out: The chronicle of a mixed methods analysis. In T. S. Weisner (Ed.), *Discovering successful pathways in children's development: Mixed methods in the study of childhood and family life* (pp. 47–64). Chicago, IL: University of Chicago Press.
- Wilde, E. T., & Hollister, R. (2007). How close is close enough? Evaluating propensity score matching using data from a class size reduction experiment. *Journal of Policy Analysis and Management*, 26, 455–477. doi:10.1002/pam.20262
- Wilson, W. J. (1996). *When work disappears: The world of the new urban poor* (1st ed.). New York, NY: Knopf.
- Yoshikawa, H., Weisner, T. S., Kalil, A., & Way, N. (2008). Mixing qualitative and quantitative research in developmental science: Uses and methodological choices. *Developmental Psychology*, 44, 344–354. doi:10.1037/0012-1649.44.2.344
- Yoshikawa, H., Weisner, T. S., & Lowe, E. D. (Eds.). (2006). *Making it work: Low-wage employment, family life, and child development*. New York, NY: Russell Sage Foundation.

APPENDIX

PROJECT DESCRIPTIONS

Project Star

The multi-district Tennessee STAR experiment randomly assigned students in a single cohort of kindergarten students (the 1986–1987 entering cohort) in 79 participating schools that each had at least three kindergarten classrooms. The students were assigned to one of three groups: (a) large classes (approximate mean of 22–24 students) with a teacher aide, (b) large classes without a teacher aide, or (c) small classes (approximate mean of 15–16 students). Kindergarten teachers in each school were also assigned randomly to one of these types of classrooms. The structure of this design meant that each school constituted a separate random experiment.

The sample of entering students across 328 kindergarten classes was approximately 6,500 students. Those students entering at kindergarten were scheduled to maintain their treatment through first, second, and third grade. However, the groups changed in significant ways over the 4 years of the experiment due to sample attrition and new students entering schools during kindergarten, first, second, and third grade. Students entering participating schools after kindergarten entrance were also randomly assigned to one of the three groups, but these late-entering students came from schools outside the sample and likely had been in larger classrooms prior to entering a participating school. Many students originally in participating schools also

moved away after 1, 2, or 3 years. As a result, over the course of the experiment, almost 12,000 children participated.

Students leaving the sample and those entering later generally had lower scores on standardized math and reading tests than those who began and remained in the original entering cohort. The 12,000 study subjects therefore consisted of some who remained in the sample all 4 years, some who entered the school later and had fewer than 4 years in an assigned treatment group, and some who left the sample after kindergarten entry and completed fewer than 4 years in a treatment or control group. There were also some crossovers (about 15%) who switched at some point from one treatment group to another.

In addition to administering mathematics and reading tests in each grade, teachers and aides completed questionnaires and time logs to document their perceptions and experiences. In the Grade 4 follow-up study, researchers collected behavioral data in addition to achievement scores. Using a 28-item Student Participation Questionnaire, Grade 4 teachers rated each pupil who had been in STAR. This instrument assesses specific learning behaviors (“engagement behaviors”) judged by educators to be important in the classroom. The instrument yields reliable, valid measures of the effort students allot to learning, initiative-taking in the classroom, and nonparticipatory behavior (disruptive or inattention).

Researchers also conducted follow-up measurements with the students at the fourth- and eighth-grade levels and after high school. At the fourth- and eighth-grade levels,

researchers collected reading and math assessment data. In high school, measurements included college entrance test taking (SAT and ACT)) and whether students completed high school.

The major results from Project STAR include the following:

- Students assigned to smaller classes all 4 years had statistically significant achievement gains of about .15 to .40 standard deviations above the mean of students assigned to the two large-class groups (with and without teacher aides).
- Gains in reading were not significantly different from gains in mathematics.
- The effect of teacher aides in large classes was small and positive but statistically insignificant when compared with large-classes without aides.
- Effect sizes were much larger for minority and disadvantaged students in small classes.
- Students assigned to small classes for 3–4 years had much higher gains than those in small classes for only 1–2 years.
- Significant effects persisted through eighth grade for students who had participated in smaller K–3 classes, though effect sizes declined somewhat from those at third grade, and effects had greater persistence through eighth grade if students had more years in smaller K–3 classes.
- Students in small classes in K–3 had higher high school graduation rates and increased incidence of testing linked to college applications (SAT and ACT).

New Hope

New Hope was an ambitious project based on two simple yet widely held principles: (a) People who are willing to work full time should be able to do so, and (b) they should not be poor as a result. The program was designed to improve the lives of low-income individuals and families by providing several benefits for parents who worked full time: an earnings supplement to raise their income above poverty, subsidized health insurance, and subsidized child care. The program also offered access to wage-paying community service jobs for people who could not find full-time work.

New Hope was run as a demonstration project from 1994 to 1998 in two inner-city areas of Milwaukee, WI, by the New Hope Project, Inc., a local community-based organization. The researchers targeted New Hope at two geographic areas with high levels of poverty, thus allowing

a more detailed analysis of program context than would be possible in a program that served a wide geographic area.

The program had only four eligibility requirements: that an applicant live in one of the two targeted service areas, be age 18 or older, be willing and able to work at least 30 hours per week, and have a household income at or below 150% of the federally defined poverty level. Participation was voluntary, and adults were eligible regardless of whether they had children or received public assistance.

Persons who met these criteria were eligible to receive 3 years of the following benefits or services:

- Help in obtaining a job, including access to a time-limited, minimum-wage community service job if full-time employment was not otherwise available.
- A monthly earnings supplement that, when combined with federal and state earned income credit, brought most low-wage workers' incomes above the poverty level.
- Subsidized health insurance, which gradually phased out as earnings rose.
- Subsidized child care, which also gradually phased out as earnings rose.

The study enrolled over 1,300 low-income adults who volunteered to participate. Half the applicants were randomly assigned to a program group that was eligible to receive New Hope's benefits, and the other half were randomly assigned to a control group that was not eligible for the enhanced benefits. From the total sample of 1,357 people, 745 people had at least one child between the ages of 1 and 10 at the time of enrollment.

New Hope program data provide information on parents' use of the program's services, as well as their job status, hours worked, and earnings. State administrative records provide data on employment and receipt of welfare and food stamp benefits. Researchers collected primary data in the form of interviews and surveys at 2, 5, and 8 years from the beginning of the experiment. In-person surveys revealed information on families' receipts of New Hope benefits, job histories, parents' employment and earnings, family functioning, and parent-child relations. For up to two "focal" children in each family, the surveys also collected information from parents, teachers, and children on school performance, psychological well-being, and behavior problems. At the 5-year interview stage, children took standardized tests. Parents were also asked about stress levels, depression, and their hopes for the future. Parents and children reported on parent-child relationships, children's experience in child care, and activities outside school.

To better understand the detailed dynamics and contexts of family life, fieldworkers drew an ethnographic sample of 44 families from the total sample of participants. They gave these families—half of whom were in the New Hope group and half of whom were in the control group—periodic in-depth interviews from the third year to the final year of the New Hope program (1998–2001) and again in 2004. The ethnographic data include extensive field notes as well as focused interviews covering a wide range of topics, including, for example, parents’ experiences with New Hope, family routines, work experiences, family relationships, child-care arrangements, and goals. Unlike surveys, these open-ended interviews and conversations allowed participants the opportunity to tell their stories. Families did not shy away from talking about difficult issues—domestic abuse, drugs and alcohol, family conflicts, and health problems. In addition to conducting interviews, the ethnographic fieldworkers participated in family routines and events including lunches, dinners, birthday parties, and trips to the mall.

Key New Hope differences between treatment and control groups included the following:

- New Hope had varying impact on work and earnings across study subgroups.
- For individuals working little or not at all at the beginning of the program, New Hope led to more work and higher earnings during the 4 years of operation but did not have persisting significant effects after the program ended.
- For individuals already working full time at the beginning of the program, the program showed no effects on long-term work or income.
- The effects on work and earnings were significant and persisted after the program ended for some individuals whose barriers to employment were addressed by New Hope benefits (e.g., child care or health insurance).
- For women without children at the beginning of the program, there were no work or earnings effects.
- For men without children at the beginning of the program, there were boosts to work and income, but only sporadically.
- Partly due to income supplements, New Hope reduced poverty substantially during and modestly after the end of the program.
- New Hope child-care subsidies increased children’s participation in center-based child care and after-school programs.

- New Hope insurance benefits led to fewer episodes of unmet medical and dental needs and some improvement in adult mental and physical health.
- New Hope improved children’s school performance, especially in reading.
- For boys, New Hope led to increased positive social behaviors and reduced behavior problems and increased engagement in school and higher education.
- For girls, New Hope had mixed effects. Parents reported improvements in their daughters’ positive behaviors, but teachers reported worse behavior for those same girls at school.

Moving to Opportunity (MTO)

The MTO demonstration program was designed to assess the impact of providing families living in subsidized housing in high poverty neighborhoods with the opportunity to move to neighborhoods with lower levels of poverty. Families were recruited for the MTO program from public housing developments in Boston, Baltimore, Chicago, Los Angeles, and New York. Researchers primarily targeted housing developments located in census tracts with 1990 poverty rates of at least 40%. The average poverty rate in these tracts in 1990 was 67%.

Program eligibility requirements included residing in a targeted development, having very low income that met the Section 8 income limits of the public housing authority, having a child under 18, and being in good standing with the housing authority. Participants volunteered to be part of the study. Families that volunteered for the program were more disadvantaged than their public housing counterparts who did not join MTO. MTO families were more likely than nonparticipating families to receive welfare and to be headed by women who were young and unemployed.

Volunteering families initially living in public housing were assigned by lottery to three groups:

- The control group received no new assistance but continued to be eligible to stay in public housing.
- The Section 8 group received traditional Section 8 voucher that enabled movement from public housing to subsidized rental housing without geographic restriction.
- The experimental group received a Section 8 voucher, restricted for one year to a census tract with a poverty rate of less than 10%.

From 1994 to 1997, 4,248 eligible families were randomly assigned to one of these three groups. Families in the treatment groups had 4–6 months to find qualified housing

and to move, using an MTO voucher. Forty-seven percent of the experimental group families and 59% of the Section 8 group families used the program housing voucher to “lease-up,” or move to a new apartment.

Baseline interviews with heads of households were conducted from 1994 to 1999, before random assignment and relocation of movers. The structured interviews focused on demographic information for householders and children and data from householders on labor force and welfare benefits characteristics.

Researchers supplemented MTO baseline surveys with state administrative earnings and welfare data. In 2002, 4–7 years after enrollment, researchers surveyed all of the household heads in the experiment, as well as school-aged children and teens in each family. They collected more comprehensive measures related to economic self-sufficiency and mental and physical health outcomes, as well as a broader range of mediating factors, to potentially illuminate the mechanisms by which residential neighborhoods may affect economic and health outcomes. In addition, there were several specialized data collections conducted with subsamples of participants. These included a subsample of children who were administered achievement tests. Researchers obtained juvenile arrest records, as well as qualitative interview data through in-depth personnel interviews and telephone conversations with teens and adults.

Results from the analyses of these data included the following differences between the three groups:

- There were no significant effects among the three groups on measures of work or earnings.
- There were no significant effects among groups on children’s achievement.
- There were significant positive effects on some measures of adult and child mental health for the experimental group.
- Boys in the experimental group fared no better or worse on measures of risk behavior than boys in the control group.
- Girls in the experimental group had improved mental health and lower risk behavior than girls in the control group.
- Adults in the experimental group had reductions in obesity but no effects on other physical health measures.

