

What does good EdTech evidence look like?

Key insights from the EdTech Evidence Board pilot

September
2026



Introduction

Across the education sector, there is growing emphasis on the importance of evidence in EdTech. Schools and colleges are increasingly expected to make informed decisions. Developers are encouraged to rush to demonstrate impact. Policymakers are seeking greater assurance about what works - and yet, there is no single, shared answer to the question: What does 'good' EdTech evidence look like?

In fact, the answer to this apparently simple question is perhaps more complex than we might think. For example, there is no universally agreed definition of what we even mean by the term 'EdTech'; there is no common language for describing different types of evidence; and no clear and consistent expectation about what exactly EdTech suppliers should be aiming to evidence, and how they might approach this. At the same time, EdTech products themselves are complex. They evolve rapidly, are used in different ways across contexts, and often combine multiple features, approaches and intended outcomes. This adaptability is on one hand a benefit, and on the other, can make products difficult to evaluate using simple or standardised approaches.

Taken together, this presents a challenge: It is difficult to answer the question of what 'good evidence' looks like because we do not yet share a clear and consistent way of talking about evidence for EdTech products.

This does not mean that good evidence does not exist. But it does mean that it is often difficult to recognise, compare or interpret evidence. In the absence of shared standards, EdTech suppliers may default to producing evidence that is determined most likely to yield market success - rather than producing evidence that would be most useful to educators.

The EdTech Evidence Board (EEB), convened by the Chartered College of Teaching and funded by the Department for Education, was established to explore how evidence gathered about EdTech products can be assessed in a structured and consistent way. As part of this work, a pilot was conducted with a cohort of 12 EdTech products currently used in UK education settings. Participating organisations submitted evidence portfolios, which were reviewed by independent experts against a defined set of criteria which together aim to provide a holistic indication of product efficacy.

This process provided a unique opportunity to examine questions such as:

- »» What kinds of evidence are organisations able to produce?
- »» How is evidence presented and interpreted?
- »» What are the common strengths and gaps that exist?

And perhaps most importantly:

- »» What is it that distinguishes stronger portfolios of evidence from the weaker ones?

This guidance attempts to draw together the resulting insights to build a picture of 'what good EdTech evidence looks like' as surfaced through the EdTech Evidence Board pilot. These insights are exemplified through accompanying case studies which aim to provide a picture of how 'good evidence' might be enacted in practice.

Whilst we hope that you find this guidance to be helpful, the insights and examples shared here should not be viewed as an exhaustive list, but rather a starting point for thinking about evidence quality in relation to EdTech. This is one of a number of resources developed as part of the EdTech Evidence Board (EEB) project and we recommend taking the opportunity to explore the full breadth of resources - including the EEB evaluation criteria - on the Chartered College of Teaching website.

Sam Evans

Head of EdTech, Chartered College of Teaching
Contact us: edtech@chartered.college

Key insights from the EdTech Evidence Board pilot

EdTech operates in a context that is materially different to most other product-led environments. While many digital products are evaluated primarily in terms of user growth, engagement or commercial performance, EdTech products are used within a system that shapes children's learning, development and life chances. As a result, this introduces a different level of responsibility.

The effects of a particular product may be positive, neutral, or, in some cases, inadvertently negative - particularly where design choices influence how students engage with content, how teachers make decisions, or how learning is structured over time. As a result, evidence in EdTech is not simply about demonstrating that a product 'works' in a general sense, but about understanding how it affects learning, for whom, and under what conditions.

This reinforces the importance of robust, well-interpreted evidence and highlights the need to place greater emphasis on the careful alignment between claims, design and impact compared to what may typically required in other product contexts.

Understanding this distinction explains two phenomenon witnessed through the pilot:

1. That many of the EdTech organisations engaging in the pilot indicated an urgency to demonstrate the impact of their products on student outcomes - in a number of cases, this urgency meant that experimental studies were conducted too soon, and opportunities were missed to gather rich data that would help to better understand - and improve - the product's use in real educational contexts.
2. Much of the available product evidence - particularly case studies, testimonials and implementation stories - sat between two different purposes: demonstrating pedagogical value and supporting commercial growth. These materials often contained potentially useful insights, but were typically designed as sales collateral. As a result, they tended to foreground positive experiences and outcomes without providing critical interpretation or acknowledgement of limitations needed to function as robust evidence of efficacy.

Shifting from 'evidence activity' to 'evidence literacy'

Learning from the pilot suggests that what determines 'good' isn't always the volume or frequency of evidence generation. In most cases, organisations were active in engaging with evidence: They had collected data; referenced research; commissioned or conducted evaluations; and developed case studies and reports. This activity is important and meaningful. It reflects a growing recognition across that sector that evidence of efficacy matters and evidence activity signals clear intent and commitment to demonstrating it.

However, activity alone does not resolve the question of quality nor does it - alone - provide the answer to the question we are looking to answer. A clearer distinction emerges when considering how evidence is *understood* and *used*. We call this 'evidence literacy'. In the strongest examples seen within the pilot, EdTech suppliers demonstrated a form of evidence literacy that ensured evidence was:

- » interpreted in relation to a clear purpose and claim
- » connected across different stages of development and use
- » centered by strategic directions and roadmaps
- » used to inform decisions.

In these examples, individual items of evidence formed part of a coherent whole.

In weaker examples, evidence-related activity was present but far more fragmented. Data, research and evaluation existed, but were often presented in isolation, not clearly linked to one another, and not necessarily used to inform product development and improvement.

This created a gap between *effort* and *impact* that provides the key distinction between two different organisational approaches:

The distinction is not between organisations that do and do not engage with evidence. It is between those that treat evidence as an activity, and those that use it as a system.

This distinction helps to explain a broader pattern observed across the pilot. Where evidence literacy was stronger, it was possible to see clear connections between:

- » what a product is designed to do
- » how it is used in practice
- » what difference it makes
- » how this is measured.

Whilst not a tangible or practical example of what good evidence looks like, what is described here is certainly one of the most important characteristics that contributed to a higher level of performance of an evidence portfolio within the pilot.

For these higher-performing portfolios it was significantly easier to demonstrate product efficacy across the multiple areas evaluated as part of the EEB review process, but perhaps more importantly, they demonstrated that adopting a more evidence literate approach can have a potentially transformational effect upon the entire operation of an EdTech product, that would not otherwise be achieved through 'evidence activity' that is deployed as a vehicle for market-growth.

The sections that follow explore what **evidence literacy** looks like in practice, considering six characteristics that may be viewed as indicators of 'good' EdTech evidentiary practice.

Good evidence...

1. is aligned to a clearly defined product purpose
2. shapes product design
3. draws on data that supports a hypothesis
4. enables valid product claims
5. forms part of a culture of continuous improvement
6. is gathered as part of a wider evidence strategy concentrated on impact

Evidence literacy in action

Potential indicators of good evidentiary practice

1. Good evidence is aligned to a clearly defined product purpose

Good evidence begins with precision. In many cases, products are positioned to support broad goals like 'raising attainment' or 'improving engagement'. These are valid ambitions, but they are not, on their own, specific enough to guide clear evidence.

Products with a strong evidence portfolio start with a clear definition of:

- who the product is for
- the specific problem it addresses
- the outcome it is intended to influence.

These portfolios also set out a clear rationale for how change is expected to happen through use of the product. This is often expressed through a **logic model or theory of change**. In stronger cases, this is more than a high-level summary; it's a working model that constrains both product design and evaluation. The clarity does more than improve description, it enables evidence.

When the problem and intended outcomes are well defined, it becomes significantly easier to demonstrate:

- **educational need** - because the problem is clearly evidenced and contextualised
- **coherence** - because different elements of the portfolio relate to a shared purpose
- **design rationale** - because decisions can be justified against a defined problem.

Where scope is clearly defined, claims about the product's efficacy and impact are also more naturally aligned to what the available evidence can support. In this sense, a well-defined problem supports **claim discipline** - the ability of an organisation to clearly define what its product is intended to achieve, what evidence objectively says about its performance and how closely related these two are to impact claims. It limits overreach and ensures that evidence is interpreted within appropriate boundaries.

2. Good evidence shapes product design

It is now common for EdTech products to reference research. Many cite established fields such as cognitive science or pedagogy. This reflects a positive shift in the sector. However, the way research is used varies significantly. In stronger examples, there is a clear line from:

research insight → **design decisions** → **user experience**

This makes it easier to demonstrate how design is informed by evidence, rather than simply aligned to it. It also strengthens coherence across an evidence portfolio. Research, design and intended outcomes are connected, rather than presented as separate components. In weaker examples, research is used to support claims but is not clearly connected to how the product works in practice.

Good evidence is not something that is added on. It is visible in the design of the product.

3. Good evidence draws on data that supports a hypothesis

Most EdTech products generate significant amounts of data. However, how strongly data performs as evidence of impact efficacy depends upon a number of factors. Across the pilot, it was common to see large volumes of data presented with limited analysis. In these cases, the meaning and purpose of the data was unclear. This made it difficult to demonstrate how users were engaging with the product, what patterns were emerging and what implications this had for development.

Strong evidence portfolios use data purposefully:

- to analyse patterns over time
- to explore variation across users and contexts
- to connect findings to specific questions.

This data analysis makes it easier to demonstrate user experience and impact, as evidence is clearly interpreted rather than simply presented. Data becomes evidence when it is used to answer a question or support a hypothesis.

4. Good evidence enables valid product claims

Across the pilot, organisations drew on a range of evidence types, including case studies and testimonials, internal usage and outcome data, pilot evaluations and quasi-experimental studies conducted over time.

Each of these types of evidence plays a valuable role; the challenge is around how they are used, and if they count as objective evidence to support appropriate product claims. In some evidence portfolios, claims about impact extended beyond what the supporting evidence demonstrated. For example, a case study may provide rich insight into how a product is used in a specific context, but offers limited basis for generalisation of use in wider contexts. Internal data may indicate patterns or associations, but does not, on its own, establish causality. A quasi-experimental design however, particularly when developed over time and across contexts, can provide stronger evidence of impact.

The challenge is not about the presence or hierarchy of evidence types, but how are they utilised. Good evidence practice makes this distinction clear and plain. Claims about impact should not be made without sufficient regard to:

- the limitations of the evidence
- the extent to which findings can be generalised
- the difference between correlation and causation.

5. Good evidence forms part of a culture of continuous improvement

The most consistent gap observed across the pilot was not the absence of data, or even the absence of impact evidence. It was how evidence was used over time.

In many cases, evidence was treated as a collection of outputs. What was less visible was how this data informed decisions or led to changes in the product itself. This made it difficult for suppliers to demonstrate how evidence influenced product design; whether improvements had been tested; and how the product had evolved.

Stronger portfolios of evidence followed a different pattern in which evidence was an established part of an ongoing cycle:



6. Good evidence forms part of a wider evidence strategy

Logic models can provide a foundational role in a product evidence strategy, but evidence can be limited when suppliers are in too much of a rush to demonstrate impact at too early a stage.

The EEB evaluation criteria are designed to recognise the range of evidence that could help to demonstrate product efficacy, and highlight where different types of evidence might provide earlier indications of product effectiveness - this kind of evidence has the potential to offer greater insight about where improvements could be leveraged to strengthen the impact of a product in the longer term.

The criteria cover six key areas:

- 1. Approach to evidence:** Does research and evidence inform product evaluation and development?
- 2. Educational need:** Is there evidence to suggest there is a sound educational need for this product?
- 3. Pedagogical design:** Is the design of the product informed by evidence?
- 4. User experiences:** Is there evidence from users that suggests the product is able to be used effectively?
- 5. Pedagogical affordances:** Does research/evidence demonstrate that the product can enhance teaching and/or learning?
- 6. Impact:** Is there evidence to demonstrate positive effects on identified outcomes for teachers and/or learners?

An evidence strategy that seeks to gather a range of evidence may be most helpful when it comes to evidencing effectiveness and the EEB evaluation criteria provide a useful framework for auditing existing evidence in order to identify gaps and inform next steps.

Bringing it together

The question of what good EdTech evidence looks like does not have a single answer. However, the findings from the pilot suggest that strong evidence shares a defining common characteristic: **coherence**.


And coherence is not achieved by addressing individual criteria in isolation. It is achieved through:

- » clear definition
- » disciplined use of evidence
- » strong connections between design, use and impact.

At its core, this is a question of evidence literacy. When these elements are in place, it becomes significantly easier to both capture and demonstrate the efficacy of an EdTech product - and consequently alignment with EdTech evidence frameworks like the EdTech Evidence Board criteria. Without these elements in place, even substantial amounts of evidence 'activity' can fail to provide real clarity or insight.

When it comes to performance against the EEB evaluation criteria, the pilot has demonstrated that stronger portfolios did not address criteria in isolation. They were structured in a way that made multiple criteria easier to demonstrate simultaneously.

The diagram below aims to show how suppliers might adopt a tiered approach to meeting the EEB evaluation criteria:





Case study one

Teachmate: Working with schools to generate meaningful evidence

This case study explores how EdTech organisations can work collaboratively with schools and educators to generate more meaningful, credible and actionable evidence about product implementation and impact.

Across the EdTech sector, organisations frequently identify school engagement as a significant barrier to evidence generation. Time pressures, competing priorities and varied implementation contexts can make it difficult to collect evidence that reflects day-to-day educational priorities. As a result, many organisations struggle to gather evidence that reflects authentic classroom use over time, limiting both the quality of evidence collected and its usefulness for product development.

In response, some organisations are developing more collaborative approaches to working with schools. Evidence generation becomes part of an ongoing dialogue with educators and schools, rather than a discrete research activity. These partnerships can provide richer implementation insights, support responsive product development and ensure evidence remains grounded in real-world practice.

Context: Introducing Teachmate

Teachmate focuses on developing AI-powered tools designed to support teacher workload, planning and classroom practice. As part of the EdTech Evidence Board (EEB) pilot, Teachmate submitted evidence relating to its product development approach, implementation support structures and wider evidence-generation practices.

Throughout both the portfolio review and subsequent discussions, Teachmate emerged as a strong example of how an EdTech organisation has sustained collaborative relationships with schools to strengthen product development and evidence-generation practices.

This case study draws on insight generated through the independent review and an interview with CEO of Teachmate Dr Jon Chippindall. Together, these highlighted four key features of the organisation's approach:

- 1. Co-creation and strong educator relationships**
- 2. Staying close to implementation**
- 3. Designing low-friction evidence generation**
- 4. Broadening perspectives through partnerships.**

Approach: Working with schools to generate meaningful evidence

1. Co-creation and strong education relationships

Fundamental to Teachmate's approach is a commitment to embedding collaboration with educators throughout the product lifecycle. Rather than treating evidence generation as a separate activity, Teachmate integrates dialogue with educators into the ongoing development, implementation and refinement of their product.

Central to this approach is the belief that schools should be active participants in the process, rather than simply recipients of products or subjects of evaluation activity. The organisation emphasises the importance of dialogue, participation and co-creation, describing schools less as 'users' and more as ongoing 'partners'. As CEO Dr Jon Chippindall explained: 'Schools don't want to be passive recipients of what you're going to come in and test - they want to be part of it.'

Rather than focusing on institutional relationships, Teachmate prioritises relationships with the educators themselves: learning how teachers adapt workflows, share frustrations, refine outputs and shape how products develop through everyday use.

Many of Teachmate's strongest educator relationships emerged through existing educator communities and ongoing professional dialogue, rather than formal evaluation structures. This relationship-led approach provides a continuous source of insight into how products are experienced in authentic educational settings.

2. Staying close to implementation

One of the ways Teachmate operationalises this is through a team of dedicated account managers, whose role extends beyond onboarding or technical troubleshooting. These relationships are an important mechanism for maintaining ongoing dialogue with schools, creating opportunities to understand how implementation evolves over time.

Importantly, this support is intentionally contextual. The organisation recognises that schools implement products in different ways, encounter different barriers and adapt tools to meet the needs of their unique workflows. Rather than viewing this variation as problematic, Teachmate sees it as one of the richest sources of implementation insight.

This approach reflects a deliberate move beyond relying solely on surface-level engagement metrics. As Jon notes, 'Sometimes the data tells you people are using something. That doesn't necessarily mean it's working well for them.'

Instead, Teachmate seeks to understand how teachers are adapting products in practice, what friction emerges and what adaptations educators make to meet the needs of their daily realities. Some of the most valuable insights have emerged from the unexpected ways teachers are using the product, which only become visible once products are embedded in real classrooms.

3. Designing low-friction evidence generation

The same philosophy informs Teachmate's approach to evidence generation. Rather than treating adaptations made by educators as implementation failures, the organisation is interested in understanding what they reveal about workflow needs, classroom pressures and teacher decision-making.

This mindset shapes how the organisation approaches evidence generation itself. Multiple opportunities for feedback have been embedded into the product workflows so that educators can share insights during everyday use, rather than through detached evaluation exercises or retrospective surveys. In practice, this creates a more continuous relationship between implementation, feedback and refinement.

At the same time, the organisation is candid about the tensions this can create. Schools frequently request different things depending on their context, staffing structures and implementation confidence, and not every request can be translated into the development of the product.

Instead, the organisation focuses on stepping back from isolated requests and identifying the broader patterns underpinning them. The challenge is not simply to gather more feedback, but to understand what that feedback actually reveals about implementation needs across different educational contexts.

4. Broadening perspectives through partnerships

Teachmate also seeks to avoid relying too heavily on feedback from a small number of highly engaged groups. To broaden its understanding of implementation across different contexts, the organisation has developed partnerships beyond its core school network.

One example includes their partnership with the University of Manchester, where Teachmate's tools have been explored within teacher education contexts, including supporting non-specialist science teachers and developing educator confidence around responsible AI use.

Importantly, these partnerships are not intended to provide validation or promotion. Instead, they function as additional opportunities to understand how different educators interact with products in practice and where assumptions or implementation barriers become visible.

Teachmate maintains that products are intended to support professional judgement rather than replace it. Teachers remain active decision-makers within workflows, adapting outputs, refining materials and exercising discretion in how products are used within classrooms. These teacher adaptations are treated as valuable evidence in their own right, offering insight into classroom realities and professional practice.

Reflections and advice for the sector

Reflecting on Teachmate's experience, when building collaborative relationships with educators, Jon emphasised the importance of depth over scale: 'I'd probably do less... better. I'd probably go deeper.'

Jon reflected that scaling rapidly across multiple areas can make it harder to maintain the depth of collaboration and implementation understanding that produces the richest insight. Going deeper within fewer areas, he suggested, may create stronger opportunities for refinement, evidence generation and meaningful partnership with educators.

Taken together, Teachmate's approach illustrates how stronger evidence cultures can emerge through sustained collaborative relationships with educators. Rather than relying on discrete research activities, the organisation has intentionally built structures that connect implementation insight, educator feedback and product refinement to interact continuously over time.

For other EdTech organisations, the key lesson is that schools are not simply places where products are deployed or evaluated. They are active partners in understanding how products work in practice and in generating evidence needed to support continuous improvement.



Case study two

Learning by Questions: Evidence as product intelligence

This case study explores how EdTech organisations can move beyond treating evidence as a retrospective validation exercise and instead use it as a form of product intelligence that actively shapes product development and strategic decision-making.

Across the EdTech sector, evidence is often positioned primarily as a mechanism for communicating impact externally. However, organisations that embed evidence more intentionally into product development processes are often better placed to not only articulate impact more credibly, but to improve the quality, effectiveness and responsiveness of their products over time. In this sense, evidence-informed development can become a pathway towards stronger efficacy.

Context: Introducing Learning by Questions

Learning by Questions (LbQ) is an assessment and retrieval practice platform designed to support formative assessment, curriculum alignment and responsive teaching. As part of the EdTech Evidence Board (EEB) pilot, LbQ submitted SATs Springboard for review: a product designed to support SATs preparation through retrieval practice, diagnostic insight and targeted feedback.

Throughout both the portfolio review and subsequent discussions, LbQ emerged as a strong example of how evidence can be embedded within organisational decision-making, rather than used solely to demonstrate impact retrospectively.

This case study draws on insight generated through the independent review and an interview with Managing Director Greg Adam to explore how the LbQ team uses evidence to interpret classroom practice, identify misconceptions, refine product design and shape ongoing development priorities. In particular, discussions around SATs Springboard highlighted a close relationship between:

- 1. Educational intent as an organising principle**
- 2. Product refinement through evidence**
- 3. Interpreting evidence in context**
- 4. Evidence maturity and organisational learning**

Approach: Using evidence as product intelligence

1. Educational intent as an organising principle

Participation in the EEB pilot provided the LbQ team with an opportunity to reflect more deeply on how evidence, implementation insight, and product development interacted over time. Rather than treating evidence as an external validation mechanism, the organisation demonstrated a more iterative approach in which evidence is used to interpret classroom practice, understand misconceptions, refine product design and inform ongoing priorities.

A consistent theme throughout the review was the organisation's focus on educational intent. The organisation places great emphasis on understanding why particular features, tools, or changes are being introduced and what educational problems they are intending to solve. Product design decisions are guided not simply by engagement metrics or technical functionality, but by whether features genuinely support retrieval practice, formative assessment and classroom decision-making in the ways intended.

This relationship between evidence and development was particularly visible in the organisation's use of pupil misconception analysis. SATs Springboard generates large volumes of pupil response data, but the emphasis is not placed on the data itself. Instead, the LbQ team focuses on interpreting patterns within pupil responses to better understand misconceptions, curriculum understanding and learning behaviours over time. Importantly, these patterns are not treated as definitive answers, but as prompts for further investigation and refinement.

2. Product refinement through evidence

The LbQ team seeks to refine their product through identifying recurring patterns in pupil responses. One example of this involved reviewing question-level performance data to identify recurring misconceptions within a particular question set. As Managing Director Greg explained, the goal is not merely to determine whether pupils have answered correctly, but to understand why misunderstandings were occurring and how the product might better support teachers and pupils in response, "We're not interested only in whether pupils got it wrong. We're interested in why."

Through this process, the team were able to make targeted revisions and introduce misconception-focused feedback. Over time, this process has also enabled subject expertise to become increasingly embedded within the product. This example illustrates how analytics can move beyond measuring engagement to generating actionable insight capable of enhancing instructional quality.

The organisation also makes a considered effort to use multiple evidence sources to inform decision-making. Teacher feedback, implementation insights, classroom observation, published research and usage analytics are interpreted together, rather than treated as isolated or definitive indicators of success. For example, teacher feedback may identify usability issues but not learning impact; analytics may reveal patterns of engagement but not pedagogical quality. Instead, by combining perspectives, the organisation seeks to develop a more contextualised understanding of product performance and development priorities.

3. Interpreting evidence in context

The LbQ team highlights the challenges involved in translating educational research into operational product decisions. Large-scale research findings, while valuable, do not always provide straightforward answers about how products should function within classrooms or how pedagogical principles should manifest within product design.

This tension is particularly visible in implementation testing. The organisation approaches different forms of product testing proportionately: minor interface adjustments may require only internal review, while changes likely to affect classroom practice need testing within authentic educational settings. As Greg reflected, “the data tells you what is happening. It doesn’t tell you why.”

One example involved classroom testing of a new feature. Internally, the feature initially appeared promising, but classroom testing quickly surfaced practical implementation issues. A simplistic reading of the analytics might have suggested that, for example, teachers simply did not want or need the feature. However, classroom observation indicated that the implementation conditions themselves were not yet functioning effectively.

This distinction highlights the importance of contextualising evidence within EdTech environments. Low usage may indicate low value, but it may also indicate poor implementation, insufficient onboarding, workflow friction or unintended classroom consequences. The LbQ team’s development approach exercises caution about jumping too quickly to conclusions from isolated data signals.

4. Evidence maturity and organisational learning

Another notable feature of the organisation’s approach was its keen awareness of the risks associated with narrow evidence collection. The LbQ team highlights the importance of testing assumptions across varied contexts, users and implementation conditions, rather than relying too heavily on the same enthusiastic schools. Indeed, the organisation’s approach is underpinned by a commitment to challenging assumptions and resisting the temptation to use evidence simply to confirm existing belief. As Greg put it, “It’s okay to be wrong. It’s okay to make mistakes as long as you dig into why”.

This openness to challenge is closely connected to the organisation’s evolving evidence maturity. Greg Adam reflected that many of the principles now formalised through theories of change and structured evidence processes had existed informally within their development thinking much earlier.

This reflects an important wider insight from the EEB pilot: mature evidence cultures do not necessarily emerge fully formed through formal evaluation frameworks alone. In many organisations, they develop progressively as teams begin making their assumptions, reasoning and intended outcomes more explicit and systematic over time.

Concluding thoughts

Taken together, LbQ’s approach to SATs Springboard illustrates how evidence can become embedded within the ongoing life of a product. Evidence is used not simply to demonstrate impact retrospectively, but to interpret misconceptions, test implementation, refine feedback, challenge assumptions and inform future development priorities. In this sense, evidence becomes more than a proof point; it becomes part of how the product, and the team behind it, learns. This is a continuous process for LbQ, as Greg put it: “We never believe we’re done.”

For other EdTech organisations, the transferable lesson is clear: evidence is most powerful when it becomes actionable. Strong evidence strategies do not simply ask, ‘Can we demonstrate impact?’ They ask, ‘What are we learning about how the product works, what assumptions are being challenged, and what should improve next?’

EdTech Evidence Board

Chartered College
of Teaching

The Professional
Body for Teachers

The Chartered College of Teaching is
incorporated by Royal Charter and a
registered charity (number 313608).

edtech@chartered.college | 020 3433 7624

