ImpactEd

Kami

Initial Evidence Report

December 2023

Contents

1. Introduction

Kami is a versatile PDF annotator and interactive teaching platform that not only engages students with accessible tools but also has the potential to saves teachers time through its ability to streamline the ways in which teachers can provide instruction and carry out assessment. Kami integrates with a number of other learning tools, like Google Classroom, Microsoft Teams, Schoology, and Canva. These integrations can help further simplify student and teacher workflows and provide new ways to help teachers monitor their students' progress.

Kami partnered with ImpactEd Evaluation in 2023 to produce three pieces of research. The first was a rapid literature review of some of the main challenges in the education sector that tools like Kami might be able to help alleviate. The second was a logic model, designed to build upon the findings of the literature review and articulate the mechanics of how Kami is intended to help both teachers and learners, and address parts of these challenges. The third is this report. Using some of the key outcomes identified in the Kami logic model, we designed an evaluation that would balance methodological robustness with speed and ease of deployment. Collecting data from teachers and students in a selection of secondary schools and Sixth Form colleges in the UK, we sought to understand how Kami users compared to non-Kami users against a number of key outcomes. This report outlines our findings.

Because this report only features data from a single point in time, it cannot answer many of the more complex questions about the true impact of Kami. For example: how does teacher workload change over time, as teachers learn to use Kami to streamline their lesson planning workflow? How does students' growth mindset change, as they learn to work more collaboratively and receive more tailored feedback? And how, ultimately, does the long-term use of Kami result in changes in crucial learning outcomes, like student attainment, emotional resilience, or post-education plans? We envisage this report constituting the first piece of work in a longer-term partnership, during which we would seek to answer some of these trickier questions. As such, we focus here on benchmarking some of the key outcomes for both teachers and students against some available comparison figures.

2. Methodology

This section outlines the methodology we followed to put together this report. Section 2.1 outlines the construction of a Logic Model and accompanying literature review. Section 2.2 discusses some of the evaluation design choices that were made, including the choices of survey measures and the selection of participants. Section 2.3 provides an overview of how we approached data collection, and how we analysed the data once it was collected. In Section 2.4, we outline some of the limitations of this evaluation.

2.1 Logic Model & Literature Review

To ensure that the evaluation design was grounded in an up-to-date understanding of the current education landscape, we decided to conduct a rapid literature review (RLR) help shape the evaluation design. One of the RLR's aims was to identify the key contemporary issues in education across the globe so that our research could position Kami's role in alleviating some of these issues. The other, more important, aim was to explore what the academic literature said about the impact of some of Kami's core features on pupils' learning. This would help us to focus on measuring the outcomes for which there was the strongest evidential support in the academic literature.

The RLR on contemporary global issues in education focused on:

- Changes in schools' budgets and funding
- The influence of socio-economic differences on pupil outcomes
- Disparities in pupils' social and emotional learning skills

The RLR on Kami's functionalities focused on the following key functionalities:

- Annotation
- Formative assessment
- Collaboration in learning
- Accessibility in learning
- Personalised learning.

It is important to note that this RLR was not a systematic literature review. A structured keyword search approach was used for both goals of the review. The review on global issues in education used Google; the review on Kami's functionalities used JStor and Google Scholar. The full RLR can be found in Appendix A2, including further details on the methodological approach.

The output from the RLR helped to inform the design of the Logic Model in two ways. Firstly, the findings on global issues in education shaped the articulating of problems in the Logic Model. Secondly, the RLR of Kami's functionalities informed what inputs and activities were

detailed in the Logic Model as well as pulling out the key intended intermediate outcomes and long-term outcomes. Kami achieved their goal of using evidence to guide their evaluation design.

2.2 Evaluation Design

The evaluation design was steered strongly by the findings of the rapid literature review and the logic model. In collaboration with Kami, we selected a small number of intermediate outcomes for both students and teachers from the logic model, and we mapped these to appropriate data collection measures. Where possible, we mapped outcomes to validated instruments from the academic literature, prioritising those with available benchmark figures. Where no suitable validated instrument was available, we designed bespoke questions to probe specific outcomes.

We employed a mixed-methods approach for teachers, collecting both quantitative and qualitative data and using each one to inform the other to build up a richer understanding of Kami's impact on these users. Owing to the difficulty of conducting qualitative research with students at short notice, we opted to collect only quantitative data for students via an online survey. This was facilitated by their teachers.

The timelines for this work meant that we needed to design an evaluation that could be deployed quickly, and that would make use of Kami's existing user base. Rather than employing a pre-post study design, which would have required surveying a group of new Kami users before onboarding, and then again, some time after prolonged use of Kami, we employed a single-timepoint evaluation design, with data being collected at one point only. Where possible, we either compared this data to available national averages, benchmark figures from the academic literature, or data collected from a small control group (in the case of teachers).

2.3 Data Collection & Analysis

2.3.1 Sampling & Recruitment

This evaluation relied on a convenience sample. We asked Kami to select up to 10 schools in which they knew Kami was well-embedded across the teaching staff, and who had been users for at least a month. Kami then approached these schools, and asked them if they would be willing to participate in data collection with us (via online surveys with both teachers and students, and additional focus groups for teachers). Four schools indicated that they would be willing to participate in the evaluation. This sampling method is a limitation of the design (cf. Section 2.4).

Table 1 – Sample sizes for the evaluation.

Group	Quantitative Data	Qualitative Data
Schools	4	4
Students	685	-
Teachers	35	5

2.3.2 Quantitative Data Collection & Analysis

Quantitative data was collected throughout November 2023, from both students and teachers at the schools that had agreed to participate in the evaluation. Links to the online surveys were sent to our contact teachers to disseminate among students and teaching staff. From the four schools, we received quantitative data from 685 students and 36 teachers.

We attempted to collect data from teachers that did not use Kami, in order to construct a small-scale control group. The control survey began with a filter question that excluded any respondent who had, indeed, used Kami before, and then asked the same questions as the participant survey, with the exception of the questions around confidence using Kami. This control survey was disseminated on social media accounts by ImpactEd Evaluation and Kami, which led to a very large number of responses. Unfortunately, the vast majority of these were spam responses. To verify that the respondents were, indeed, school teachers, we had asked them to provide their school email address, and incentivised them to do so by entering them into a prize draw for a £20 Amazon voucher. After excluding the spam responses, as well as any response that did not provide a valid school email address, the control group size was only 4.

Quantitative data analysis consisted of simple descriptive statistics (primarily the calculation of sample means) on both aggregated data (for subscales and scales), and item-level data for individual questions. Where scales included both positive and negative polarity questions, we adjusted the data to ensure that aggregation remained consistent.

Since we had also collected a selection of demographic data for pupils (such as gender, year group, and school), we segmented the data to examine differential impact. For teachers, we segmented the data by duration and frequency of Kami usage.

Wherever possible, we compare scale-, subscale- and item-level averages to either national averages (mean scores from all-time responses to the survey on ImpactEd Evaluation's School Impact Platform). We also draw tentative comparisons to the teacher control group, but the small sample size here is a clear limitation of the evaluation (cf. Section 2.4).

In Section 4.1, we use a survey measure called the Self-Directed Learning with Technology Scale. Although there is a benchmark figure available for this scale in the academic

literature, we elect not to use this benchmark in this study because we do not believe it to be an externally valid comparison for the context of this research. The benchmark figure available was constructed from a sample of Singaporean high school students, and we believe that the UK and Singaporean educational contexts are not sufficiently similar to license direct comparisons between data collected in either setting.

It was not possible to conduct statistical significance testing in this report, owing either to one-shot nature of the data (i.e. we only have one set of outcome data, so there is no difference against which to assess significance), or owing to the small size of the sample group (it is not best practice to conduct statistical significance testing with a samples smaller than n=30).

2.3.3 Qualitative Data Collection & Analysis

Qualitative data was collected throughout November 2023, from teachers at the schools that had self-selected to participate in the evaluation. To aid recruitment, we offered each participant a £20 Amazon voucher. Both of these factors (self-selection and incentivisation) introduce a possibility of bias (cf. Section 2.4).

We collected data from six participants across two focus groups and one 1-to-1 interview. Participants came from three schools, and covered a range of teaching experience and subject specialisms, with three being Science teachers and two being Maths teachers. The remaining participant was a trainee teacher who had used Kami very little. For this reason, we excluded them from the analysis.

The remaining participants had been using Kami for at least two months, with four of the teachers having started using Kami in September 2023 and the other teacher having been using Kami for over 5 years. Almost all the teachers expressed that they used Kami daily.

Once collected, we transcribed the audio file using an audio transcription software, and then cleaned the transcript to ensure fidelity. We employed a deductive thematic approach to our analysis, using as our coding framework the set of outcomes in which we were interested (namely, teacher workload, passion for teaching, confidence providing differentiated instruction / feedback, and so on). We present the findings of the qualitative analysis on an outcome-by-outcome basis throughout Section 3.

2.4 Limitations

- 🔺 **One-shot data collection.** The data in this evaluation was collected at one time point, after schools had been using Kami for a sustained period. As such, we cannot use this data to understand how any of the outcomes examined have changed over time. This, in turn, limits our ability to attribute any impact to Kami itself.

- **Lack of control group / small sample size control group.** As noted above, although we constructed a small control group for teachers, the data collection here was challenging and when we controlled for spam responses, the resulting control group is very small. For student outcomes (Section 4), we do not have any control group, though we have compared to available benchmark figures (from our own wider dataset, or from academic studies) where possible.

- **Selection bias among participating schools.** The schools participating in this evaluation self-selected, rather than being selected randomly for participation. This introduces the possibility of bias into the data, and future evaluation work could consider adopting a more robust approach to recruitment.

- **Diversity of school types.** The sample here comprises secondary schools, sixth form colleges, and both in-person and online schools. As such, it is perhaps not a representative sample of the school landscape. It also means we should be wary of making school vs. school comparisons (hence our decision not to conduct school-level subgroup analysis).

3. Impact on Teachers

This section examines the trends in the data we collected against three outcomes. First, we examine how long Kami users spent on different teaching-related tasks, and compare these figures to a small comparison group of teachers that do not use Kami. Next, we look at how passionate Kami teachers were about teaching. Lastly, we look at data from both the teacher survey and the pupil survey to understand whether teachers using Kami were confident providing differentiated feedback and instruction to their students.

3.1 Teacher Workload

3.1.1 Overview

One of the intermediate outcomes in Kami's logic model is that *the amount of time that teachers spend planning lessons decreases.* To understand how long teachers that used Kami spent on planning lessons, we asked them to indicate how long they spent, in the average week, on a range of teaching-related tasks. The responses to this question are shown in Figure 1.

Frequency of responses to time spent on teaching related tasks (n=35)

| Task	0 hours	1-3 hours	4-6 hours	7-9 hours	10-12 hours	13-15 hours	More than 15 hours
Communication with Parents / Guardians	11	20	1	1	1	1	
Professional development activities	10	20	4	1			
General administrative work	1	18	7	4	1	3	1
Marking student work	3	15	10	1	3	1	2
Planning or preparation of lessons	2	11	12	4	3	1	2

Figure 1

To calculate the average number of hours spent on each task, we converted each response to a number by taking the midpoint of the chosen bracket (e.g. '4-6 hours' became 5 hours). The mean number of hours spent on each task is shown in Figure 2, where we compare teachers using Kami and teachers not using Kami.

Average number of hours per week spent on different teaching tasks for Kami users and non-Kami users (user n=35; non-user n=4)

Task	Teachers not using Kami	Teachers using Kami
Communication with Parents / Guardians	2.25	2.23
Professional development activities	2.00	1.94
General administrative work	3.50	4.91
Marking student work	6.75	4.77
Planning or preparation of lessons	8.25	5.51

Hours spent per week

■ Teachers not using Kami ■ Teachers using Kami

Figure 2

It is noteworthy that although users and non-users spend comparable amounts of time on certain tasks that do not involve Kami (such as communication with parents, or professional development), Kami users spend substantially lower amounts of time on the tasks that Kami is well-suited for: marking student work, and planning and preparing lessons. **Kami users spent 33% less time planning and preparing lessons than non-Kami users, and 29% less time marking student work than non-Kami users.** This trend is reflected in the qualitative data with one teacher going as far to say that using Kami has enabled him to sometimes spend no time planning lessons.

We should be wary of drawing firm conclusions from this trend for two reasons. First, because we are looking at data from a single timepoint, we cannot confidently say that there has been a decrease in time spent on these tasks associated with the use of Kami, or whether these teachers have always been faster. Second, the control group is extremely

small: only 4 teachers, and this small control group prevents us from conducting any meaningful statistical significance testing. As such, we cannot be entirely confident that this is a genuine difference between the two groups. However, the fact that this trend is so pronounced, and that the results align well for other tasks where Kami would not be used, lends some credibility to the result.

3.1.2 Subgroup Analysis

To understand how Kami usage affected workload, we segmented the data into categories according to the number of times each respondent used Kami in the average month. The results are shown in Figure 3. There are a number of trends worth commenting on.

- **There appears to be a consistent 'sweet spot' for Kami users who used Kami 6-10 times per month, or approximately twice per week.** These users consistently spent the least time on all teaching related tasks, with the exception of professional development, where they were able to spend more time than their peers. This could have interesting implications for Kami's onboarding / embedding guidance, and may warrant further investigation.

- **Teachers that used Kami 6-10 times per month spent 54% less time marking student work than teachers that never used Kami in the average month**: Teachers that never used Kami in the average week spent, on average, 6.50 hours per week marking student work. Teachers that used Kami more regularly (6-10 times per month), spent only 3.00 hours per week marking student work. This number jumps back up to 4.36 hours per week for teachers using Kami more than 10 times per month.

- **Teachers that used Kami 6-10 times per month spent 29% less time planning lessons than teachers that never used Kami in the average month**: Teachers that never used Kami in the average week spent, on average, 5.63 hours per week preparing lessons. Teachers that used Kami more regularly (6-10 times per month), spent only 4.00 hours per week marking student work. This number jumps back up to 5.93 hours per week for teachers using Kami more than 10 times per month.

Average number of hours per week spent on different teaching tasks by Kami usage (n=35)

Task	Never	1-5 times per month	6-10 times per month	More than 10 times per month
Communication and cooperation with parents or guardians	3.00	2.43	1.00	2.21
Professional development activities	1.13	2.57	2.17	2.00
General administrative work (including communication, paperwork, and other clerical duties)	5.63	5.86	3.00	4.79
Marking / correcting student work	6.50	4.86	3.00	4.36
Individual planning or preparation of lessons, either at school or outside of school	5.63	5.57	4.00	5.93

Figure 3

3.1.3 Qualitative Findings

Overall, **teachers expressed that Kami had helped them save time** and some of them identified ways that Kami could help them save time in the future. Before examining which of Kami's functions enables teachers to save time, it is important to note that many teachers highlighted that although Kami saved them time in a range of ways, it sometimes led to them spending more time on their role. Teachers reported that they repurposed their saved time on more productive activities, that benefitted pupils even further, but sometimes they spent more time on their role than they would have before Kami; crucially, however, Kami enabled them to better support pupils. One teacher highlighted that **it was hard to imagine**

teaching without Kami and that Kami was instrumental in his capacity to complete all his teaching responsibilities whilst still having time to raise his family.

Almost all the teachers interviewed highlighted that Kami had **saved them time in planning their lessons**. The consensus was that although Kami's functionality was not used directly in the activity of planning a lesson, **planning was easier and faster when they knew they would use Kami in their class**. One teacher went as far to say that knowing they would use Kami's interactivity meant that **sometimes they did not spend any time at all planning lessons**. Other teachers expressed similar thoughts; knowing that Kami provided them with the freedom to do more in their classroom, and **not having to think about the logistics of their teaching, made it easier to plan lessons**.

Teachers also reported that **Kami saved them time by enabling them to see pupils complete their work live and increasing the ease of delivering feedback to pupils**. One teacher expressed that simply not having to ask pupils to share their screen saved time. Teachers expressed they saved time marking work through Kami because they did not have to wait for pupils to submit their work, they could do it from a distance, **they could identify common mistakes and commonly misunderstood concepts faster**, and they could **copy and paste feedback where those common mistakes occurred**. Teachers also reported that pupils were then able to rectify mistakes immediately; this implied **that Kami also saved pupils' time**. One teacher quantified the amount of time they had saved by using Kami to mark pupils' work.

> *"[The time taken to prepare lessons] is probably half the amount of time it would have taken me before."*
>
> *Teacher 1*

Using Kami as a resource centre and editor was also another way that teachers reported Kami helped them save time. Three teachers highlighted that Kami had saved them a lot of time by **not having to print resources for pupils**. They also reported that Kami's functionality enabled them to use old resources they had not been able to use before in a digital format which saved them time in not having to design a new resource. Teachers also communicated that Kami's **PDF editor and capacity to put documents together** had also saved them time. They also detailed **that Kami made it easy for pupils to access a wide range of resources**, such as recorded lessons and extension tasks, at any time which saved teachers' time because pupils could independently access materials. One teacher also mentioned that using the worksheets on Kami had saved them time; another expressed that he could understand how Kami's templates could be helpful in helping new teachers' save time.

Teachers also mentioned a few other ways they had perceived Kami saving them time. One teacher mentioned Kami's ability to integrate with other applications. **Another teacher highlighted that Kami's OCR (Optical Character Recognition) software had helped them save**

time trying to manipulate text in images and another teacher said Kami had helped them save time by distributing assignments.

3.2 Passion for Teaching

3.2.1 Overview

To understand teachers' attitudes towards teaching, and to probe how these might differ between teachers that use Kami and teachers that do not use Kami, we used the Engaged Teachers Scale.[1] This is a 16-item Likert scale instrument probing teachers' feelings towards their work, their students and colleagues, and teaching in general. Both the group of teachers using Kami and the comparison group completed this survey. The mean score on this instrument for both groups is shown in Figure 4.

Mean Engaged Teachers Scale score by experimental group (user n=35; non-user n=4)

|	5.95	5.27

Teachers using Kami   Teachers not using Kami

Figure 4

Kami users scored slightly higher passion for teaching than non-Kami users (+0.68), indicating that **teachers using Kami were, on average, 13% more passionate about teaching than teachers not using Kami.**

---

[1] Klaassen et al. (2013)

Mean scores for individual items from the Engaged Teachers Scale relating to teaching, in general, are shown in Figure 5. Kami users substantially outperformed non-Kami users on all four questions. This was reflected in the qualitative data collected; all teachers expressed in one way or another that Kami had made teaching less stressful and more enjoyable.

Mean item scores for Engaged Teachers Scale items relating to teaching, by experimental group (user n=35; non-user n=4)

- "I love teaching." — 5.47 / 4.75
- "I am excited about teaching." — 5.39 / 4.00
- "I feel happy while teaching." — 5.67 / 4.25
- "I find teaching fun." — 5.50 / 4.00

■ Teachers using Kami   ■ Teachers not using Kami

Figure 5

Mean scores for individual items from the Engaged Teachers Scale relating to teachers' attitudes towards their colleagues, are shown in Figure 6. Kami users outperformed non-Kami users on all four questions, to varying degrees.

Mean item scores for Engaged Teachers Scale items relating to colleagues, by experimental group (user n=35; non-user n=4)

Figure 6

Mean scores for individual items from the Engaged Teachers Scale relating to teachers' attitudes towards their work, are shown in Figure 7. Here, the overall trend is not so pronounced. Kami users outperform non-Kami users on some questions, but not others. Interestingly, one of the statements that Kami users agreed less with, in comparison to non-Kami users, was "while teaching, I work with intensity". Although this appears negative, at first glance, it could be argued that not feeling the need to work with such intensity reflects the fact that the workloads of Kami users was, on average, lower than that of non-Kami users. This is speculative, however, and data would be required to confirm this conjecture. Focus groups and interviews with teachers using Kami did reflect the view that Kami allowed teachers to be more flexible and less constrained in their teaching which could make them feel like they are working with less intensity.

Mean item scores for Engaged Teachers Scale items relating to work, by experimental group (user n=35; non-user n=4)

Figure 7

Mean scores for individual items from the Engaged Teachers Scale relating to teachers' attitudes towards their students, are shown in Figure 8. Here, Kami users scored higher than non-Kami users on three out of the four questions. On the other question, the two groups scored equally highly.

Figure 8

3.2.2 Subgroup Analysis

We segmented the data by usage of Kami to understand how usage effects teacher engagement and passion for teaching. The results are shown in Figure 9

Figure 9

The data suggests that there is only a minor difference in engagement between the different levels of usage. Those who never used Kami in the average month at school scored on parity with those who used it over 10 times in the average month.

3.2.3 Qualitative Findings

Analysing data from the focus groups and interviews with teachers, **Kami has helped increase teachers' enjoyment of teaching as well as making it easier for them to teach, but it is not clear that Kami has helped teachers increase their passion for teaching**. One teacher reported that Kami had helped them to enjoy teaching again because it had **enabled them to support pupils again rather than before they had used Kami where they were just assigning them tasks**. Multiple teachers highlighted that using Kami had made them feel **less restricted and constrained** and that it had **helped them streamline their work**. Using Kami's functionality meant they **spent less energy worrying about the logistics** of delivering their sessions and more time teaching. They reported that this **made their job much easier and less stressful**. One teacher highlighted that being able to deliver immediate feedback within **Kami had helped him feel like a more effective teacher in areas he had previously perceived as an area of weakness**. Another teacher communicated that he had always enjoyed teaching and that Kami had not increased that inherent enjoyment he felt.

3.3 Differentiated Instruction & Differentiated Feedback

3.3.1 Overview

We asked teachers using Kami and teachers not using Kami two pairs of custom, 1-5 Likert-scale questions to gauge their feelings towards providing differentiated feedback and differentiated instruction for their students. Each pair consisted of a positive polarity question (where higher scores were 'better'), and a negative polarity question (where lower scores are 'better'). The pair of statements for differentiated instruction, and the average scores for each group of individuals, are shown in Figure 10Figure 10.

Mean scores for statements related to differentiated instruction, by experimental group (user n=35; non-user n=4)

Statement	Teachers not using Kami	Teachers using Kami
"When students don't understand the instructions I've given them, I struggle to find other ways to explain what I'm asking them to do."	2.25	2.28
"I feel well equipped to provide differentiated instruction to each one of my students that is tailored to their needs."	4.25	3.81

Figure 10

Both groups had near-identical average scores for the first statement (2.28 for users, vs. 2.25 for non-users). However, Kami users were less likely to agree with the second statement, "I feel well-equipped to provide differentiated instruction to each one of my students that is tailored to their needs". Users scored, on average, 3.81 for this question, compared to non-users who scored, on average, 4.25. This suggests that **teachers who used Kami were 10% less likely than teachers who didn't use Kami to feel well-equipped providing differentiated instruction to their students.** However, the data does not allow us to be confident about the direction of the causal effect here. Rather than Kami use driving a lower level of confidence provided differentiated instruction, it could instead be the case that teachers who worry about providing differentiated instruction to their students are more likely to seek out tools like Kami. This pattern is reflected in the qualitative data; although some teachers provided some instances where they had been able to use Kami to differentiate instruction, such as providing extensions to higher-ability pupils, they were not as forthcoming with examples of differentiated instruction as they were when it came to differentiated feedback, suggesting that perhaps Kami is not being used in this way. Further research would be required to fully understand this trend.

The pair of statements for differentiated feedback, and the average scores for each group of individuals, are shown in Figure 11.

Mean scores for statements related to differentiated feedback, by experimental group (user n=35; non-user n=4)

"When students don't understand the feedback I've given them on their work, I struggle to find other ways to communicate what I mean."
- 1.50
- 2.39

"I feel well equipped to provide differentiated feedback to each one of my students that is tailored to their needs."
- 3.75
- 4.03

■ Teachers not using Kami ■ Teachers using Kami

Figure 11

Kami users were less likely than non-users to agree with the statement "I feel well-equipped to provide differentiated feedback to each one of my students that is tailored to their needs", scoring, on average, 4.03, compared to 3.75, on average, for non-users. This suggests that **teachers who used Kami were 7% more likely than teachers who didn't use Kami to feel well-equipped to provide differentiated feedback to their students.** This data is reflected in the experiences described by teachers in focus groups and interviews; teachers communicated that seeing pupils' work being completed live meant that they felt better equipped to provide differentiated feedback because they could deliver it to pupils much quicker.

Interestingly, there is a much more pronounced difference in the first statement. When asked the extent to which they agreed with the statement "when students don't understand the feedback I've given them on their work, I struggle to find other ways to communicate what I mean", users scored, 2.39, on average, indicating mild disagreement or neutrality. Non-users, on the other hand, scored 1.50 on average, indicating stronger disagreement. On the face of it, this suggests that Kami users are more likely to struggle to give feedback in different ways, when the way they've chosen doesn't work. However, it is worth noting the different implications of the statement for the different groups. Kami users, after all, have

many more modes of feedback available to them. For a teacher that doesn't use Kami, the statement might imply that they have tried one or two favoured methods of feedback, and so they may be aware of others that they can try if required. For teachers that use Kami, however, the statement might instead imply that the teacher has already tried all the different options for differentiated feedback that Kami provides, and so they may struggle to think of yet more approaches.

This hypothesis is particularly compelling when drawing upon the qualitative data, in which teachers told us that Kami improved their ability to provide differentiated feedback. Many teachers highlighted that the feedback tools in Kami were incredibly effective in a range of ways. Teachers communicated that the voice notes feature enables pupils to explain their thought-process which facilitated teachers to pinpoint the blocker in the pupils' learning and then providing specific feedback or changes in the pupils' approach. Additionally, voice notes enabled pupils who were less confident in reading and writing of English in comparison to their oral skills to express what they did or did not understand more clearly to their teacher.

One can imagine how this heightened ability to draw out pupil feedback and for teachers to tailor feedback when using Kami might lead Kami users to score lower on this question, since it seems conceivable that someone with access to a very wide range of options will struggle more than someone with more limited options to come up with a new approach.[2] It is not clear, however, why this logic should work for this pair of questions, but not the other (Figure 10). This narrative is speculative, and we would need to conduct further research to understand whether this is truly the case.

3.3.2 Subgroup Analysis

To understand how teachers' confidence providing differentiated feedback and instruction changed with their usage of Kami, we segmented the data into usage categories and graphed the average score in response to the statements above for each category. The results are shown in Figure 12.

In the case of differentiated instruction, there is a clear linear trend whereby confidence providing differentiated instruction increases as Kami usage increases. Those who never used Kami in the average month scored lowest on this question (3.22), whereas those who used Kami more than 10 times in the average month scored highest on this question (4.29). This suggests that teachers who consistently and regularly use Kami feel, on average, 33% more confident when it comes to providing differentiated instruction to their students.

---

[2] Put another way, the *marginal* difficulty of coming up with a new method for differentiated feedback is higher for those who already have a higher number of methods available to them.

In the case of differentiated feedback, the trend is less linear, with those using Kami 6-10 times in the average month scoring slightly higher than those using Kami 1-5 times in the average month. However, those who never used Kami in the average month scored lowest on this question (3.44), whereas those who used Kami more than 10 times in the average month scored highest on this question (4.36). This suggests that teachers who consistently and regularly use Kami feel, on average, 27% more confident when it comes to providing differentiated feedback to their students.

Mean scores for feeling well-equipped to provide differentiated feedback and instruction by usage of Kami in the average month (n=35)

"I feel well equipped to provide differentiated instruction to each one of my students that is tailored to their needs."
- 3.22
- 3.57
- 3.83
- 4.29

"I feel well equipped to provide differentiated feedback to each one of my students that is tailored to their needs."
- 3.44
- 4.14
- 4.00
- 4.36

■ Never ■ 1-5 times per month ■ 6-10 times per month ■ More than 10 times per month

Figure 12

3.3.3 Qualitative Findings

Teachers **clearly expressed that they had found Kami helpful in delivering differentiated feedback** whilst they were **more tentative that it had helped them deliver more differentiated instruction**.

Teachers reported that Kami had allowed them to **provide feedback anywhere and anytime**. This meant that they were able to deliver feedback extremely quickly; at times, teachers expressed that **they were able to deliver live feedback**. This then facilitated interactive feedback with pupils, creating **unique dialogues with each pupil** and therefore facilitating differentiated feedback. Teachers reported that this immediacy of the differentiated feedback had **encouraged pupils to act more quickly on their feedback** as it was fresher in their mind, meaning that **feedback was more effective and helpful to pupils**.

Teachers specifically mentioned two tools within Kami that enabled this differentiated feedback. **Kami's ability to allow pupils to voice note their thoughts on their own work meant that pupils were able to provide their own feedback on their work**, sometimes identifying their own problems. Pupil voice notes also **enabled teachers to pinpoint where pupils' thinking became confused**, which again facilitated super-targeted feedback and support for pupils. Voice notes also helped less competent English writers express their understanding to teachers, meaning that the **teacher could understand whether errors are occurring due to their ability to write clearly in English or whether it is because they do not understand the concept**. Again, this helps teachers provide appropriate and specific feedback. One teacher reported that using screenshots to demonstrate modelling was another way that Kami had helped them tailor feedback to specific pupils.

Some teachers reported that Kami had facilitated their capacity for differentiating instruction. One clear way it helps teachers to achieve this goal is that **it can be used in a class of varying ability**. Kami means that pupils can access need-specific resources, they can go through resources at their own pace, and can extend their work easily, although this functionality was more likely to be accessed by higher ability pupils. One teacher highlighted that **the range of tools Kami has to offer when teaching means that teachers can tailor tool use to pupils' needs**. They used an example of using a visual and interactive tool that could pull shapes apart to help explain compound shapes to lower-ability pupils. Another teacher expressed that being able to see pupils' work live in Kami meant that they would differentiate pupils' abilities in real-time. One teacher reported they had successfully used Kami to differentiate assignments, another teacher reported that they had not successfully been able to do this.

4. Impact on Students

This section examines the trends in the data we collected against three pupil outcomes from the Kami Logic Model (cf. Section 2.1 and Appendix A2). In Section 4.1, we look at how students taught using Kami scored on measures relating to collaboration and self-directed learning. Section 4.2 explores how motivated and excited about learning students taught using Kami felt. Section 4.3 looks at how students taught using Kami scored on measures relating to school engagement.

4.1 Collaboration & Self-Directed Learning

4.1.1 Overview

To understand the levels of collaboration and self-directed learning among students that used Kami, we used two subscales from a validated instrument that probed self-directed learning and collaborative learning respectively.[3] Together, these subscales made up an 8-item Likert-scale survey, scored on a scale of 1-7. The mean scores for the two subscales are shown in Figure 13.

Figure 13

Students using Kami above the scale midpoint on both subscales: 4.51 and 4.77 for self-directed learning and collaborative learning respectively. This corresponds to scores of

---

[3] Lee et al. (2014)

58.5% and 62.8% respectively. There is no valid comparison data available for this measure, and so we cannot say more about Kami users' scores relative to non-users.

Item-level mean scores for students using Kami are shown for the self-directed learning subscale in Figure 14 and for the collaborative learning subscale in Figure 15.
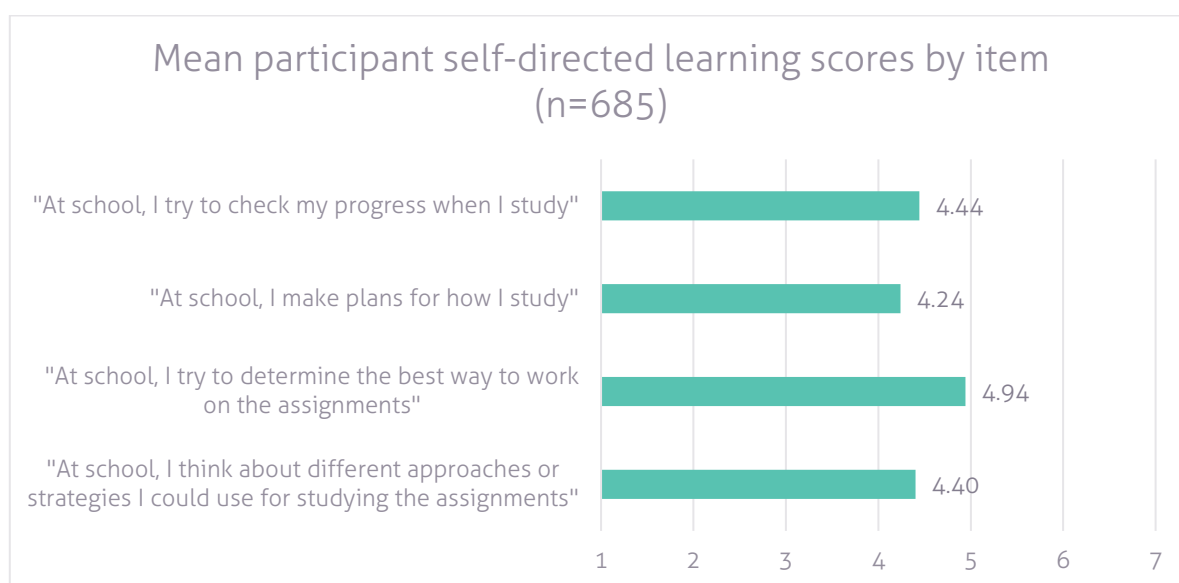
Figure 14

Figure 15

4.1.2 Subgroup Analysis

To understand the differential impact of Kami on self-directed and collaborative learning, we segmented the data by year group and gender. The results for year group are shown in Figure 16, and the results for gender are shown in Figure 17.

Figure 16

There are only minor differences in scores for collaborative and self-directed learning across the different year groups, and no clear trend emerges from the data overall. The year group with the highest scores is Year 11, and the year group with the lowest scores is Year 8. The data likely supports there being no meaningful difference in Kami's differential impact on collaborative and self-directed learning for different year groups.

Figure 17

Similarly, there are only minor differences when segmenting the data by gender. Male participants scored lower than female and other participants (4.54 for male, vs. 4.67 for both female and other). The data likely supports there being no meaningful difference in Kami's differential impact on collaborative and self-directed learning for different genders.

4.2 Excitement & Motivation for Learning

4.2.1 Overview

To understand the levels of collaboration and self-directed learning among students that used Kami, we used the intrinsic value subscale from the Motivated Strategies for Learning Questionnaire (MSLQ).[4] This is a 9-item subscale, scored on a 1-7 Likert scale. The mean scores for the subscale, along with the ImpactEd School Impact Platform national average, are shown in Figure 18.

---

[4] Pintrich & De Groot (1990)

Mean scores on the MSLQ Intrinsic Value subscale (user n=685; benchmark n=41,674)

Figure 18

Kami users scored marginally higher than the benchmark (the national average for all pupils that have ever taken this questionnaire through the ImpactEd School Impact Platform). **On average, students that used Kami had 4% higher motivation for learning than the national average.**

Item-level averages are shown in Figure 19. Although there are no item-level benchmark figures available, it is interesting to note that certain items score substantially higher than the scale-level benchmark, and certain items score substantially lower. High-scoring items include "It is important to me to learn what is being taught in my lessons", "Understanding my subjects is important to me", and "Even when I do poorly on a test, I try to learn from my mistakes". This last statement may be of particular interest to Kami, given that one of the longer-term outcomes in the logic model (cf. Appendix A2) not examined here is around developing learners' resilience and growth mindset. This high score may suggest that Kami students are more determined to continue trying despite setbacks, but further research would be required to confirm this. Low-scoring statements included "I prefer schoolwork that is challenging so I can learn new things" and "I often choose topics I will learn something from, even if they require more work".

Mean scores on the MSLQ Intrinsic Value subscale by item (n=685)

| Item	Score
"It is important for me to learn what is being taught in my lessons."	5.57
"I prefer schoolwork that is challenging so I can learn new things."	4.18
"I like what I am learning in my lessons."	4.70
"I think I will be able to use what I learn in one subject in other subjects."	4.89
"I often choose topics I will learn something from even if they require more work."	4.55
"Even when I do poorly on a test I try to learn from my mistakes."	5.07
"I think that what we are learning this year is interesting."	4.73
"I think that what I am learning in my lessons is useful for me to know."	4.92
"Understanding my subjects is important to me."	5.51

Figure 19

4.2.2 Subgroup Analysis

To understand the differential impact of Kami on self-directed and collaborative learning, we segmented the data by year group and gender. The results for year group are shown in Figure 20, and the results for gender are shown in Figure 21.

Figure 20

There are notable differences in mean motivation by year group. In KS3 and KS4 (i.e. Years 7-11), motivation peaks in Year 7 and Year 11, but falls slightly between these years. This may be explained by the fresh motivation of starting secondary school, and the impetus of passing GCSE exams in Year 11, and may not be a product of Kami itself. Similarly, Year 12 and Year 13 motivation scores are substantially higher than those of other year groups. However, at this stage, students are no longer in compulsory education, and so there is likely a survivorship effect taking place here: students choosing to continue to post-16 education are likely to be those who had high levels of motivation for learning in the first place. In sum, there are clear differences between motivation at different year groups, but these are perhaps better explained by external factors, rather than by Kami. This is supported by the fact that all figures are very similar to the appropriate benchmarks. The exception to this is the Year 12 figure, where Kami users outperform the benchmark considerably, and the Year 8 figure, where Kami users fall short of the benchmark considerably.

Figure 21

Interestingly, in a reversal of the trend we saw in the previous section, here male students outscore female and other students by some margin (5.69 for male students, 4.82 for other students, and 4.78 for female students). When comparing these figures to the national averages for gender, we see that although female students are essentially on parity with the benchmark, male students using Kami score substantially higher than the benchmark. No benchmark data was available for students with other gender identities. This suggests that **male students using Kami are 23% more motivated at school than the national average.**

4.3 School Engagement

4.3.1 Overview

To understand the levels of engagement among students that used Kami, we used the behavioural and emotional engagement subscales from the School Engagement Scale.[5] Together, this comprised a 10-item scale, scored on a 1-5 Likert scale. The mean scores for this scale, along with the ImpactEd School Impact Platform national average, are shown in Figure 22.

---

[5] Fredericks et al. (2005)

Mean school engagement scores, participant vs. benchmark (participant n=685, benchmark n=75,392)

Students using Kami — 3.37
Benchmark — 3.30

Figure 22

Students using Kami score 2% higher than the national benchmark (3.37 vs. 3.30). When looking at the data on an item-level, the difference in mean scores is quite apparent (Figure 23).

Mean scores on the School Engagement Scale by item (n=685)

Item	Mean score
"My classroom is a fun place to be"	2.92
"I am interested in the work at school"	3.30
"I like being at school"	2.92
"I feel excited by the work in school"	2.74
"I feel bored in school"	3.21
"I feel happy in school"	3.04
"I get in trouble at school"	1.66
"I follow the rules at school"	4.31
"When I am in class I just act as if I was working"	2.57
"I pay attention in class"	3.91

Figure 23

Interestingly, although students responded positively to statements like "I follow the rules at school", "I get in trouble at school", and "I pay attention in class", they responded less positively to statements like "I feel excited by the work at school". On average, they disagreed slightly with this last statement, which suggests that despite widespread use of Kami in the classroom, the students are still feeling unexcited by their work at school.

4.3.2 Subgroup Analysis

To understand the differential impact of Kami on self-directed and collaborative learning, we segmented the data by year group and gender. The results for year group are shown in Figure 24 and the results for gender are shown in Figure 25.

Figure 24

As might be expected, the trend for school engagement very closely tracks the trend for motivation, with there being small upticks in school engagement at the start of KS3 (Year 7) and at the end of KS4 (Year 11), possibly driven by starting secondary school and sitting external exams respectively. There is also a generally higher level of school engagement in Years 11 and 12, which one again would expect since students in Further Education will likely have higher levels of engagement. The school engagement of Kami users closely tracks the school engagement national averages, with the exception of Year 12, when Kami users again substantially outscore the national average, and Year 8, when Kami users score substantially below the national average.

Mean school engagement by gender (n=667)

Figure 25

Again, the trend for school engagement by gender compares directly to the trend for motivation by gender. Female students achieve on parity with the national average (3.33 vs. 3.31 respectively), whereas male students outperform the national average by some margin (3.67 vs. 3.29). This suggests that **male students using Kami were, on average, 12% more engaged at school than the national average**. There were no national averages available for students with other gender identities.

5. Conclusion

One of the current challenges for the education sector that our literature review focused on was deepening spending cuts and limited school budgets. As those financial constraints get tighter, it is more and more crucial for schools to weigh their procurement choices against the available evidence to ensure that every penny is put to work in the pursuit of improving student outcomes and life chances.

This report has built on the work in our literature review and Kami's logic model in order to make tentative first steps towards building a compelling picture of the impact of Kami on a range of important student and teacher outcomes. Although our evaluation design needed to strike a balance between practicality and rigour, we have uncovered some important first glimpses into the potential of a tool like Kami. One of these relates to teacher workload. It is encouraging to find that teachers using Kami spent 33% less time preparing lessons and to hear anecdotes about how Kami enables teachers to achieve a better work-life balance, increasing the time they can spend with their families, or decreasing the time they spend having to prepare for each lesson. Another is passion for teaching, where we saw that teachers using Kami outperformed those not using Kami on 14 out of 16 questions around school engagement and passion for teaching, and were 13% more passionate about teaching, on average, than non-Kami users. Although our findings on differentiated feedback and instruction were more mixed, our qualitative findings strongly suggested that teachers found the tools provided by Kami of great use in the classroom, and we would be excited to dig into these trends in more detail.

Our conclusions on student impact are more mixed. Although the above-midpoint scores for self-directed learning and collaborative learning may be encouraging, it is difficult to say more about Kami's effect here without valid comparison data. This is an area where we would strongly recommend additional data collection, as well as qualitative research to understand how students feel. On the other hand, however, students taught with Kami scored slightly more highly than benchmark for motivation for learning (4%), and a reasonably strong gender effect was observed here. Again, this outcome would benefit from more focused data collection as well as targeted qualitative research with pupils to understand how Kami impacts their learning.

The findings presented here have uncovered a number of interesting new avenues for exploration in future evaluation work, and we hope to be able to build upon the methodology and findings of this report in order to help Kami take the next step on their impact journey.

Appendices

A1. Glossary

Evaluation terminology

Academically validated measures

These are scales to measure social and emotional skills linked to academic achievement and long-term life outcomes that have been developed and peer reviewed by academic researchers within the fields of education and psychology. These have been developed to ensure:

- ▶ Predictive validity. These skills have been shown to be closely related to desirable life outcomes such as educational achievement, employability and earnings potential, or long-term health and life satisfaction. (In psychometrics, predictive validity is the extent to which a score on a scale or test predicts scores on some criterion measure. For example, the validity of a cognitive test for job performance is the correlation between test scores and, say, supervisor performance ratings.)

- ▶ Construct validity. The measure tests for the skill that it says it does, as defined in the literature.

- ▶ Test-retest validity. The results stay the same when tests are repeated.

Control Group

A control group is composed of individuals who did not participate in the programme but who closely resemble the individuals who take part in the programme in other ways. A control group is used to get an indication of whether a change in results over the course of the programme can likely be attributable to the programme itself, or whether results were likely to change over time in any case. Also known as a comparison group.

Evaluation

An evaluation is set up to measure the impact of a particular programme. This will involve monitoring the programme over a specified period, for one or more groups, in order to evaluate the progress participating pupils make. One programme can involve multiple evaluations, and we recommend gathering data across multiple time points to ensure valid and reliable results are generated.

Outcomes

We use outcomes to refer collectively to any social and emotional skills and academic attainment scores that are being measured over the course of an evaluation.

Participants

The group of individuals participating in the evaluation, and not forming part of a control group.

Programme

This could be any intervention, project or programme run in school with the aim of improving pupil outcomes or life chances. ImpactEd works with schools to build evaluations of their programmes in order to better understand whether they are having their intended impact.

Social and emotional skills

The term 'social and emotional skills' refers to a set of attitudes, behaviours, and strategies that are thought to underpin success in school and at work, such as motivation, perseverance, and self-control. They are usually contrasted with the 'hard skills' of cognitive ability in areas such as literacy and numeracy, which are measured by academic tests. There are various ways of referring to this set of skills, such as: non-cognitive skills, twentieth century skills and soft skills. Each term has pros and cons; we use social and emotional skills for consistency but we recognise that it does not perfectly encapsulate each of the skills that come under this umbrella.

Statistical analysis terminology

Statistically significant

A result has statistical significance when it is very unlikely to have occurred given the null hypothesis. In other words, if a result is statistically significant, it is unlikely to have occurred due purely to chance.

P-value

A p-value is a measure of the probability that an observed result could have occurred by chance alone. The lower the p-value, the greater the statistical significance of the observed difference. Typically a p-value of ≤ 0.05 indicates that the change was statistically significant. A p-value higher than 0.05 (> 0.05) is not statistically significant and indicates strong evidence for the null hypothesis; i.e. that we cannot be confident that this change did not occur due purely to chance.

A2. Kami Logic Model & Literature Review

kami **Logic Model** **ImpactEd** Evaluation

Problem	→	Inputs / Activites	→	Intermediate Outcomes	→	Long-Term Outcomes	→	Overall Aim

It is extremely challenging for teachers to provide **high quality, differentiated instruction** across different groups in the classroom.[1]

This means that pupils commonly receive the **same learning experience**[2], regardless of their level of academic attainment, their preferences regarding how they best attain, give evidence of their knowledge and receive feedback on their work, as well as any special educational needs they may have.

This challenge is exacerbated by several contextual factors that currently make it difficult for teachers to adopt appropriate solutions for differentiated instruction:

- 'Tech fatigue' makes teachers and learners reluctant to use 'yet another' tech solution.
- **Budget cuts**[3] put tech solutions out of financial reach for many schools.
- **Teacher workloads** make it hard for teachers to find the time to adequately tailor instruction and feedback.[4]
- **Lack of support and training** for embedding tech solutions.

How does Kami help address this?

Kami provides a **uniquely structured, collaborative digital environment** that:

- is **easy for teachers to understand** and incorporate into their teaching practice.
- Provides a way for teachers to **transform any kind of content** into an effective, consistent learning workflow.
- Provides **rich annotation features and collaborative learning tools**, which give students a range of ways to represent and improve their knowledge.[5]
- Allows teachers to quickly **provide richer and more targeted feedback** to learners, in a way that suits their learning style.[6]
- is both **device- and context-agnostic**, accessible to a range of users and school settings.[7]
- **Integrates fully** with a range of digital tools that teachers already use and know.
- Comes with **free, inclusive access to a dedicated training team** to help teachers use and embed Kami in their school.

For teachers

- Teachers have a greater understanding of how they can use Kami to **save time**.
- The amount of **time** that teachers spend **planning lessons decreases**.
- Teachers feel better equipped to provide **differentiated instruction and differentiated feedback** to their students.[8]

For students

- Students receive instruction that is **more personal** and more closely tailored to their needs.
- Students receive feedback that is **more timely and more meaningful**, allowing them to improve more quickly.
- Students' **excitement and motivation** around learning increases.[10]
- Students become **more deeply engaged** in learning (e.g. longer duration.)[11]
- Students' learning becomes **more self-directed and collaborative.**[12]

- Teachers' **passion** for teaching improves.
- Teachers' ability to provide **richer, more differentiated instruction** to their students improves.[9]

- Students receive a **more satisfying learning experience**.[13]
- Students develop more of a **growth mindset.**[14]
- Students are empowered to become **more independent**, taking **more ownership** of their participation and learning.[15]
- **Disparities** in formative indicators of learning between students at different levels of disadvantage **become smaller.**[16]

As a result of using Kami, providing differentiated instruction and effective feedback becomes easier and does not place undue burden on teachers.

This means that all students have access to media-rich, engaging, fulfilling learning experiences, irrespective of aptitude, physical classroom context, language context or other circumstances.

Kami - Rapid Literature Review

This document is designed to provide insight into the following areas to guide Kami and ImpactEd's Logic Model workshop. It is structured as follows:

- Part 1: Contemporary Global Challenges in Education
- Part 2: Evidence for the Usefulness of Kami's Features
  - Role of **annotation** in learning
  - Role of **formative assessment** in learning
  - Role of **collaboration** in learning
  - Role of **accessibility** in learning
  - Role of **personalisation** in learning
- Appendix: Outline of Methodology

Part 1: Contemporary Global Challenges in Education

In order to make the outputs of our evaluation as applicable as possible to the education sector, we want to ensure that Kami's evaluation design is informed by the current global challenges in education. To this end, we have conducted a rapid literature review identifying these key problems faced by schools worldwide. Ahead of the Logic Model workshop, Kami will be able to help us identify how features of Kami App begin to address these issues and ultimately lead towards more positive outcomes for a range of pupils and students across the sector. This mapping exercise will allow us to design an evaluation whose desired outcomes address issues that school staff members can relate to.

Budgets and funding

There is clear evidence there has been funding cuts and budget squeezes across education globally. According to the Institute for Fiscal Studies, school spending per pupil in England has fallen by 9% in real terms between 2009/10 and 2019/20, the largest cut in over 40 years. Education budgets globally have been further squeezed as a result of Covid-19 and rising energy costs. Whilst in the US, Congressional Republicans are pushing for a legislation that would result in a 22% cut in schools with low-income students and students with disabilities, which could force a workforce reduction of up to 108,000 teachers, aides or other key staff.

Unsurprisingly, cuts in education funding have severe consequences for pupil outcomes. Tighter budgets means being able to afford fewer teachers which have a direct impact on pupils' learning. Fewer teachers result in larger classroom sizes and, in turn, a decrease in personalisation of learning. Tighter budgets also lead to a reduction in the number of extra-

curricular activities, shortened school hours, a narrowed school curriculum and a reduction in SEND support.

Socio-economic differences

Another key global issue that schools are facing is that a pupil's socio-economic status continues to play an influential role in their success at school (EPI). The general trend is that pupils from wealthy backgrounds achieve better results than their less affluent peers (Ballard Brief). Some recent research suggests that the level of inequality has not been addressed, continuing unchanged (EPI, Ballard Brief), and some research suggests that educational inequality has actually widened (TES). In the UK, some research indicates that although some progress had been made in narrowing this gap, Covid-19 has reversed this progress and that it will take another 10 years to return to pre-pandemic levels of inequality (House of Commons Committee of Public Accounts).

Social and emotional learning

The final key problem is the worrying trend of pupils experiencing poor mental and emotional health (KFF). The literature highlights that although poor mental and emotional health may have been exacerbated by Covid-19 (KFF, U.S. Department of Education, NCES), research by WHO and UNICEF indicates that poor mental wellbeing among students was, in fact, already a severe problem well before the pandemic (European Commission). In the UK, survey results from the National Health Service indicated that there had been an increase from 11.6% to 17.8% of 5-16 year olds being identified as having probable mental health problems (NHS). There is some hope, however, following the UK government's report of the state of schools in 2022 claiming that mental health levels are back to pre-pandemic levels but that anxiety levels are higher than the previous year (UK Government).

Part 2: Evidence for the Usefulness of Kami's Features

We have conducted a rapid literature review on the role of the following features of Kami: annotation, formative assessment, collaboration, accessibility, and personalised learning. Although we discuss each feature separately, it is important to note that there is interplay across various features. For example: annotation can facilitate collaboration, formative assessment can facilitate collaboration and personalised learning, and personalised learning can facilitate accessibility. In the table below, we present some emerging findings from the rapid literature review.

Functionality	Emerging Themes

Annotation
 	
Link to issues:
 	
Annotation sharing (and knowledge sharing) could be beneficial in addressing the inequality in academic achievement between wealthier pupils and less affluent peers.
	▶ Annotation had a positive impact on a wide range of pupil outcomes such as learning achievement (Hwang et al., 2015; Su et al., 2015; Shadiev et al., 2015), approach to reading (Porter-O'Donnell, 2004; Zywica et al., 2008) and learning vocabulary (Xu, 2010; Akbulut, 2007), a range of meta-cognition outcomes (Porter-O'Donnell, 2004; Zywica et al., 2008; Akbulut, 2007), and collaboration among peers (Hwang et al., 2015).
	
	▶ However, annotation does not always increase achievement in reading comprehension (Akbulut, 2007).
	
	▶ The quantity of pupil annotations (Su et al., 2015; Hwang et al., 2015), as well as the frequency of pupils reviewing their own annotations (Su et al., 2015), were both good predictors for learning achievement.
	
	    ○ Lai et al. claims that multimedia environments that provide annotations are most effective because perhaps students can select annotations that fit their needs (Lai et al., 2011).
	
	▶ Annotations were also found to facilitate collaboration.
	
	    ○ Pupils found annotating other pupils' homework solutions beneficial (Su et al., 2015).
	
	    ○ Lower-achieving learners reinforced their learning by reading higher-achieving learners' annotations (Hwang et al., 2015).
	
	▶ Another apparent theme was the positive implication of annotations on pupils' reading:
	
	    ○ Annotation made pupils realise that reading is a process, forces them to slow down, prevents skimming, and makes them more active in their reading (Porter-O'Donnell, 2004; Zywica et al., 2008).
	
	▶ Annotation seemed particularly beneficial to learning new vocabulary, making pupils reflect more on what they'd read (Xu, 2010; Akbulut, 2007), as well as helping pupils visualise in their learning (Zywica et al., 2008) .
	
	▶ Finally, one key theme of the impact of annotations on pupils is that it helped pupils in their independence (Zywica et al., 2008) and made them feel more

positive about their learning (Akbulut, 2007). It did not, however, have an impact on their cognitive load (Lin et al., 2016).

Formative assessment

Link to issues:

Formative assessment could have a positive impact both on improving outcomes for less affluent pupils as well as improving pupil's social and emotional learning.

▶ Formative assessment is a powerful way to influence and improve student learning and achievement (Mirriahi et al., 2016; Shih et al., 2012; Cauley et al., 2010; Baron, 2016; Clark, 2012; Spector et al., 2016), as well as positively impact a range of non-cognitive outcomes (Mirriahi et al., 2016; Nolen, 2011, Vassilakis, 2009/2010, Cauley et al., 2010; Clark, 2012).

▶ Formative assessments facilitate students sharing their knowledge with each other (Vassilakis, 2009/2010).

▶ Formative assessment also facilitates personalised learning (Shih et al., 2012; Spector et al., 2016) as the teacher can use outputs from formative assessments to improve their instruction (Cauley et al., 2010; Volante et al., 2011).

▶ There is a large body of evidence on the positive impact that formative assessment has on pupil motivation (Mirriahi et al., 2016; Nolen, 2011; Cauley et al., 2010; Vassilakis, 2009/2010; Clark, 2012). It also seems to have some positive benefits on pupils' persistence (Nolen, 2011) and engagement (Vassilakis, 2009/2010).

Collaboration in learning

Link to issues:

Collaboration could also have a positive impact on reducing inequality between disadvantaged pupils and their peers.

▶ There is some indication that cross-peer collaboration can improve student learning (Hwang et al., 2015; Smith, 2019).

    o As mentioned above, sharing annotations between pupils, and annotating other pupils' homework is seen as beneficial (Hwang et al., 2015).

▶ The impact of collaboration on pupil learning seems to be highly context dependent and is likely to vary from student to student (Kuo et al., 2015; Chen et al., 2015).

▶ Collaboration seemed to increase pupils' motivation (Smith, 2019).

▶ Collaboration seems to be effective because students with higher-quality work were proud to guide other students, and students with lower-quality work were proud to contribute to higher-quality work (Vassilakis, 2009/2010).

▶ Staff seemed to struggle to incorporate tech-based collaboration in their teaching (Wardlow et al., 2015).

Accessibility in learning

Link to issues:

With restricted budgets having a negative impact on support available to pupils with SEND needs,

▶ The literature indicates that technology has made education more accessible in some ways (Traphagan et al., 2010) using technology has also created new blockers around accessibility (Fichten et al., 2019; Hollins et al., 2013; Chu, 2014) and that more work needs to be done to address these newly formed gaps (Kumar et al., 2016) .

▶ Some key issues around technology and accessibility were:

    o The benefits of using mobile devices be outweighed by their capacity to distract some users (Hollins et al., 2013) .

it's clear that successfully accessible EdTech products could help plug this gap.	o The problems posed by the widespread use of PDFs, since these files are often not readable by screen-reading software (Fichten et al., 2019).
	o The problems posed by using videoclips without captioning (Fichten et al., 2019).
	o Learning Management Systems often are not capable of increasing the time allowed to complete a specific task for users who are usually permitted to have extra time in educational activities (Fichten et al., 2019).
	o The combination of digital (e.g., device-based) and non-digital (e.g., real-world based) learning can create a high level of cognitive load for students (Chu, 2014).
	▶ One solution for increasing readability of websites is by utilising highly contrasting colours (Bray et al., 2007).
Personalised learning *Link to issues:* *The level of personalisation granted by an EdTech product may resolve the implications of a teacher having limited time to spend on each pupil, an issue derived from budget cuts.*	▶ In the literature, that there is still a need to define personalised learning before investigating its impact (Hartley, 2007; Lee et al., 2018). Some authors view the concept as incoherent (Hartley, 2007). Others state that personalised learning can be characterised as follows: learning analytics recorded while students interact with a learning system, and these are used in a predictive algorithm to suggest next learning task based on individual students' data (Bienkowski et al., 2012).
	▶ Personalised learning appears to produce learning gains for pupils (Chen et al., 2021). The literature shows that in specific subjects, such as mathematics, using a digital personalised learning is better than a more traditional paper and pen approach (Shih et al., 2012).
	▶ Personalised learning is useful in giving teachers insight and overview of their pupil cohort. It can help teachers identify pupils at risk (Sharples, 2019) as well as identify the contextual factors involves (Herodotou, et al., 2019).
	▶ Personalised learning can help to create accessible and robust learning profiles (Spector et al., 2016) as well as meeting the needs of everyone (Rienties et al., 2019) because it gives students the choice as to how they can most effectively learn (Lai et al., 2011) .
	▶ There's some evidence around how personalised learning can improve pupils' engagement, self-efficacy, and confidence (Chen et al., 2021).
	▶ There are suggestions that for personalised learning to be effective it needs to be designed based on learning theories (Zhou et al., 2018) and that it needs to integrate recordkeeping, planning, instruction, and assessment (Lee et al., 2018).
	▶ There are some broader concerns that personalised learning is another off shoot of the marketisation of education (Hartley, 2007; Hartley, 2008) and that this type of learning will privilege the middle classes (Hartley, 2008).

45

Appendix – Outline of Methodology

Part 1: Contemporary Global Challenges in Education

1. Key phrases were searched in Google, both from UK and US region, to identify key issues being faced by schools.
   a. (Those scanned articles are collected here: ICG, My Tutor, The Ed Advocate, EduStaff, Education Inspection, The Guardian, Cambridge University, Aim A Litte Higher, London Councils, and Our World in Data.)
2. Once key themes were identified, key phrases focusing on key themes identified in previous search were search on Google (Budget and funding, Socio-economic differences, and social and emotional learning.)
3. A range of sources was then collected for each key theme and conducted a thematic analysis drawing out what kind of impact these challenges were having on pupils.

Part 2: Evidence for the Usefulness of Kami's Features

1. For each key functionality that Kami wanted to investigate, using key phrases (e.g., impact of annotation on pupils' learning), JSTOR's database was searched for relevant articles.
2. Key findings were noted for each relevant paper for each functionality.
3. A thematic analysis was conducted on key findings for each functionality and written up.

Partners and supporters

United Learning
The best in everyone™

THE YOUNG ACADEMY

CHALLENGE **PARTNERS**

SUPER BEING

nesta

TeachFirst INNOVATION PARTNER 2018/19

LEEDS BECKETT UNIVERSITY

MANCHESTERTSA
Teaching School Alliance

Improving pupil outcomes by working with schools to address the evaluation deficit.

Get in touch
hello@impacted.org.uk

.ıIImpactEd