

NCSER Webinar  
Evaluation of State and Local Education Programs and Policies

Jackie Buckley:

Thank you for joining the webinar on Evaluation of State and Local Education Programs and Policies. I am Jackie Buckley at the National Center for Special Education Research and I'm joined by Allen Ruby from the National Center for Education Research. For shorthand, we'll say State/Local Evaluation when we refer to this topic area. This is one of two topics being competed under the Partnership and Collaborations Focused on Problems of Practice or Policy grant program, which has a CFDA number of 84.305H.

Today, we are going to give a brief overview of IES and its mission, particularly for those of you who may not have submitted grants to IES before.. We'll talk about the requirements of the grant program -- the general requirements, as well as the specific one. We'll discuss the information that you need to include in your project narrative related to the five sections of your Project Narrative: Significance, Partnership, Research Plan, Personnel, and Resources. We'll talk about other important sections of the grant application, as well as some information about preparing and submitting an application.

The legislative mission of IES is threefold. First is to describe the condition and progress of education in the United States. Second is to identify education practices that improve academic achievement and improve access to education opportunities and third, we are tasked with evaluating the effectiveness of federal and other education programs.

This is an organizational chart of IES and so, across the bottom, you see the four centers within IES. We are led by a director who is advised by the National Board for Education Sciences. The director and the national board are both appointed by the president and confirmed by the senate. The four centers at IES include the National Center for Education Evaluation which does contracted' evaluations. and supports the regional labs, the What Works Clearinghouse, the National Library of Education and ERIC. NCEE's work contributes to the legislative mission of evaluating the effectiveness of federal and other education programs.

Next is the National Center for Education Statistics which supports national data collections such as the Common Core of Data and national surveys such as ELS and NAEP. , NCES's work contributes to the legislative mission of understanding the progress and conditions of education,. Next, highlighted in blue, are the National Center for Education Research, which again, is where Allen sits and the National Center for Special Education Research where I work and these are the two research grant centers within IES which include the State/Local evaluation grants.. Within those two centers, program officers are available to provide advice on the application process, comments on the content of the work you're proposing, as well as answering general questions about submitting applications.

What you also see on this chart which is important to point out is our Standards and Review Office off on the left side. We have a separate office at IES that is responsible for the review process, soliciting reviewers, getting reviewers to review applications for us, organizing the review panels managing that process. The Standards and Review Office does that and what that

allows program officers to do is to be able to work with applicants around issues related to their application and talk through substantive issues with you, research and design issues, as well as general application submission questions.

The IES grant programs within the National Center for Education Research and National Center for Special Education Research are to answer four questions. One is, what works to improve student education outcomes so that we can disseminate them. Second, what does not work, so we can stop using it. Third, what works for whom and where, so we can use it with the appropriate people, in the appropriate places. Fourth, why does it work? So, we can understand how to improve education and we can build on this understanding.

The director sets the institute's priorities, which are also approved by the board. Partnerships have become of growing importance in these priorities in order to do two things really. One is to help focus research on the issues that most concern policymakers and practitioners and second, to help researchers communicate their findings in useful ways. Policymakers and practitioners are to have a strong role in these partnerships in setting their research agenda. And there should be constant communication between the researchers and the practitioners to discuss what's being found and what further research should be done.

So, the State/Local Evaluation grants are to support new or established partnerships that will evaluate education programs or policies from pre-k to adult education that are implemented by state or local education agencies. The program and policies should be of high importance to the state or local education agency.

There are three general requirements that you must meet when submitting an application to the State/Local Evaluation program. First, you must focus on student education outcome. Second, the work is to be done by partnership between a research institutions and a state or local education agency. Third, you must evaluate an education program or policy implemented by a state or local education agency.

Let's talk about these three general requirements in more detail. IES is focused on research that improves the quality of education for all students. All research much address the education outcome of students. And these student outcomes can be grouped under academic outcomes, as well as social and behavioral outcomes that support student success at school. Academic outcomes can include measures of academic achievement such as, test scores (e.g., standardized tests, end-of –course test, graduation tests) and grades , and measures of academic progress such as course and grade completion, high school graduation and dropout, postsecondary access, progress and completion.

Social behavior outcomes can include a variety of outcomes, such as social skills, learning strategies, goal setting, self-regulated learning, student attitudes, motivation, and student behavior, such as attendance and disciplinary outcomes.

These academic outcomes and social and behavioral outcomes are for students from pre-kindergarten through post-secondary and adult education. And you can address any grade or grade range within those boundaries. So, for example, you could focus solely on fifth grade.

You can examine all grades in middle school. You could look at that transition between high school and college. In addition, these outcomes are for students without disabilities, as well as students with or at-risk for disability.

If you're going to research student with or at-risk for disabilities, you need to take a look at the website that's posted here because there are very specific requirements for how IES defines at-risk. Applicants who are proposing to study children at-risk for developing disabilities must make this determination of at-risk status on an individual child basis and may include, for example, factors used for moving children to higher care in a response to intervention model, or instruments used for individual child screening. Evidence consisting only of general population characteristics, such as labeling students or children to be at-risk for disability simply because they're from low income families or English language learners is not sufficient. So, I encourage you to read the definition that is online and certainly follow-up with me to answer any questions that you have about working or addressing students with at-risk for disability.

IES also identifies the student education outcomes by the age or the grade range or the education setting that the students are in. In the pre-kindergarten group, the required outcomes are primarily those that would be considered school readiness. So, pre-reading, pre-language, vocabulary, early math and science and any of those social and behavioral competencies are necessary to engage in the schooling process. If you're interested in pre-kindergarten, a research must address center-based pre-K programs.

From kindergarten to grade 12, there are two categories of student outcomes. One is the learning outcomes in the major academic subjects of reading, writing, math, and science. You could explore outcomes such as standardized test scores, end of course exams, exit exams, and grades.. You could examine progress in the education system, such as completing a course, completing a grade, being retained, graduating from high school, dropping out and progress through postsecondary education. Also included are the social skills, attitudes, and behaviors that support this type of learning and progress.

For postsecondary education, which is defined as baccalaureate and sub-baccalaureate programs, the student outcomes are access to, persistence in, progression through, and completion of the education program – be it a certificate or degree program. Additional outcomes can be used with students in developmental or remedial education programs including achievements in reading, writing, English language proficiency and math. Adult education students are defined as those who are 16-years old or above and outside the K through 12 system. Adult education students are expected to be in such programs as adult basic education, adult secondary education, adult English as a second language, or high school equivalency test preparation such as the GED program.. For these students, the student education outcomes include achievement in English, math, and English as a Second Language as well as access to, progress through, and completion of their adult education program.

If the student population you are studying includes students with or at risk of disabilities, you are encouraged to also include outcomes accepted under National Center for Special Education Research Grants Program. Those outcomes include developmental outcomes for younger students, e.g, cognitive, communicative, linguistic -- social, emotional or physical developmental

outcomes for young students. . For older students, these include functional outcomes that we think of that improve educational results and the transition to employment, independent living, or postsecondary education for students with disabilities.

The second requirement for all applications is that the application must come from a partnership. At a minimum, the partnership must include one research institution and one education agency. The Institute does not endorse a specific model of research partnership, however, the Institute views research partnerships as going beyond two very typical and common forms of cooperation between research institutions and educational agencies that we often see. The first one being where a researcher is hired by an education agency to perform a specific research service and reports the results to the agency. Second, you may see where a researcher has an initial research interest, gets permission from the agency to carry out that research within agency schools, but that's the extent of that partnership. The Institute expects you to go beyond these two types of researcher-practitioner cooperation to a model where both institutions are actively involved in establishing the research and working together throughout the evaluations.

Again, each institution must contribute at least one on principal investigator, one PI. Each can contribute other co-principal investigators or co-investigators. Overall, IES grants have one PI/Project Director so you will have to decide which of the PIs will be the grant PI (the other will be a co-PI) The PI from the research institution should have experience in the program or policy being evaluated and in evaluation work.. The PI from the education agency must have decision-making authority for the program authority – they don't have to be the superintendent (or someone with final decision-making authority) but they must have some oversight for the implementation of the program or policy. The application should include a joint letter of agreement from the two or more organizations setting up the partnership. You can think of this as a memorandum of understanding showing that both sides agree to the work proposed. Other members of the partnership can also provide separate letters of agreement agreeing to their roles in the evaluation.

The partnership may be new or existing. If there is an existing partnership and there's been previous collaboration, the partnership will be stronger if you include the same personnel as have been involved in this prior collaboration (not only the same institutions), in order to show continuity of personnel.

Let's discuss what types of organizations are eligible to be partners. The definition of a research institute is very broad. To be eligible to be the research institution partner, you have to show that the institution has the ability to do this type of research. And so, it can include non-profits and for-profit organizations. It could include public and private institutions and agencies, such as colleges and universities or research firms.

Eligibility to serve as the state or local education agency partner is a bit more complicated. This slide describes possible state education partners. The examples presented are all eligible as they are overseeing some aspect of education- early learning, elementary, secondary, post-secondary, higher education, or adult education. States are not uniform in the types of public agencies they establish to oversee a specific area of education so you may want to contact a program officer if your state has a unique agency involved in a particular level of education. Also, if there is a

separate agency looking at a particular level of education but the main education is also involved, you may want to include both. Further, if you are looking at a student transition that includes two agencies (e.g., from pre-K to K-12 from K-12 to postsecondary, from education in the juvenile justice system to the regular K-12 system), it can be helpful to include both agencies. For this grant program, IES uses a broader definition of education agency than the Elementary and Secondary Education Act, which focuses on the primary agency responsible for supervision of public elementary and secondary school.

There are a number of organizations that can serve as a local education agency partner. The main local education agencies are public school districts. In addition, there are community college district, tribal education agencies, state and city postsecondary education systems, and the adult education providers as defined under the Workforce Innovation and Opportunity Act can also serve as a partner when there's no state or local education agency for adult education. For this grant program, LEAs do not include non-public organizations that oversee or administer school, such as education management organizations or charter management organizations. These groups can be additional partners, but you must include the public state or district agency that oversees the schools that are involved.

There are some very small local education agencies that contain one school. Those can apply as a partner, but the peer reviewers may consider that work less significant than projects that involve multiple schools.

You can also have additional partners. Partnerships may include additional partners if you think they will increase the quality of the research. So, you'll want to show that any additional partners that you bring on have similar interests. For example several education agencies may use a similar program or policy so it makes sense to include multiple education agencies and it may also increase the significance of your proposed work.

Non-education state and local agencies may be useful partners as long as the education agency is also a partner. If you're looking at a specific population, say you're interested in foster children, you would have your education agency, but you might also want to consider having your social service agency that works with foster children involved. So, for example, there's a project that combines Connecticut Department of Education and the Connecticut Department of Justice and the Yale University Center on Children – together they're examining the education of court-involved youth. You can also include more than one research institution. Again, the point is to argue that they have shared interests and they will make unique contributions to the work.

Non-research organizations are often part of the partnership such as interest or community groups. They may have strong their interest in the issue. They may collect useful data for the evaluation. They're may provide services or have critical links to the community and other organizations. For example, there's a project that's seeking to link early childhood education with the kindergarten system. It has United Way as a partner because of the United Way role in supporting early childhood education in that district and its links to key stakeholders.

Keep in mind though, as you have more partners, the coordination becomes more difficult. So, it's important to address how that partnership will work together, how you'll keep in contact and

importantly, how the partnership is going to make decisions. One type of partnership that we would not recommend is the inclusion of multiple education agencies where the only similarity is that they have worked with the same research institution, rather than share a program or policy.

The third requirement is that you need evaluate an SEA or LEA program or policy. This program or policy should be of high importance to the education agency. It should represent a substantial modification of the agency's practice or other agency's practice. So, you're evaluating something that goes beyond the existing practice. The program or policy should have been adopted by the SEA or LEA and that they have the final authority over its use. It should not be a researcher's intervention that the agency is allowing to be tried out, for example, in the district or the state. We have other grant programs that would support that that kind of work. And the program or policy is implemented as everyday practice. So, the district isn't receiving additional support for implementing the program or policy or pilot testing with a limited population.

We have other grant mechanisms within the National Center for Education Research and the National Center for Special Education Research that can support work that doesn't fit within the State/Local Evaluation grants. So check your research idea against the requirements of the State/Local Evaluation requirements. First, if you're not looking at student outcome, then IES in general, is not the appropriate funding agency for your research. No matter what issue it is that you want to address, no matter what level, you must be able to address student outcomes. Second, if you more time and effort to build a partnership and carry out some initial research before you're ready to do a full evaluation, then we would have you consider the Researcher-Practitioner Partnerships grants, which is the other topic under the 84.305H Request for Applications.

If the program or policy you want to evaluate is not implemented by an SEA or LEA, or it's not of high importance to the SEA or LEA, we would want you to consider one of the two main research grant programs - the Education Research Grant Program (84.305A) or the Special Education Research Grant Program (84.324A) that could provide grant money to for such evaluations. In addition, if a program or policy needs further development work before being fully implemented and evaluated, these two main research grant programs have funding for such work. Please talk to the program officers so we can help you figure out the best fit for your research idea.

To restate: the purpose of State/Local Evaluation is to carry out a joint evaluation by a partnership of a research institution and a state or local education agency of a specific education program or policy of key importance to the education agency. It is critical to identify the specific education program or policy you intend to evaluate in your application- do not propose that the partnership will determine what programs or policies it will evaluate after receiving the grant. The evaluation should contribute directly to state or district decision-making, -how the state or district will use the findings from the evaluation to decide on changes in the program or policy, e.g., reform the program, revise it, adapt it, broaden it perhaps, or in other cases, scale it back or reduce it based on the evaluation results.

State/Local Evaluation projects are also to foster longer term research partnerships. So, as an example, the state and local project in Michigan includes the Michigan Department of Education with the University of Michigan and Michigan State University. Midway through the grant, they decided that the partnership was beneficial to everyone, so they created a whole new organization called the Michigan Consortium for Education Research and that consortium is developing new research ideas and identifying other funding sources to support work of interest to all the parties.

To restate what the partnership should do during the grant: First, before applying the partnership should identify an education program or policy implemented by a state or local education agency that is of high priority to the agency and intended to improve student education outcomes. Second, during the project the partnership will evaluate the overall impact on student education outcomes, impacts on important sub-groups, as well as any intermediate impacts that are of importance. Also, the partnership must check on the fidelity of the implementation of the program or policy. Findings from this check are often of great interest to the agency partner and they can use these findings to revise implementation as needed. The partnership is also to determine the costs of the program.

Past projects have covered a wide range of programs and policies. There is no particular area to focus on for to the State/Local Evaluation grants. IES has made 22 awards to evaluate programs and policies that include expanded pre-kindergarten programs, professional development programs for pre-kindergarten teachers, gifted education programming and policy, dual-immersion language instruction, student retention policy, increasing academic high school curriculum requirements, high school exit exams, high school diagnostic testing for college readiness, mentoring for college enrollment.

To repeat, the focus of the State/Local Evaluation grants is not on specific areas of education, but on programs and policies of high concern to the district or the state and the results from which can then be used for the agency to make further decisions in this topic area. What are the expected products of the grant? We want you to come away with causal evidence of the impact of a clearly specified program or policy implemented by that state or local education agency, including the overall impact, impacts under a variety of conditions, and impacts for subgroups. We want to see conclusions or revisions on the theory of change that guide the program or policy which will contribute to our theoretical understanding of education processes and procedures.

If the evaluation finds a beneficial impact, then the evidence will support continuation of the program or policy and also perhaps promote its use by other districts or states. If you do find that the program is beneficial, the project should identify the tools, procedures, and support used to foster full implementation in order to help others adopt it. On the other hand, you might not find a beneficial impact. So, if that happens, you will want to examine the possible reasons why not beneficial impact was found – was it a failure of implementation, something similar occurring in the comparison group, or did the program or policy not function as theorized (e.g., expected intermediate outcomes did not occur). These findings might lead to potential revisions of the program or scaling back of its use.

Under a State/Local Evaluation grant, you can give feedback to the state or the district as you're doing the evaluation. The purpose is to both learn about the intervention, but also to inform the state or the district so that improvements can take place. Such information, which is provided early in the grant, is an important part of your dissemination. Also important to dissemination are the impact findings and the financial costs of the program and policy as these are critical for the agency's decision-making regarding the program or policy as well as other agencies' decision-making to adopt it.

Now, I'm going to move on to the project narrative on what goes into the actual application and the narrative of your research.

Allen Ruby:

Hi, this is Allen Ruby and I'll be discussing the project narrative portion of your application which is the substantive part of your application. It has a 25-page maximum. It is what the peer reviewers focus on and it contains these five sections, Significance, Partnership, Research Plan, Personnel and Resources. We'll just walk through each one.

The purpose of the Significance section is to justify why this evaluation should be done. You're asking for a grant and an amount of time to do this work and in this section, you make the argument why it needs to be done. So, the first thing you need to do is describe what the actual education program or policy is that you want to evaluate. What are its components, what are its processes, evidence that is either to be implemented or has been implemented in the past and will continue to be done and how it differs from existing practices in the same location or in different locations.

After you've described the program or policy so that people understand what it is you're going to evaluate, you want to also address its implementation. As Jackie spoke several times, it has to be implemented by a state or local education agency. This grant program is not to evaluate a researcher-developed an intervention and the state wants to see does it work. For that type of evaluation you should go to Education Research Grants or Special Education Research grants programs. Under this grant program, the state or the district has agreed to implement the program but may not have the funds or expertise to evaluate it. If the education agency is going to start this program next year, that's fine as well, as long as it starts by the time the grant is awarded. But then you'll also have to provide evidence that implementation will actually take place. For example, identify dedicated funding streams for implementation or legal authority to start the program, identify the office actually set up to actually implement the intervention and who's actually going to implement it. Also, the program or policy should be rolled out under routine conditions.

You should provide a theory of change. How is this program or policy theorized to lead to the final outcome of improving student outcomes? It may go through several steps to get there. It may be an intervention to improve principals' roles at school, which is then to improve teachers' instruction, which will then improve student outcomes. You want to map this out showing the intermediate outcomes in the process. And you may want to diagram this using a logic model.

This is often a very quick and easy way for the reviewers to see how the intervention is supposed to look and act.

You also want to give the rationale for testing the impact of this intervention. It may be a very widely used intervention in the state or district, but nobody's ever taken a look at its impact. It may be because you have an ongoing intervention that's worked in the past, but the agency has added a new component to it or is implementing it in a new way, so you want to test if that the change has a beneficial impact on students. It may have been evaluated before, but now there's an opportunity to use a much more rigorous design to test it. Or there may be a very strong theoretical justification for why it should work, although the empirical work hasn't been done yet. Maybe the state or district is trying something so new that other districts or states would be interested in seeing does it really work. Or maybe this district or state is using a widely used intervention that everybody else is using and it's just accepted, but it's never been tested. The theory of change and the rationale are probably not going to be as well developed as if you were evaluating a researcher developed intervention because oftentimes states and districts adopt interventions, policies, and programs with less theoretical justification for them. But there still should be a simplified theory of change, indicating why this intervention should work.

Next is the Partnership section. You are to apply as a partnership of a research institution and a state or local education agency, so there's a separate section where you discuss the partnership. First, you should describe the partners. Who are the research institution and the educational agency? Are there other members of the partnership taking part and why are they included? What are their common interests? How did the partnership come together? What was the decision process to propose the state and local evaluation project? Did it grow out of past work? Or did the education agency need a research partner to evaluate a new program or policy. Do the partners have a management structure and procedures in place to keep the project on track and ensure the quality of the research? The partnership grants are to give education agencies a greater role and voice in deciding what research is done and how that research is done. This should be reflected in the management structure and the procedures used from the project. These are not projects where the agency says, "Okay, evaluate this program that we're implementing, then come back and tell us what the findings are." You can do that under our regular grant program. Under a partnership grant, there should be a structure where the researchers and the agency are constantly coming together to discuss how is implementation of the evaluation going? Do changes have to be made? We don't want to go four or five years down the line until the agency gets feedback. So, there should be feedback very early on that the agency can take advantage of and that the researchers can learn things as they go along. Also, agencies sometimes change interventions over time and the researchers need to be kept up to date on any changes which may impact the evaluation.

One major issue is data, data sharing and data housing. Who's going to have access to data? How is it going to be kept secure? It's probably going to be partly confidential data. You have student level data and school level data, so it's important to show that data storage and sharing decisions have been jointly made and that the agency will have easy access to the data even after the project ends. There are opportunities here. You can place data managers and researchers in the agency. If the agency doesn't have a strong data management ability, so you can build a capacity of the agency to either collect the data or to maintain, to clean the data and to merge the

data. And many of these projects expend an time doing just that for those agencies that don't have that capacity.

There's also an expectation that there will be some agency capacity building to use research coming out of this partnership. How to use research and gain an understanding on how research is done and I want to be very clear on those two points. That doesn't mean that the agencies have to learn how to do the research itself. That can be up to the agency's decision. We know local and state agencies vary quite a bit in their research capacity. Some have very little ability. Some have very strong research offices. So, we want the agency to identify what is its interest in capacity building and then the project can build in some general and some specific efforts to address those.

As agencies should be very involved in the decision-making process should help the agency learn more about how research is done and how it can be used. The actual research design should be well understood by the agency. The agency should understand the value of this type of design, how it's done, and the reliability and validity of the evidence provided from it. There also may be specific skills an agency will want to learn. An agency may tell you, "We're not very good at doing surveys. As we do this project, we're going to have a survey in the project and we'd like you to train some of our people in survey implementation." And that's fine. Sometimes the agency will say, "Well, we actually need help in the data management and generating our reports from that data. We've got these massive data sets for the whole district, but we aren't generating school-level reports," and the research institution can take on some of those roles to help the agency carry out work it's not capable of doing.

The peer reviewers take Partnership very seriously. One of the major criticisms they give of these grant applications, is it doesn't look like a partnership. It looks like a research grant. The work is being driven by the researcher and they're just using the district to gather the data. The district is seen as a subject rather than a partner. If that's the approach, please don't apply here because you're not likely to fit well-reviewed.

So let's move onto the Research Plan now. Here's where you describe the actual evaluation you'll be doing. Open it up with your research questions or hypothesis. You may have described these earlier in the Significance section to argue what the significance of the work you're doing, but it's good to restate them here because that reminds the peer reviewers of specific questions you want to ask, and then you can map out how your research plan will lead to answering these research questions.

So we're going to walk through this set of issues that should be addressed in your research plan. You want to define the population you're looking at and how your sample and sampling procedures will allow you to generalize to that population. You want to be looking at a population that's important enough for the district to be concerned about the results for. If you're excluding any groups, you want set out your exclusion rules and justify them.

You do want to talk about how participation will be maintained and attrition will be reduced. Causal designs are threatened when attrition gets too large, and so you want to give expected

response rates and how you expect to obtain them. If you claim an expected 85% response rate to surveys, you want to justify how you're going to obtain that rate.

Next, let's discuss the research design. Again, you're doing a full evaluation here, so you want to be able to make causal conclusions. So, you want a very strong research design. If there are any weaknesses, openly discuss them. The peer review panel will have people who understand different designs, and if you skip over an obvious weakness, that will count against you because they'll think you don't understand your design because you didn't realize a weakness. So openly discuss it, and talk about how you may be able to offset it to some degree. You want to have a baseline comparison from the start so that we know the intervention and comparison groups are similar as they start out. You want to check for any biases that may come from attrition, either overall attrition or differential attrition. You want to know your counterfactual. So if you don't know it now, I mean, you can argue with the business it's usual counterfactual, but you are going to want to address what is actually going on with counterfactual because that's very important for understanding why you don't find impacts.

By doing all these things, you have the potential to meet the What Works Evidence Clearinghouse evidence standards. The reports based on IES-funded evaluations, such as the State/Local Evaluation grants, are submitted to the What Works Clearinghouse, and they evaluate the quality of the evidence provided by the designer's presentation. So it is worthwhile to take a look at the site on the What Works Evidence Standards for Design. The preferred design is the randomized controlled trial. If you propose a RCT, you should note the unit of randomization and justify the choice of that unit. You want to describe the process you'll use for random assignment and then how you'll maintain its integrity. And what we've learned from these projects is that one good way to help maintain integrity is, very early on in the project, talk with the education agency, talk with the folks on the ground who are implementing the program or policy, about the research design -- what practices should be followed in order to maintain the integrity of the design.

There are many different ways to approach a randomized control trial, and the different ways often lead to other issues you need to keep in mind. In some places, everyone's just included into the process. You'll have -- everybody takes part, and then they're randomized in to the control group, so it's a mandatory participation. What sometimes happens here, though, is that there are schools, teachers, or students who don't want to take part. So you run into problems of treatment fidelity, so it's important to think about ways to maintain fidelity.

Another approach is to only include volunteers in the RCT, where schools or teachers want to take part. For example a state or district may say, "We have enough money for 20 schools to receive the program this year so we'll randomly accept half of those who apply and use the other half as the control group." Here, you kind of get the opposite problem. The volunteers want the intervention. If they don't get in, they may not stay in the control group. They may leave the control group or they may seek an alternate source of the intervention they can get on their own. This happened in a large scale study we did on social and character development where -- this is why it's so important to look at comparison group practice. What we found is that the -- over 70 percent of the teachers who were in the control group were implementing alternative materials similar to the treatment programs. So it was not a surprise when the interventions were not

found to be effective. There was little difference in the level of implementation in the treatment and control schools. One approach to this problem is to provide something else for the control group. For example, a study looking at the impact of literacy PD in OH preschools offered the control group PD in a different subject area.

Lotteries can be used, and the lottery applicants, of course, want to be in. So you have the same issue with volunteers. They may leave the study to seek the intervention at another school. You'll often up with differential attrition and it is not possible to follow the leavers, and sometimes the less motivated are left behind in the control group. There may not be alternatives you can offer but there may be ways that you can follow some of the leavers.. For example, a study in Portland looking at dual immersion schools that selected applicants by lottery was able to follow some of the control-group leavers who had left the district by using state administrative data.

A staggered rollout can be used when everyone is to get the intervention but the agency cannot provide it in the same year but promises it in a future year. For example, a project studied the use of formative assessments in Indiana schools under a design in which the state decided it could pay for the program for all schools over multiple years. Schools applied to receive the program and were randomly selected to receive it that year or the next year with the next year group becoming the control group. The major issue that comes up is whether 1-year is enough time to implement the program with fidelity or is more time needed. If it can be fully implemented in one year then it's a good design, but for some interventions two to three years are really needed to get the intervention up to a level of implementation we'd expect to have an impact off of outcomes. And such a staggered rollout may be giving an unfair evaluation of the intervention because it will never be able to evaluate it at a time when that intervention's working at a high enough level expected to have impacts on students..

Another approach is to look at a variation of the program or the policy. Here, everybody is getting something, but the treatment group is getting a new component or some other addition to the program or policy. The question is whether the new component or new way of implementing it expected to have an impact? Is it large enough to have an impact? If it is then that's an interesting study for to the agency to do. For example, an ongoing study in Tennessee's statewide teacher evaluation program is examining the impact of a teacher mentoring program which is being rolled out as a new component.

One issue that education agencies have with RCTs is that they are most interested in making sure the maximum number of students' receive the new program or policy. One approach we've seen used is the randomization of a small third group that is not included in the study. Instead, if treatment students dropout, members of this third group take their place and receive the intervention. In this way, the maximum number of students receive the intervention but the design is not compromised as it would if control group students were allowed to take the treatment.

If you can't use an RCT, you should justify why, and there are good reasons why you can't do one (maybe everyone is receiving the intervention). So you will propose using a quasi-experimental design. As you do so, you should discuss how you will minimize or model the

selection biases. Regression discontinuity designs (RDDs) are often proposed under the State/Local Evaluation grants. When you propose an RDD, you should discuss whether the assignment variable can be manipulated, because that will threaten the design. For example, an evaluation of gifted education under this program found that IQ test scores were being manipulated for higher SES students, but not for lower SES students. Project using state standardized tests find less evidence of manipulation. You may also consider other well-designed quasi experimental studies. Some comparative interrupted time series designs may be considered strong enough. However, your typical matching design that match using only the most commonly available variables, such as ethnicity or race ethnicity or free and reduced lunch are often not considered strong enough by the peer reviewers because they're not seen as being able to control for the unobserved factors that may be part of the selection.

When using a quasi-experimental, there may be a need to work with the education agency in understanding the design and implementing it correctly. For example, when using a regression discontinuity design, it's very important to work with the agency and the schools on student assignment so that students above or below the cut-off did not get into the intervention. For a State/Local Evaluation project in Oregon looking at middle-school literacy interventions, the methodologist went into the school to discuss what an RDD design means and why the cut-off score had to be respected. In addition, the project allowed schools to assign about 5% of students based on its own determination outside the cut-off point. Students who teachers felt would benefit or not from the program were assigned accordingly and were not included in the study. This approach helped to maintain school support for assignment using the cut-off score for the majority of the students.

Statistical power is very important. If you have a high enough power, it avoids the issue that if you don't find a significant result it's not because your sample size was small. You should describe your power analysis in detail and justify the method you used to calculate the power, and include all your assumptions. There are reviewers who will redo your power calculation, and if they can report on in panel that, "Yes, I got the same results as here," that raises confidence in you application, especially in contract to a comment that, "I can't replicate this power now." That raises huge concerns for the panel.

You should provide a power analysis for the main impact analysis and for any confirmatory subgroup analysis you're doing. If you provide a minimum detectable effect size, you should discuss the practical meaning of it. Is such a change important? Now, if you're reporting a high minimum detectable effect size, there's immediate reaction by some of the reviewers that that's not possible to obtain. So if you have a reasonable practical impact, then they say, "Oh, in that case that's not so unreasonable." And conversely, if you report a very low minimal detectable effect size, the initial reaction is that it's so low, does it have any meaning, or are you just way overpowering the study, and spending too much money on your sample? But if there's a practical meaning to that effect the peer reviewers are going to say, "Well, that's reasonable to have such a large sample"

Your outcome measures should be outcome measures that are highly relevant to the agency - the districts and schools and teachers. So these are often ones found in administrative data, very much like the ones Jackie talked about early on; measures of achievement such as test scores,

measures of progress through the school system such as graduation, and social and behavioral outcomes such as attendance, disciplinary referrals or special education placements. You can use researcher-developed measures as well, if you think they're going to help you understand how the intervention works or does not work correctly, but these should not be the primary outcome measure. You want to address the reliability of your outcome measures. And you want to show how all your measures fit well into the theory of change. These are measures of the outcomes that your theory of change has predicted will be affected by the program or policy.

You should have measures of any intermediate outcomes that the program or policy is expected to create and that are to lead to the final student outcomes. An example is an evaluation of a 12th-grade course in Florida to help students who tested not college ready in 11th grade be prepared to enter college without needed remediation coursework. The project examined this program when it was voluntary and when it was mandatory. Looking at the voluntary program time, they found there was almost no 12<sup>th</sup> grade course-taking going on at many schools. So, we wouldn't expect to see any impact on students' preparedness to go to college if they weren't taking the course from this intervention. So it's very important to determine whether the intermediate outcomes expected are found.

You want to describe how you will examine any important moderators and mediators. You may want to examine variation in outcomes, e.g., by subgroup, by site. Some of these may be identified in the theory of change but others may be identified from your analysis. For example, a study in Massachusetts looking at high school graduation tests found that they had a negative impact on high school completion for students who were from low SES families. These students were less likely to retake the test, versus higher SES students making them less likely to complete high school and less likely to go on to college. As you describe your mediator and moderator analyses, you will want to differentiate which are confirmatory and which may be more exploratory.

IES supports the use of mixed methods in the evaluations it support. We define mixed methods as a combination of quantitative and qualitative work. Mixed methods can help identify possible reasons for why impacts are found or not found as well as providing qualitative findings explaining the intervention and how it is implemented.

We've supported a number of useful evaluations that have found impacts for sub-groups that were not in the original theory of change. For example, one found an impact on girls in high school but not on boys. Another one found an impact on Hispanic girls. What happens often, then, if there's no follow-up work built in to look at why different sub-groups have different outcomes, it becomes a matter of speculation. It may be, –speculation that boys have other labor market opportunities than girls.. But at that point we're not sure and we need further research to determine if it's correct. We would prefer that some of that research was done under this grant program to understand any relationship is found, so building in follow-up research to do that kind of work after your initial finding is a helpful. Earlier we also mentioned the need to understand what level of fidelity of implementation occurred and what the comparison group practice was in order to understand the impacts of the program or policy. So, we encourage you to create a team that can do both types of work can interact together.

I think another interesting example of how this can have an impact very early on is a fairly new study looking at dual credit in high school math courses in TN. There was skepticism at the agency level that the two-day professional development program for the dual language math teachers, would have an impact on what was actually taught in the classroom.. The project did a survey asking teachers to list their syllabus and what they were teaching at specific times of the semester, and they found a large difference between the amount of time spent on specific topics in the treatment and control groups with what was being covered in the treatment group was the material to be covered in order to receive college credit. And that has changed how the agency views the project.

You want to describe what your fidelity measures are, capture the core component of the program which you described earlier in the Significance section, and are there -- what are their psychometric properties? Showing adequate fidelity of implementation is key to understanding if the overall evaluation is truly measuring the impact of the program or policy as expected or that it has revealed a failure to adequately implement the intervention. You may need to develop decent measures of fidelity before doing the evaluation. You should discuss how the fidelity data will be analyzed, and will it be included in the final analysis. You might want to do a treatment on the treated analysis that includes fidelity. And, as I mentioned, we have found that many of the agencies find the initial fidelity study to provide some of the most useful information they get, and that it helps to build the partnership. You often can't give evaluation results right away, but you can give fidelity results. And when these are useful, the agency is more interested in continuing the project.

For example, an evaluation of voluntary pre-K in Tennessee found large variation in quality of instruction which has shifted the state's policy discussion from access to quality of instruction.

Understanding comparison group practice is also important to understanding the impacts of a program or policy. If there is a clear treatment contrast, there is greater confidence in the findings.

Next, you'll lay out your analysis plan - how you actually look at the data. You want to link your analysis back to your research questions. How does this analysis answer those research questions? For the quantitative work, you want to describe the models you'll be using. It's helpful to write out the models you'll be estimating, explain what each of your variables and coefficients mean and which are the most important. For the models you should address how cluster of date (e.g., students in schools) will be handled.

You should describe the software you'll be using for the qualitative, again, matches the index, summarize and interpret the data, and how are you using the qualitative and qualitative? Are they going to be for separate complementary analysis, or combined analysis?

You should also provide separate analysis plan, for the examination of mediators and moderators, and fidelity of implementation, and comparison group practice? And finally, you should describe how you will check the effects of attrition, overall and differential on your results.

Finally, you should describe how you will do a cost analysis where you document the financial cost of the program that you're implementing (you do not have to provide the costs of the program in your application but describe how you will obtain it). It should be detailed enough for another state or local education agency to use. If you want, you can do a cost-effectiveness or cost benefit analysis, but that's not required.

Let me just say a few words on the Personnel section. What you want to do in the Personnel section is say, "I have this person who can do this job that I've laid out in the Research Plan. They have these qualifications to do this work. They're going to be on the project for X percent of the time, which is enough time to do it. You want to target your discussion of each person to the work they're doing on this project.

Again, if someone has worked in similar types of partnerships, that's helpful to note. From the IES perspective this is a very large, often complicated grant. The peer reviewers are going to expect that a PI has experience in managing such grants. Also, you want to be ensuring objectivity of the evaluation. The evaluator should not be the same as the implementer of the program or policy. You may have implementers on this project doing things such as looking at fidelity of implementation or comparison group practice, but they should not be doing the impact evaluation.

For Resources, you want to describe the resources all the institutions are providing, what's going to be available, and that you have the capacity to manage the grant. The joint letter of agreement is considered a major sign that the resources are in place. If districts and schools are taking part, how are you going to document that they're willing to do this? And we often get applications where the district says, "Yes, we're on board," but the schools say, "We don't have time to do this – the state or district requires us to do some other program." So the more you can show that every level is interested in the study, the stronger the justification for resources.

If you're using administrative data, we want a letter from the organization that holds that data saying it will be released with that study. Now, we know some state agencies aren't going to provide that letter until the grant is given, but you should at least get a letter from them saying, "We will give the data if the grant is received and the project follows the process obtain the project." But there will be concerns if their data-holder's support for the project is not even discussed.

Regarding dissemination, these grants are expected to provide causal results with high interest not only within the district or state but across the country. So we are expecting broad dissemination of results. You should discuss what resources are available for dissemination and how you're going to carry out the dissemination to the different audiences.

Now, I'll just go through some of the last sections of the application. You have four appendices and a budget plus a budget narrative. Appendix A is used if you are resubmitting an application. You have three pages to discuss how you responded to reviewer comments. That's the only thing you should put here, and it's very important to address those comments because one of your reviewers is probably going to have been a reviewer in the past.

Appendix B is where you can put tables or charts with additional information on the evaluation or the partnership. You may give a timeline of the project. You show the management structure of the partnership. You may give examples of the measures you'll be using in the evaluation (e.g., tests, survey items).

Appendix C is where you can put materials from the program or policy – e.g., curriculum material, the language of the policy. The difference between Appendix B and Appendix C is that Appendix B materials are in regards to the evaluation while Appendix C materials are in regards to the program or policy being evaluated.

And Appendix D has no page limit. This is where you provide the joint letter of agreement from the partners, any other supporting letters from other organizations involved, and letter from consultants. The more specific these letters are, the more confidence the reviewers have that the partnership understands its roles and responsibilities. This is where you would provide just a letter from whoever has the data, if you're using any administrative data, saying that that data will be provided to the project.

The maximum project length is five years, and we have a maximum award of \$5 million. The funds can only be used for evaluation purposes. You can't use them for implementation of the program or policy. , You should ask for a budget based on the size or the scope of the project (not the maximum award), so you'll include a detailed budget form for the overall study and for any of the sub-awardees, and a budget narrative to describe how you intend to use the funds.

You must submit your application by August 4th, 4:30 and zero seconds p.m. Please submit a few days earlier to avoid problems with an overloaded server on August 4 and so that you have time to respond to possible error messages regarding your submission. Your project can start anytime from July 1st to December 1st. Thank you very much, and that's the end of today's webinar. We look forward to hearing from you.

Jackie Buckley:  
Thank you.

Alan Ruby:  
--

[end of transcript]