

Making Research Matter for the SEA

Nathaniel Schwartz
Tennessee Department of Education

November 2015

© 2015 Edvance Research, Inc.

State education agencies (SEAs) have significantly expanded their reach in the past few decades, driven partly by standards-based reforms that make states the primary actor in school accountability and partly by educator evaluation initiatives that grant states new authority to define and measure instructional quality.

This broader state role means that the decisions of today's SEAs are more consequential and demand more thoughtful analysis than ever before. Yet few SEAs are meaningfully poised to respond to these demands. To date, 21 states have no defined research or analysis office; those that do vary widely in their capacity to take on strategic analysis.

This creates a difficult dynamic. The call for evidence is only growing, as is the recognition that data matters for strategic decision-making. But because state agencies often have little capacity for such work, research and evaluation tends to get outsourced or avoided altogether.

In 2012, the Tennessee Department of Education established an office of research and policy to coordinate internal and external research across the agency. Before this office was created, analytic capacity was scattered throughout the organization; no team had clear ownership over the work. The office's establishment heralded our state's intent to have outcome-focused research help drive organizational policy. But it also forced us to rethink how to conduct research and analysis within a state education agency.

Many state departments often view rigorous research as a nice-to-have supplement, but in the meantime, decisions must get made with or without definite evidence. To change this, agency officials must be able to expect in-time analyses that offer clear direction in decisions where the path is not already determined. Success cannot be defined by the research alone but also by the extent to which the research office changes the agency's trajectory in measurable and meaningful ways.

As we have moved down this path in Tennessee, we have faced choices about how to take on the work most effectively. This essay aims to make these decision points explicit. What defines the research agenda? What are the products? And what structures make it possible to achieve measurable results? The decisions we have made are not necessarily right for every agency, but the questions are ones that each SEA will likely face in prioritizing research-driven practice.

RESEARCH THAT COUNTS: CRAFTING AN AGENDA TO INFORM THE POLICY CYCLE

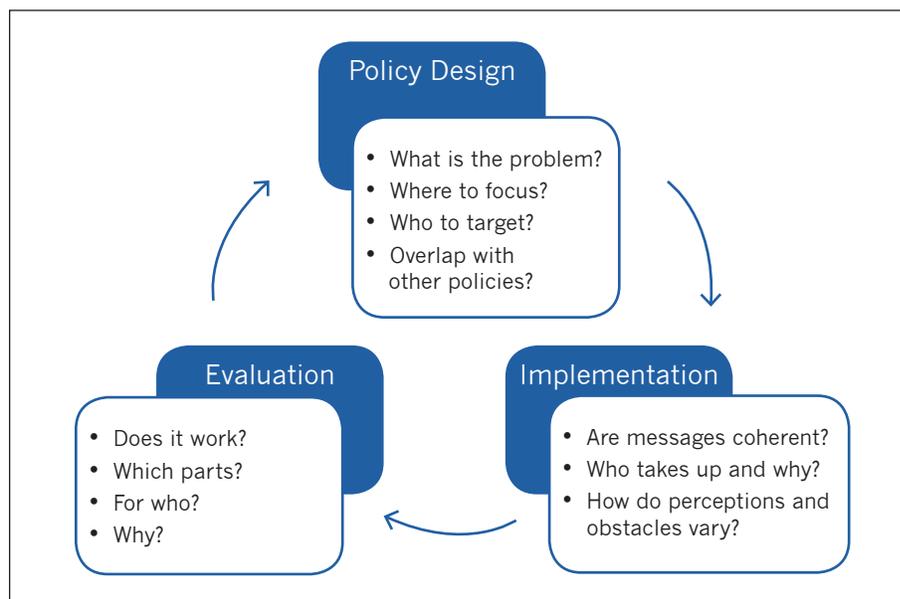
Researchers within the SEA confront an extravagance of data riches. Where most researchers struggle to gain access to a data set that might allow meaningful analysis, those of us in state departments stand at the base of a waterfall with data rushing past us from every angle.

Ever-expanding data and a broad array of state initiatives make for more research questions than even the most heavily staffed research office could hope to answer. This creates a continual dilemma about how to identify the highest priority topic areas and how to decide where we can get the greatest payoff for our research investment.

Policymakers and observers commonly call for ever-more program evaluation. Although we in Tennessee field internal requests for impact evaluations more frequently than any other demand, we do not think research can focus only on program evaluation if we want the research to drive policy. Such a single-minded focus severely limits research’s influence over practice, as explained further below.

Figure 1 represents a typical policy cycle within an education department. The cycle begins with policy design and moves through implementation and evaluation. But for most policies, the starting point is fairly irrelevant since the cycle constantly iterates among all three components.

Figure 1. How Research Fits in the Policy Cycle



When our research and policy team meets with department members about possible analyses, the discussion often jumps almost immediately to questions about program impact (represented by the “evaluation” box in Figure 1 above). This is understandable since these meetings are usually with officials who have launched a series of policy initiatives and want to know whether these initiatives have made a difference. But confining the research office’s work to the back end of a policy means research cannot inform policy and practice at other key junctures, such as policy design and implementation.

Impact analyses are necessarily backward-looking and often must take place considerably after the fact to be able to draw on outcomes such as state test scores and/or teacher observations. By the time an evaluation can be completed, a policy has often already been in place for several years and state officials are already thinking about an entirely new set of initiatives. Moreover, without careful setup on the front end of the policy roll-out, it can be difficult or impossible to determine the extent to which a change in outcomes represents a true causal effect of the policy.

In order to influence policy on the front end, research within a state department must focus equally—if not primarily—on the other two boxes in Figure 1: policy design and implementation (see Figures 2 and 3 for concrete examples of projects and outcomes aligned with each type of question). The sections below expand on this and offer concrete examples of some of our most successful projects at each stage of the policy cycle.

Forward-Looking Policy Design Analysis

Within our Tennessee office, we often refer to work on the design end as “landscape analysis.” By this, we mean analysis aimed at giving department personnel a clearer picture of a particular issue in order to drive initial policy efforts and pilot strategies.

For example, after the Tennessee legislature set aside a relatively small grant to fund pilot strategies aimed at increasing student success on Advanced Placement (AP) exams, we did a landscape analysis of AP testing across the state, focused on understanding the factors that had kept low the numbers of students who actually earned college credit through AP exams. We **identified trends in the type of obstacles students faced in different schools**; in some schools, AP-ready students were rarely placed into AP courses in the first place, while other schools did better at placement but had few students who chose to pay for and sit the AP exams. By analyzing the landscape, we could offer a framework for classifying schools that prompted a series of differentiated, small-scale pilots aimed at helping remove obstacles we identified. We are now evaluating these pilots and choosing which to expand in coming years.

Similarly, **a simple investigation into the typical math course progressions that students followed in different areas of the state** yielded the surprising fact that enrollment in 8th grade algebra I (versus 9th grade algebra I) had fallen over several years from nearly a third of the cohort to less than one-sixth, severely limiting the number of students reaching advanced math by 12th grade.

This study started as a design-focused analysis to inform curriculum decisions for our division of curriculum and instruction, but its results triggered an immediate shift in school and district accountability policies. And it prompted prolonged discussion with districts about how to ensure schools make course-placement decisions that push qualified students forward rather than hold them back.

Such design analyses primarily aim to better define the problem our SEA seeks to resolve. Done well, these analyses can sufficiently narrow the scale of the issue and make it concrete so the agency can ultimately make a real difference in the field. If students are not completing the prerequisites for advanced coursework, a statewide program aimed at training new calculus teachers will have little payoff. If schools are not counseling enough qualified students into AP classes, covering the cost of AP tests is unlikely to yield major benefits.

The best landscape analyses also create the conditions for further research on the effectiveness of department programs. A well-defined problem coupled with deep knowledge about the range of challenges and outcomes across the state makes it far easier to propose defined initiatives to solve the problem. These can then be rolled out and rigorously tested in comparison to other proposed alternatives. In our agency, our initial forays into data landscapes often wind up driving the creation of future programs. These data landscapes then serve as necessary precursors to most of the research that follows.

Ongoing Implementation Analysis

In the same way landscape analyses provide the foundation for policy design, implementation analyses serve as a crucial benchmark of progress on the path toward program effectiveness. Yet implementation research is the SEA's most difficult kind of research. State administrative data tends to offer little in the way of meaningful implementation feedback. And few state departments have the personnel to undertake prolonged, qualitative research.

While our research office is constrained in assessing implementation, we have found that we can offer valuable program development feedback through several relatively limited efforts that do not require the time and travel costs typically associated with deep, qualitative implementation studies.

Most importantly, we have benefitted from an effort launched several years ago with the Tennessee Consortium on Research, Evaluation and Development at Vanderbilt University to **annually survey all teachers and administrators in the state**. Reliable survey data allows us to monitor crucial changes in educator perceptions over time, as well as determine the extent to which the field is aware of and acting on state guidance around key initiatives. For example, the survey has tracked both the extent to which Tennessee’s teacher evaluation system is being faithfully implemented according to department guidelines, as well as how often teachers feel they receive feedback that helps them improve their practice. Data demonstrating success in the first realm and challenges in the second has helped push the department toward initiatives aimed at helping districts use evaluation as an improvement tool. Several efforts around teacher evaluation have been launched in ways that allow rigorous research on their effectiveness.

We have also reaped payoff from “take-up” studies. These simple, descriptive analyses map the extent to which department initiatives reach their desired audience. For instance, when Tennessee offered leadership trainings through its regional offices, department officials noted considerable statewide interest, but few realized these trainings wound up reaching principals or assistant principals in 80 percent of Tennessee schools. Coupled with stats on the take-up of teacher trainings as part of the Common Core transition, as well as teacher subscription rates to regional reading instruction courses, this data allowed our office to map the extent to which state messages had penetrated the school level.

Such studies (which we supplement, when possible, with the work of external researchers who have greater capacity for qualitative work) offer low-cost options to understand the ways state programs develop over time. Like the design-focused research described previously, these studies aim to bring research to bear on department policy long before program evaluation could provide results.

Backward-Looking Impact Analysis

Under the right circumstances, impact evaluations can be a smart use of department resources. But circumstances help dictate how helpful these evaluations can actually be. Below, we outline some considerations to keep in mind.

Impact evaluations are much more likely to yield valid measures of a program’s effect if they are planned as part of a policy roll-out rather than after the fact. Pilot programs can often provide ready-made opportunities for meaningful evaluation, assuming that someone is thinking about these opportunities during the pilot set up. For instance, in Tennessee the law

required us to create dual-credit courses, offering high school students the chance to receive college credit by passing statewide end-of-course exams. Early on, it made sense to limit some of these courses to a select set of schools. We gathered a pool of interested schools and randomly selected some from each region to create our initial pilot sample. This kept the initial pilot a manageable size and created a gold-standard research design that would allow us to directly measure the effects of the program on student outcomes (as we would have some schools with the program and some schools without, as a control). Recognizing the research design's strength, **the Institute of Education Sciences (IES) awarded \$2 million in 2014** to the project partners, including the Tennessee Department of Education, to support ongoing research.

When stakes are high, we recommend passing impact studies to trusted, independent researchers to ensure results are buffered from organizational politics. Evaluations are judgments of program effectiveness, and therefore have winners and losers. Even the most research-driven organization will struggle at times to come to terms with negative evaluations of popular programs. In the case of the dual-credit courses study described above, we enlisted the help of nationally known researchers at the University of Michigan and the University of North Carolina, whose work in this area ensures the final findings will receive serious consideration from both practitioners and academic researchers.

Although state departments administer many programs, usually only a handful are positioned for evaluation results to have an immediate impact on departmental strategy. In these cases, we advocate doing whatever it takes to allow rigorous research to weigh in on the strategy. For example, as part of its Race to the Top grant, Tennessee offered statewide summer teacher trainings through an innovative peer coaching system. As the grant wound down, department staff questioned whether the trainings—though popular—were worth future investment. Our team's **rigorous impact evaluation** found that teachers who attended the state trainings both received higher classroom observation scores and saw better student results on state tests. While the effects were not huge, they were enough to deliver a clear cost-benefit payoff, and they helped drive the decision to continue training teachers beyond the federal grant.

Choosing the right program evaluations to undertake means choosing only those questions for which the department has:

- A. the right set of decision points;
- B. the opportunity to apply appropriately rigorous methods, and;
- C. access to data on meaningfully aligned outcomes.

Meeting this criteria demands constant communication between research staff and department leadership, which in turn requires explicit attention to the relationship between the department's research and operational divisions.

RESEARCH THAT MAKES A MARK: DRAWING ACTION FROM RESULTS

Just as the agency's research agenda needs to speak to the agency's needs, so do research results. Asking the right questions and conducting the right analysis is only a piece of the puzzle. Too often, researchers see the analysis as the endpoint rather than the beginning.

What does it take to make findings meaningful for department leadership?

First, analyses must be done within the right time frame. This does not preclude a lag between project conception and results. Certain questions are likely to be meaningful for long-term department strategy; investing in long-term research on these topics can be worthwhile. Others will only catch the department's attention if they are linked to particular decision points. The key is to fit the research to the time-frame rather than the other way around. We have seen projects our researchers thought fairly uninteresting take on outsize influence simply because they appeared at a time when people were hungry for our answers. Similarly, we have done other projects that our office saw as hugely important and watched the work languish because we developed the findings at a time when no one was ready to take ownership over the results.

Second, research must be framed to speak to department policymakers. In other words, department personnel must be convinced that the findings relate directly to what they do.

One cannot assume that the questions will speak for themselves. From the outset, the research must be framed to explicitly suggest the kinds of "takeaways" that various department groups might draw from it. In other words, researchers bear the responsibility to explain why findings in one direction or another matter for department policy.

As findings develop, team leads within the department must be given time to process the findings and integrate them into their worldviews.¹ This can only happen through discussion and engagement. Again, it becomes incumbent on the research team to ensure that these department discussions are taking place—to ensure that they not just pass on findings but also launch a reflection on the results that will culminate in department action.

Third, presentation matters. Often, this gets reduced to a discussion about a report's length and the attractiveness of its graphics. In my job, I am frequently

reminded that department personnel have so much on their plates that they won't bother to read anything that is more than a few pages long. Though often true, the simplification around report length sidesteps a lot of other important considerations that help make research influential within a state education department.

After much trial and error, our office of research and policy has adopted several informal principles that guide the translation of initial research and analysis into a product that can drive policy. These are by no means revolutionary; they represent much of what people have said for years about actionable research. However, consciously abiding by the principles has pushed us to make our work both more digestible and more immediately impactful.

In our view, **presenting research to drive policy means constructing an actionable narrative, illustrated with meaningful figures, in order to build lasting frameworks that organize departmental thinking and policymaking within a particular area.** In the space below, I describe these principles in more detail.

A Clear Storyline: Policies often get driven by small-scale, anecdotal evidence. This can be frustrating, particularly for researchers, but it is not surprising when we consider that policymakers are necessarily in the business of communication and communication thrives on stories.

We have come to believe that making the findings deliver the same communication power as individual anecdotes is key. Research results must function as stories that lead to a clear conclusion. If it is impossible to put together a list of bulleted takeaways that elucidate a series of potential project implications, the researcher has failed to take the final crucial step to make the research relevant.

During our weekly research team meetings in Tennessee, we listen to research analysts present the studies they are working on and press them with pointed questions to succinctly convey what they are trying to say and what makes it meaningful. Through this process, we work to distill each project to the storyline that offers a clear reflection on department policy. Each eventual product comes with a clear set of driving ideas and implications. The point is not to force every product to only a single page of results, but rather to force every product to make clear takeaway points that convince both researchers and practitioners that the project has offered something concrete to facilitate department decision-making.

Meaningful Figures: This is not a call for sophisticated infographics. Every moment spent deciphering a graph's intricacies is a moment lost to understanding the graph's implications for the department's work. If

department staff are accustomed to stacked bar graphs created in Excel, there is no need for anything more. Eventually, as the audience expands, more complex graphic design is worth considering. But we have found that the initial pass should come in a language familiar to the audience.

An exception to the rule of sticking with what people initially know is when certain types of difficult-to-understand graphs or figures can add value to multiple projects. In these cases, it can be worth investing the time to help decipher them for the given audience and regularly include these types of graphs across presentations, allowing department personnel to become accustomed to their look and feel. One example from our work in Tennessee comes from the common observation—that is not specific to Tennessee—that statewide averages tend to mask considerable variation at the district and school levels. To make this picture visible, we present bar graphs, such as Figure 2, where each bar represents a single district (or school). The example figure shows variation in district-level teacher retention. This kind of figure allows easy comparison between individual districts and the state average, it raises a series of questions about the districts (or schools) that fall at the extreme ends of the x-axis—and it can easily be altered to include more information, as in Figure 3, which compares different districts’ ability to retain their most versus least effective teachers. Simply because of the number of bars, such graphs can be initially confusing to a new audience, however, we have used them frequently enough across presentations to make them immediately comprehensible to our audience within the Tennessee Department of Education.

Figure 2. Overall Retention Rates by District (All Teachers)

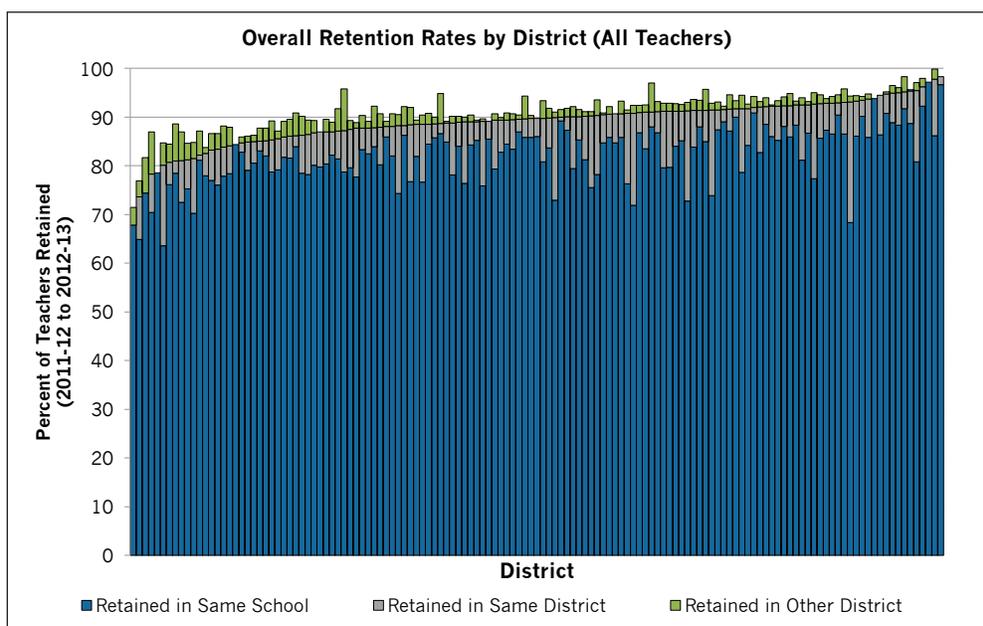
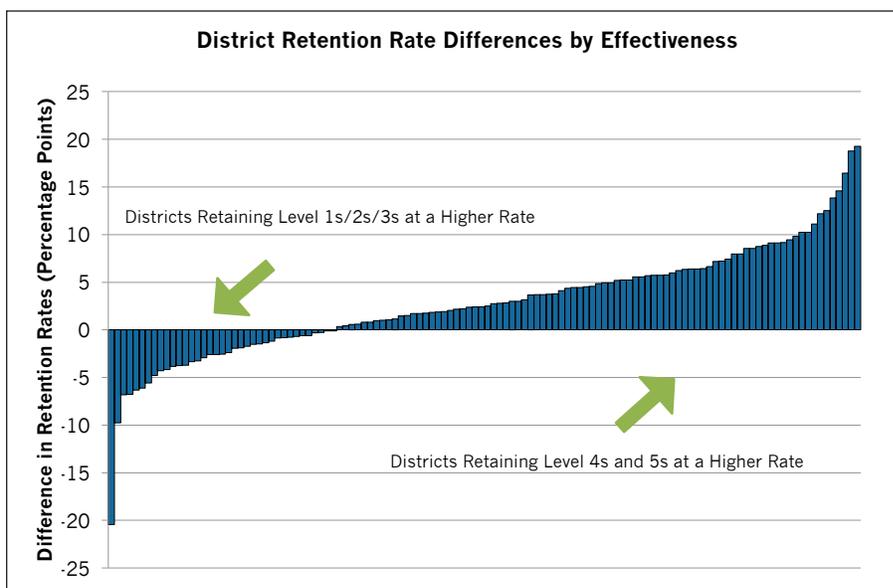


Figure 3. District Retention Rate Differences by Effectiveness



Frameworks for Anchoring Main Ideas: This principle has prompted considerable discussion across our research team, since the meaning of “framework” shifts from project to project. Yet increasingly we agree that the models and frameworks research can provide for approaching a policy dilemma often add as much value as the quantitative results that fit inside the research. In the best cases, the logic model offers organizing principles for defining an issue that can live on in people’s minds long after they have forgotten the exact numbers that accompanied the principles.

For example, the earlier noted landscape study of Advanced Placement testing across our state developed the idea of an AP-ready student (identified by 8th grade test-score results) and classified schools into buckets based on the types of AP obstacles their students encountered. “Low preparation schools” were high schools whose feeder schools never produced enough AP-ready students to justify AP course offerings. “Low access schools” had AP-ready students but no AP course offerings. “Differential enrollment schools” had low-income, AP-ready students enrolled in AP courses at far lower rates than non-low-income, AP-ready students.

The model included six buckets of schools in all, each of which defined an access problem in a particular way. What developed was both a concrete analysis—X percent of schools are low access schools, Y percent of schools are differential testing schools, etc.—and a broader way of thinking about targeted school supports. As we wrote in a sidebar to the AP report, “Too often, our policy solutions and interventions are crafted as one-size-fits-all policies that fail to differentiate based on the highly variable data coming from individual

schools and districts.” By creating a data-based diagnosis tool for school problems in the AP realm, the research work prompted a discussion on how we could use similar techniques in other circumstances to more accurately identify intervention needs.

Our AP intervention organizing model, and others like it, have taken on a life of their own after initial presentation, recurring in a host of different discussions and prompting different department groups to consider their own work in new ways—even outside of the initial project arena. The right models make individual projects generalizable, providing new vocabulary and new guidelines for confronting recurring policy problems.

STRUCTURING THE WORK

The agency has to commit to make research possible. And this requires supporting arrangements: personnel, resources, and management structures. What does this look like and what does it take to make such support available within an SEA bureaucracy?

Team Dynamics

It is entirely possible for a state department to prioritize research without consolidating its researchers into a single team, but the group structure impacts the type of work that can happen.

Before Tennessee moved to a single-office structure, department divisions carried out high-quality individual pieces of analysis, but no one was responsible for focusing on developing research to inform department-wide strategy. The old structure meant Tennessee lost the ability to deploy research as a surveying tool across programs and divisions. This in turn contributed to our department’s already silo-like nature.

The lack of a unified research office also translated into little to no oversight of the external research partnerships formed during this time. Research partnerships tended to look more like one-sided transfers where Tennessee provided data and researchers carried out their own analyses for their own purposes. Giving a single team oversight over the research agenda has made it far easier to seek out partnerships that answer departmental needs rather than the other way around.

At the same time, a separate research team brings its own challenges, the biggest being a lack of predefined connection with the research’s end-users. In Tennessee, we see building these connections as a discrete component of our work. Each research team member serves as a liaison to a department division, regularly meeting with division staff and seeking feedback on the direction of their research.

Equally important, research team members look for opportunities to perform “technical assistance” for their given division. Technical assistance here refers to the many data tasks that always need doing in a state agency—tasks that would not necessarily classify as research but that require some skill in manipulating data. By pitching in on data-related tasks for different divisions, research team members build the relationships that allow them to meaningfully connect their research to the department’s operational work. Without these personal relationships built on mutual benefit, our research team’s influence would likely be much diminished.

One issue to note: It is easy for technical assistance requests to overwhelm the actual research and analysis work. In Tennessee we have found that the department has an insatiable need for people who can answer questions about data and combine files into new spreadsheets. The quantity of federal reporting always outstrips the quantity of individuals trained to create the necessary spreadsheets; demand for data to be sliced in new ways is constant. While it is important for the research group to provide immediate technical assistance benefit to different department divisions, it is equally important for the research office to be insulated from the department’s many operational data tasks. Without some buffering, the research group will soon struggle to keep up with the day-to-day work of compliance and reporting and fail to serve the strategic purpose for which it was created.

Building and Financing the Team

Making research relevant is difficult work. In Tennessee we have benefitted greatly from ready access to well-trained researchers from institutions like Vanderbilt University. Yet the people we look for are not only those with strong methodological skills (or even with the ability to deal with large and messy data sets). Instead we look for people who are both ready to engage in rigorous analysis and who value the “everything else” described in the above sections that must accompany research to make it relevant.

Happily, the number of applicants who fit this description seems to be on the rise. Organizations like the Strategic Data Project have seen a steady increase in applications and their programs’ cohort size. The Institute of Education Sciences encourages training through pre-doctoral fellowship programs that require focusing on applied research and practitioner collaboration. In Tennessee, we have received a steady stream of well-qualified applicants for every open position and have benefitted from a series of interns from area Master’s in Public Policy programs who seek useful practicum experiences.

But even with a strong research team ready to do the work and an agency that prioritizes research, we in Tennessee still face the reality of limited funding and staffing. One way we have been able to combat this is to explicitly place

our research in service of various federal cost objectives and thus parcel out federal dollars across individual researchers. Each member of our team logs hourly personnel activity reports noting the projects they have worked on and the program link. For example, at the end of the year, one team member may wind up having worked 20 percent of his or her time on Title II initiatives, 30 percent on ESEA Title I initiatives, 40 percent on IDEA initiatives, and another 10 percent on state activities.

Even where federal funding might not be a possibility, the strategy of splitting funding for a research team across the areas where the analysts actually add value remains an important principle for us in Tennessee. Among other advantages, it forces the research team to actually make good on its value proposition. Since research team members are getting paid from buckets of money that other department teams control, our researchers must constantly ensure agency members feel like they receive direct benefit for the dollars spent.

Balancing the Work Across Internal Personnel and External Partners

Even among SEAs with strong research offices, the approach to external partnerships is often quite different. Some SEAs choose to rely on their own analysts to complete most of the work; others farm all projects out to research firms and academic partners. The place where an SEA falls on this spectrum has considerable implications for the type of staff needed to staff the research team.

In Tennessee, we aim to place ourselves somewhere near the center point, claiming most of the quick turnaround and forward-looking analysis for our internal team, and then passing off longer-term evaluations to external researchers. Under this theory, external researchers serve a role that is meaningfully different than that of even the best internal staff members. Research partners offer an independent look at state initiatives, they bring new ideas into the department, and they often use more rigorous (and slower) methodologies than a rapid-response state team. The challenge that states face is finding ways to balance these gains with the additional challenges of taking research outside the agency.

While it is simple enough to define a set of products for research that is created by an internal team of researchers, it is far more difficult to communicate these expectations to a varied set of research partners that are drawn from a wide set of research firms and academic institutions.

Even the best of research partners do not operate within the agency, and so they will tend to be less adept at framing questions and findings in a way that feels relevant and influential than an internal team that communicates with its

constituents on a day-to-day basis. Moreover, research partners face different cost structures and incentives, all of which can interfere with demands of timeliness and/or presentation.

Managing this process while containing the amount of time spent on this management has proved to be a considerable challenge for our team in Tennessee, and we feel like it is an area where we have much to learn from other states. Notwithstanding some striking successes with individual researchers, we continue to struggle to reach a point where the process of partnership—and the considerable cost of time and effort—feels like it consistently pays off in terms of research and policy impact.

DEVELOPING THE WORK

While we have logged several successes tying department practice to research in recent years, our team believes that our current structures fail to meet our agency's needs in two central areas. First, we are not set up to track recurring metrics over time and to support repeated analysis of these metrics, even as we increasingly produce data points that we would like to track longitudinally. Second, we struggle to connect the dots across research projects (internal and external) within the same areas, contributing to a sense that we are undertaking a series of disconnected studies rather than producing definitive work within a few high-priority focus areas. This final section details our diagnosis of these missing pieces and our strategy for moving from where we are to where we hope to be.

Our team often jokes that we only start new projects, never finish old ones. Successful work produces new metrics and new ways of looking at data that in turn generate further appetite for tracking these metrics over time. For example, analyses investigating the landscape of human capital management across the state (teacher retention, equity in student-teacher assignment, etc.) produced a series of data breakdowns within these areas that department personnel hope to keep tabs on over time, both at the state level to monitor progress on strategic priority areas, and at the district and school levels to target interventions and support district strategic planning. Yet if our group committed to producing yearly analyses or even yearly data reports following human capital metrics at the state, district, and school level—and if we were to meet similar demands in other areas such as discipline disparities, chronic absenteeism, and several others—we would quickly lose the ability to undertake new projects. Moreover, at some point, the work of tracking chosen metrics and analyzing these data starts to look less like research and analysis and more like progress monitoring—a different, though equally vital, organizational need.

To answer this need, we are investigating the possibility of creating a small offshoot group to help coordinate department-wide strategic progress monitoring. With assistance from each division, this team would help ensure that we follow the right metrics and use these as levers to drive our work. The goal would be to take the indicators and tools developed on the research side and put them into the hands of an expanding state audience. Over time, some metrics might develop into established ways of looking at state educational data, at which point they would migrate to our state's online report card. Others may serve a particular need at a moment in time and then eventually drop out of use. Importantly, the team's charge would not be to create platforms and structures just to monitor recurring strategic metrics over time (we expect that such data points on their own would be only marginally useful). Instead, the team would work closely within the department and with our regional offices to directly engage others in understanding these metrics and to draw on their expertise in order to surface new ideas for current state strategies.

Our second major challenge has been to connect the myriad research efforts going on simultaneously through our office and external research partnerships. We are in the midst of two internal projects that look at the connection between our state teacher evaluation system, teacher feedback, and instructional improvement. Two other external research projects take on the same topic from different angles.

Each of these studies is likely to produce interesting results and we hope each will drive agency action. But each project feels like an individual venture independent of the others, which we believe significantly lessens the total impact. We have not yet managed to create a system or process where we regularly look across current projects to ensure that each ongoing study takes advantage of the other researchers' information. And we have not developed a way to tell the broader story—to our department or the public—of what we have learned in Tennessee from all the research on teacher evaluation and instructional improvement.

CONCLUSION

As our office continues to develop, we are looking for ways to make each research project contribute to a broader picture of educational practice and policy in Tennessee, potentially with help from outside research organizations. We aim to synthesize findings across our research portfolio to offer integrated recommendations about state and district action. But this work demands both a different skill set and a greater number of positions than our office can currently fund. We join other SEAs in struggling to balance priorities as we identify new project goals.

ENDNOTES

1. For a thoughtful discussion of the time and effort it takes to integrate research results into agency operations, see Melissa Roderick, John Easton, and Penny Bender Sebring, *The Consortium on Chicago School Research: A New Model for the Role of Research in Supporting Urban School Reform* (Chicago: University of Chicago, 2009).