



Is My Ed Tech Tool Making a Difference?

An Entrepreneur's Guide to Using Research to Improve Products and Measure Impact


newschools
venture fund



by Tonika Cheek Clayton and Cameron White

CONTENTS

03	INTRODUCTION
04	GETTING STARTED
06	DEFINE INTENDED IMPACT + LOGIC MODEL
08	ITERATE ON USABILITY + FEASIBILITY
10	EVALUATE STUDENT OUTCOMES
13	SHARE YOUR ED TECH IMPACT STORY



Introduction

Every ed tech entrepreneur wants to develop an amazing tool that may one day become a household name. Many also genuinely want to support positive outcomes for students, educators and parents. At NewSchools, we invest in ventures that care deeply about impact. Through our ed tech accelerator NewSchools Ignite, we help entrepreneurs use research to generate data that can inform their product and business strategies. Across the development cycle, we have observed how research can uncover valuable insights into how well products are working, for which users, under which conditions.

Based on these experiences working alongside our portfolio, we've developed a guide designed for entrepreneurs at any stage of their research journey. This guide examines a range of research types, noting those that are appropriate at various stages of development, on small or large budgets, and within varying timeframes. Others in the ed tech research community have designed similar frameworks. However, our guide is designed specifically for entrepreneurs, highlighting the types of value research can create while taking into account the realities of resource-constrained startup environments. We believe it will also offer useful information for educators and funders.

There are four key research practices explored in this guide:



Define the intended impact and create a logic model. Identify the student outcomes you believe your product supports, document how use of your product supports these outcomes, and gather feedback on your ideas from a diverse group of stakeholders.



Iterate based on feedback from usability and feasibility testing. Observe how the product is used in a real setting by teachers and students, and refine product features and supports based on what you learn.



Evaluate evidence of student outcomes. Develop research questions and timelines that align with your product and business roadmaps. Be mindful of key considerations including product stage, costs, and return on investment.



Share what you learn about impact - both the celebratory insights and the tough lessons. Synthesize different types of evidence to describe how your product can support student outcomes, focusing on key audiences like educators and funders.

One of the most important points we want to make is that many types of research have value, and there is often something that can be learned and integrated into strategic decision-making. A user feedback session can uncover valuable information, as can a randomized controlled trial, or any of the many options that fall in between. We also want to emphasize that the entrepreneur's research journey is not linear. Research practices don't always fall perfectly into a sequential order, and each can have value as an independent undertaking.

Getting Started

So, how can an ed tech entrepreneur check in and collect evidence of impact in the midst of an iterative product development cycle? Many ed tech developers aspire to know whether their product is working as intended, yet only a small percentage of education interventions - technology-enabled or otherwise - have gathered rigorous evidence of positive student outcomes. The What Works Clearinghouse website provides a central repository of scientific evidence on “what works” in education to improve student outcomes. Yet only about 20 percent of the more than 440 K-12 interventions listed on the website have collected evidence that they produce “positive” or “potentially positive” student outcomes.¹

Sequencing and timing are important. Conducting an efficacy study before having confidence about how a product is being used in classrooms is probably not the best plan. However, it’s crucial to keep student outcomes top of mind, even at an early stage. The ultimate goals of early-stage research can be defined in many ways, and there are a number of approaches to measuring progress. Students and teachers might use a product daily and genuinely enjoy it. But that doesn’t necessarily mean it’s having a meaningful impact. To avoid such uncertainty, ed tech developers can make use of a range of study designs, each offering different types of evidence and representing varying levels of rigor. (We will discuss these options in greater detail later in this guide.)

Through our ed tech accelerator NewSchools Ignite, we invest in tools that support student learning and integrate research services into our investment strategy. We identified two external research partners to support our ventures. WestEd conducted product reviews and small-scale studies, and Empirical Education offered student user demographic reports. We build the cost of their work - an average of approximately \$40,000 per investment - into our venture support. The content and lessons learned in this guide emerged in part from this work, as well as from conversations with ed tech researchers, entrepreneurs and funders.

NewSchools Ignite launched six ed tech challenges from 2015 to 2018, focused on products addressing critical student needs in Science Learning, Middle and High School Math, English Language Learning, Special Education, Early Learning, and the Future of Work. Through these challenges we funded 84 small-scale research studies, conducted by WestEd, primarily focused on generating formative product feedback.

¹ Of the 448 K-12 interventions listed on the What Works Clearinghouse website (as of August 2018), 90 (20%) have submitted research that shows a “positive or potentially positive” effect. The remaining interventions have “mixed” or “no discernable” effects, or have submitted “no evidence.”

For ed tech ventures outside our portfolio, we recognize this level of early-stage research may be cost prohibitive. Although research is becoming part of the selection criteria used by many educators and funders, it is not routinely earmarked in investments. Many funders are beginning to realize the importance of evidence, and they should be willing to help pay for it. After all, research can enhance an investment by demonstrating how well the product is working while also supporting user and revenue growth.

We also know it's challenging for ed tech developers to find research partners that can provide useful information in the context of a fast-paced, iterative product development cycle. To support continued innovation, the sector needs more researchers who have a deep understanding of education and research as well as the realities of entrepreneurs' day-to-day experiences.


For example, fully-powered randomized control trials (RCTs) are often cited as the "gold standard"² of education research, yet they are frequently too expensive and slow-moving to be practical for many ventures. By the time results of a large-scale study are in, often the product has changed, the needs have changed, and the landscape has shifted. In recent years, several organizations have explored the potential value of "rapid-cycle evaluation,"³ designed to "quickly determine whether an intervention is effective" while also enabling "continuous improvement." Yet even this type of research requires significant

financial and human capital resources, so it's important to consider its costs and benefits as part of a sustainable research strategy.

Part of our intention is to create more empathy and understanding between entrepreneurs, educators, researchers and funders, since their goals are generally aligned – even if there are significant differences in perspectives. We see a need to share more insights, best practices and lessons learned among these stakeholders. If an ed tech developer can be certain a product is working, and continues gathering evidence as the product matures, he or she will be better positioned to communicate about impact.

The four key research practices are outlined in greater detail on the following pages.

Ed Tech Research Journey



² U.S. Department of Education (2003). Identifying and Implementing Educational Practices Supported By Rigorous Evidence: A User Friendly Guide. Retrieved from <https://www2.ed.gov/rschstat/research/pubs/rigorousetid/rigorousetid.pdf>

³ Asher, A., and Cody, S. (2014). Smarter, Better, Faster: The Potential for Predictive Analytics and Rapid-Cycle Evaluation to Improve Program Development and Outcomes. Retrieved from https://www.brookings.edu/wp-content/uploads/2016/06/predictive_analytics_rapid_cycle_evaluation_cody_asher.pdf



Define intended impact + logic model

The first step in the research process is defining the student outcomes you believe your product can support. Optimally, these outcomes should be aligned with the goals of potential users or purchasers – students, teachers, school leaders and/or district administrators. Next, developers can begin to consider short- and long-term metrics related to these intended outcomes. For example, if the ultimate goal of a product is to enhance student literacy, what should be happening in the short-term as an indicator of progress? Will the student perform better on an interim assessment? Or maybe a tool is designed to increase student engagement. In that case, is the goal purely to optimize screen time, or will the teacher be able to observe students interacting more during class discussions? If they are defined early,

Across NewSchools Ignite’s first six challenges, we observed a range of intended impacts including various academic, social-emotional, and career and college-ready indicators.

these intended outcomes (and, eventually, related evidence) can be integrated into the product’s overall value proposition.

After defining intended outcomes, ed tech developers can begin to outline how access to and use of the product is connected to these outcomes. By explicitly defining potential use cases, the developer makes clear what is required to access the product, and how it should be used in order to achieve the desired results. For example, what needs to be true about the learning environment in terms of devices, student-teacher ratio, etc. for the product to function as expected? How often and for how long should the product be used during an ideal use case? Even though actual usage may vary considerably, it is helpful to have a baseline for comparison to product analytics and observations from classrooms.

Considerations of access, use and outcomes can then be formally documented through a logic model, which describes a product’s “theory of change” through the lens of potential inputs, activities, outputs, outcomes and impact. It also produces a series of testable if-then hypotheses that can lay the groundwork for future research.


Various team members (and possibly external stakeholders) can contribute to developing and refining the logic model. It can be a messy process, as we've found that even people within the same organization can have different ideas about how and to what ends a product is used. However, it can also be a helpful forcing mechanism to get everyone on the same page, and help clarify priorities for product development.

If, for example, a venture is developing a product that increases science content knowledge, the logic model needs to reflect that goal. If a team is not aligned on a logic model, they might end up with bells, whistles and features that don't actually support what they set out to do in the beginning - help students improve their scientific understanding. A similar line of thinking can be applied to a wide range of potential outcomes

valued by students and educators. For example, some tools are designed to support improved attitudes toward learning. In this case, the team would need to develop a theory of change that addresses engagement, in addition to academic gains.

Consider the logic model a living document; it will evolve as the team learns more. But here is a critically important point: This research practice (and all the others) will likely be repeated multiple times over the life cycle of the product. What if a team develops a math tool designed to help students learn decimals, but finds their instructional approach actually is better suited to help students learn fractions? That's not a bad thing; it's an important insight. And, it could signal a reason to revisit the logic model and redefine the tool's intended impact.

Logic models: Ensuring your product has impact



There should be evidence these causal links exist.

Image Source: [WestEd](https://www.wested.org)

Kao, Y., Matlen, B.J., Tiu, M., & Li, L. (2017). Logic models as a framework for iterative user research in educational technology: Illustrative cases. In R.D. Roscoe, S.D. Craig, & S.C. Douglas (Eds.), *End-User Considerations in Educational Technology Design*. Hershey, PA: IGI Global.



Iterate on usability + feasibility

Once an ed tech developer defines a product’s intended impact and documents its theory of change via a logic model, it’s important to begin collecting user feedback and other data to understand how this vision aligns with actual usage. To this end, usability testing generates feedback about user interface and experience (UI/UX) based on 1:1 testing in a controlled lab setting.

Building on this foundation, feasibility testing provides information about the implementation of a product by teachers and students within an authentic learning environment. In addition to helping improve a product, understanding usability and feasibility is a prerequisite for conducting research on student impact.

Research studies focused on formative feedback

Study type & cost range*	Sample and setting	Potential research questions	Potential funding sources
Usability \$1,000-20,000	1:1 user testing (e.g. students, teachers, and/or parents) in a controlled lab setting	Is the product intuitive and easy to use? Are users able to use the product’s features as intended?	Self-funding Foundations and impact funders
Feasibility \$20,000-60,000	A complete product implementation is tested in authentic learning environments (e.g. classrooms)	How do students/teachers use the product in the classroom? What support materials and guides can be provided to help facilitate use in the classroom? What are the barriers to classroom implementation?	

Table Definitions. *Usability:* Usability studies test whether the core features of the product are usable by the intended end user. *Feasibility:* Feasibility studies test whether the product can be used at scale by the intended end users in an authentic educational context.

*Study costs vary widely and depend on a range of factors including: product type, whether research is conducted internally or externally, researcher reputation, geographic location, sample size, sampling methods, types of data collected, measurement tools used, and types of analysis performed.

At this stage, an ed tech entrepreneur should be working to understand whether the product can be used and implemented successfully within its intended learning context. Depending on a venture's internal capacity, these tests can be conducted by team members or by an external research partner, and can be either self-funded or supported through investments from foundations or impact funders. In any case, to generate the most valuable feedback it's essential to ensure the product is tested with a demographically diverse set of users, including (where applicable) English Language Learners and students with disabilities.

Across studies of products funded through NewSchools Ignite, evidence suggests making products easier to use can have a positive impact on student learning.

Usability and feasibility testing can also provide valuable information about potential improvements, which can be integrated into a venture's product roadmap and value proposition. It is even possible that, by observing the behaviors of new users, developers may uncover new ideas for use cases or potential outcomes. In order to reach scale, a venture needs to understand not only who they are building for, but also the various benefits and costs associated with use and implementation. For example, how much time does it take teachers and students to learn and use the product? What technical and human capital resources are needed? In addition to helping developers become more fluent in the day-to-day realities of users, answering these types of

questions can also uncover potential supports like lesson plans or onboarding guides that reduce friction as the number of students and teachers using the product in the classroom increases.

Compared to summative research, which focuses on measuring outcomes, formative research is meant to gather data that can inform product and business strategy, and is relatively inexpensive and low-risk. Indeed, best practices in user testing overlap quite a bit with principles from design thinking and lean product development. Early-stage usability and feasibility testing can help protect future investments in research by ensuring outcome data is not distorted by friction that could have been addressed through product design. As a team grows, the venture will likely add more people to the product team who are engaging in this type of work.

Whether a venture conducts its own user testing or uses an external resource, it is important to try to mitigate potential biases as much as possible. For example, an internal researcher may be predisposed to ignoring critical feedback because of their closeness to the product. On the other hand, external researchers must also take care to learn about the developer's goals and the nuances of a product they did not design. Their distance from the product or the users could create another type of bias. Ideally, team members can leverage best practices from the research community, partnering with external researchers when it is advantageous while always looking for opportunities to build their own internal research capacity. This helps sustainably integrate research practices into the team's culture as well as its product and business strategies.




Evaluate student outcomes

As the product matures, if it is performing well, there should be positive indicators that suggest it is worth investing more time and money into the product, including more rigorous research. At NewSchools, we define “rigorous evidence” as a randomized controlled trial (RCT) or quasi-experimental design (QED) study, conducted by an external researcher, that demonstrates positive student outcomes. The federal K-12 education law, Every Student Succeeds Act (ESSA) also stipulates that studies meeting its evidence standards must

be “well-designed and well-implemented,” which places additional requirements on the study design.⁴ That being said, many types of evidence have value. For example, a pre-post measurement demonstrating statistically significant gains might give a venture more confidence in investing in more rigorous research, whereas a case study documenting a successful implementation might help potential users understand the value the product can bring in classrooms like their own.

FORMATIVE
Actionable feedback

SUMMATIVE
Evidence of outcomes



Study designs range from formative (on the left) to summative (on the right), with both rigor and cost increasing as you move from left to right (Image source: WestEd).

Image Source: WestEd

⁴ U.S. Department of Education (2016). Using Evidence to Strengthen Education Investments. Retrieved from <https://ed.gov/policy/elsec/leg/essa/guidanceusesinvestment.pdf>

Summary of research studies focused on measuring student outcomes

Study type and cost range*	Prerequisite evidence needed	Why conduct this research	Evidence level**	Potential funding sources
Pre-post or Generic controls \$10,000-50,000	Evidence of usability and feasibility	Communicate potential impact to school and district decision-makers Create evidence base that makes future research grant applications (e.g. ED/IES SBIR) more competitive In some designs, surface feedback that can drive product improvement	Preliminary (Does not meet ESSA evidence standards)	Self-funding Foundations and impact funders
Correlational (statistical controls) or Randomized control (underpowered) \$45,000-250,000	Evidence of usability and feasibility Fidelity of product implementation Preliminary evidence of positive student outcomes	Communicate potential impact to school and district decision-makers (including those that require alignment with ESSA standards) Create evidence base that makes future research grant applications (e.g. IES Goal 3) more competitive Feedback to tweak product/ implementation in preparation for more rigorous/expensive studies	“Promising” (ESSA Tier 3)	Self-funding Foundations and impact funders Federal grants
Quasi-experimental (well-matched comparison groups) \$27,000-800,000	Successful usability and feasibility studies Fidelity of product implementation Preliminary evidence of positive student outcomes	Possible inclusion in What Works Clearinghouse “with reservations” Create evidence base that makes future research grant applications (e.g. IES Goal 3) more competitive Feedback to tweak product/ implementation in preparation for more rigorous/expensive studies In some designs, possible to compare performance among demographic subgroups	“Moderate” (ESSA Tier 2)	Self-funding Foundations and impact funders Federal grants
Randomized control (fully powered) \$250,000-3 million +	Successful usability and feasibility studies Fidelity of product implementation Promising or Moderate evidence of positive student outcomes	Differentiation - very small percentage of products offer evidence that meets this standard Possible inclusion in What Works Clearinghouse “without reservations” “Gold standard” of education research	“Strong” (ESSA Tier 1)	Federal grants Foundations and impact funders

Table Definitions. *Pre-post:* Pre-post studies examine changes in an outcome measured before and after an intervention. *Correlational (statistical controls):* Correlational studies examine whether changes in one variable correspond to changes in a second variable. *Quasi-experimental (well-matched comparison groups):* Quasi-experimental studies compare outcomes for intervention participants with outcomes for a comparison group chosen through methods other than randomization. *Randomized control (fully powered):* Randomized control studies randomly assign participants to an intervention or control group, in order to measure effects of the intervention while minimizing bias and other external factors. *Generic controls:* Generic controls studies compare performance results (not growth) of a treatment group to nationally or regionally accepted benchmarks or proficiency goals. *Statistical controls:* Statistical controls are methods (e.g. multiple regression analysis, fixed effects, propensity scoring) that compare treatment group performance to that of an “equivalent” population. *Well-matched comparison groups:* In a strong QED, the comparison group will be close to a mirror image of the treatment group. *Underpowered RCTs:* Underpowered RCTs have a lower probability of detecting an effect on student outcomes. *Fully powered RCTs:* Fully powered, well-designed and well-implemented RCTs provide the highest degree of confidence that an observed effect was caused by the intervention.”

*Study costs vary widely and depend on a range of factors including: whether research is conducted internally or externally, researcher reputation, geographic location, sample size, sampling methods, types of data collected, measurement tools used, and types of analysis performed.

**To meet evidence levels as defined in the federal Every Student Succeeds Act (ESSA), studies must be “well-designed and well-implemented”, with sampling methods and student growth that generate sufficient statistical power to indicate that the product has a positive effect on student outcomes.

As rigor increases, so generally does the cost of research. While investment in ed tech continues to grow, there will likely be additional resources available for research, but entrepreneurs need to be able to make the case that research is aligned with product and business goals.

There are limited opportunities to fund ed tech research with external capital. The most common options are through government- or foundation-sponsored grants. Some later-stage ventures choose to self-fund studies, and impact-focused investors have demonstrated a growing appetite to fund certain types of research. One barrier is that the return on investment can be difficult to calculate since it is unclear how much educators use different types of evidence to help inform their purchasing decisions. However, even entrepreneurs and investors more focused on financial returns should consider the potential value of research over the long term as a potential differentiator.

Most entrepreneurs naturally think about research in terms of the evidence necessary to communicate the product's impact to potential purchasers or funders. At this juncture, entrepreneurs need to ask the right types of research questions, and plan studies alongside product and business timelines, so the data generated are useful and relevant. For many years, impact was defined primarily in terms of gains on math and reading test scores. Today, ed tech impact can be defined more broadly, so it's essential to work with customers to understand the outcomes they value. For example, many educators are interested in how the product has worked specifically in classrooms, schools or districts similar to their own (demographically, geographically, or otherwise). Often, a blend of evidence types - including outcome data as well as qualitative feedback - is most effective.





Share your ed tech impact story

Once the evidence is synthesized, it's time to communicate how the product is impacting student outcomes. With respect to ed tech impact, even mature well-resourced companies may not have a good sense of what they should be sharing, with whom or how best to share it.

When communicating with an external audience about research, it's important to distill your evidence down into concise statements describing how a study was conducted and what evidence it produced. These statements should be aligned with, and ideally reinforce, your product's overall value proposition. If a study was conducted by internal team members, be transparent about that, using it as an opportunity to describe your venture's commitment to building internal research capacity. If you complete a study with support from an external researcher, it is often helpful to request their input and/or feedback to ensure statements accurately reflect the research conducted.

Before sharing your impact story, always start by determining the objective and the audience. For example, if the goal is to use research findings to make a pitch to investors, the message will be very different than if the goal is to compel more teachers to adopt the tool. The distribution channels must also be tailored to match the audience. Students, educators and parents all consume information in different ways, and in different places. If you are trying to appeal directly to students, it's important to try to meet them where they are, explaining the value your tool can bring to their lives. If you are trying to reach adult users, consider purchasing micro-targeted web-based ads or doing direct outreach to bloggers who are respected among your target audience.

Potential outcome statements and descriptions by study type

Study type	Potential outcome statements*	Potential description of sampling methods and data collection techniques
Usability or Feasibility	"[Venture] worked with [researcher] to understand user perspectives and integrated [user] feedback to improve product usability and/or feasibility."	"[Users] were recruited via [venture and/or researcher] contacts."
Pre-post	"The study showed statistically significant gains in students' content knowledge as well as [other outcomes], setting the stage for future explorations of product efficacy."	"During the intervention, two teachers implemented [product]. Data were collected through a pre-post student content quiz and [other data sources]."
Correlational	"The correlation of product usage and student performance on [measure] show promise that the product has a potential impact on performance."	"The study makes use of student level data collected from [school district] matched to student usage data from [product]."
Quasi-experimental	"Students who used [product] at recommended dosage saw additional growth of [X]% compared to the control group, with an effect size of [Y]."	"A true, group-randomized, experimental design was used to control for most threats to internal validity. [Users] were randomly assigned into treatment and control conditions."

*Actual outcome statements are dependent on study design and results.

It is also critical to think about the venture's brand identity, and how research fits into it. Is the goal to be known as a venture with research at its core? Or does the company define itself in a different way, with research used as a way to support a broader brand promise? Regardless of where the company fits on this spectrum, here are some guidelines for communicating:

- Avoid embellishment and be sure any statements or claims about the product are accurate and can be substantiated.
- Make optimal use of the channels the company can control, such as Twitter, Facebook, LinkedIn, the company website and blogs. Publish issue briefs or white papers with tight overviews. Use analytics to track reach and ROI.
- Use language that is accessible and friendly. Avoid excessive use of jargon, acronyms and complex language. Focus on top-line findings and keep it simple. Research need not be obtuse.

- Remember the power of first-person testimonials and storytelling to bring the research story to life.
- Share research findings at conferences about education research, ed tech or PreK-12 education – being sure to tailor the message for different audiences.
- Use third-party validators such as other researchers and thought leaders in education to amplify the message.

Once you reach this stage, it is often helpful to reflect back on the other three research practices. How might feedback from various stakeholders influence how you talk about the outcomes your product supports? Does the way you describe your potential impact align with the use cases that are most popular with users? What additional evidence could be used to strengthen your impact story?

Conclusion

Call to Action

Through three years of making ed tech investments, NewSchools has had the opportunity to work closely with 74 ventures creating products to support Science Learning, Middle and High School Math, Early Learning, Special Education, English Language Learning, and the Future of Work. Each cohort was focused on developing different types of tools, yet there have been notable similarities in their experiences when it comes to research.

Our investments to support research have surely created value for our ventures and their products' users. We hope this guide - which distills much of what we've learned about ed tech research - extends our collective knowledge to entrepreneurs beyond our portfolio, as well as to educators, researchers and funders who are investing resources into this important work. Together, we can help more students and teachers have access to technology-enabled tools that are accessible, easy to use and implement, and effective at supporting positive outcomes. We believe the following research-related actions will strengthen the ed tech sector overall:

- **For entrepreneurs:** Make it a priority to collect evidence. Remember that regardless of your budget or the product's maturity, many types of research can be valuable - and not just to you, so share what you are learning. These valuable insights provide evidence of how well and under what conditions your product is working.

- **For educators:** Think of yourself as a partner, not just a user or customer. Be clear about what you and your students need, and consider providing critical feedback to help entrepreneurs refine their products. Ed tech developers want to create tools that make a positive impact, but they have to be in close communication with educators to know whether they are achieving that goal.
- **For researchers:** Put yourself in the seat of ed tech developers and educators. The product development cycle is dynamic and fast paced, and classrooms are full of students who are literally changing and growing every day. Ed tech entrepreneurs need information that is timely and actionable. It can also be helpful to bring some creativity when thinking about how and when to measure progress across various dimensions of growth.
- **For funders:** Fund ed tech research, and build it into your overall investment. Early-stage ed tech ventures have the double challenge of needing to show impact while having limited resources and capacity available to measure it. As you learn from their work, it will broaden your own knowledge base, and generate insights with potential to inform the entire sector.

NewSchools Venture Fund

1616 Franklin Street
Oakland, CA 94612

www.newschools.org