The Invisible Burden:
Teacher Beliefs & EdTech Adoption in Singapore Primary Schools

Submitted by CPT31 / PFG08 / 2026 MLS Sem1

- Lua Shixiong
- Michael Sim Moh Ngiat
- Ngiam Song Wee
- Ong Shi Han
- Wong See Wan

Abstract

We examine how Singapore primary school teachers perceive educational technology as a burden vis-à-vis benefit, and which dispositional and structural factors shape adoption. Using chatbot-mediated interviews with 61 teachers, responses were coded across multiple dimensions including facilitating conditions, volition, pedagogical orientation, and perceived student impact.

In general, teachers' attitude towards EdTech is frequently ambivalent rather than binary. Findings reveal a system stable at functional adequacy, with blended efficacy as the strongest predictor of student impact. Also, teacher agency emerges as the bridge between supporting factors and classroom practice. Notably, years of service does not predict EdTech adoption. Most teachers simultaneously believe in EdTech's potential alongside genuine reservations about workload, impact on students, and pedagogical fit; most teachers carry this emotional complexity as an unremarked feature of their professional lives.

The study then identifies barriers to meaningful EdTech integration and provides ICT Key Personnel with implications pertinent to their individual schools.

Keywords:

EdTech adoption, teacher beliefs, primary education

Contents

1 Introduction

Educational technology ("EdTech") now occupies a significant mindshare in the Ministry of Education's ("MOE's") vision for teaching and learning. Yet the challenge facing schools is no longer whether technology is available, but it is how EdTech is taken up in everyday classroom practice. This study examines how teachers in selected Singapore primary schools perceive the time, effort, and demands involved in EdTech adoption, and how these perceptions may shape the extent and quality of classroom use.

Primary schools provide a particularly revealing setting for this inquiry. Unlike secondary schools, the vast majority of primary schools in Singapore do not operate a 1:1 Personal Learning Device ("PLD") programme. EdTech in classrooms therefore remains largely mediated by the teacher, who bears the brunt of deciding when EdTech is appropriate, adapting it for younger learners, integrating it as part of lesson flow, troubleshooting issues, all while managing EdTech use in real time. For School Leaders and Information and Communication Technology ("ICT") Key Personnels ("KPs"), EdTech adoption is therefore not simply a matter of access or encouragement. It is a question of how system-level support translates into classroom practice.

This study adopts a leadership lens. It seeks to understand the factors shaping uneven uptake, including how teachers experience EdTech as a burden, benefit, or both. By surfacing these perspectives, the study aims to provide ICT KPs and School Leaders with a clearer basis for responding to the realities of EdTech adoption in their own classrooms, in their own schools.

1.1 The Policy-Practice Gap

Singapore has invested significantly in EdTech infrastructure and implementation. In September 2023, MOE announced the EdTech Masterplan 2030 which set out three priority areas:

1. Strengthening students' digital literacy
2. Leveraging emerging technologies such as Artificial Intelligence ("AI") to better customise learning, and
3. Building a culture of sharing technology-enabled practices across schools.

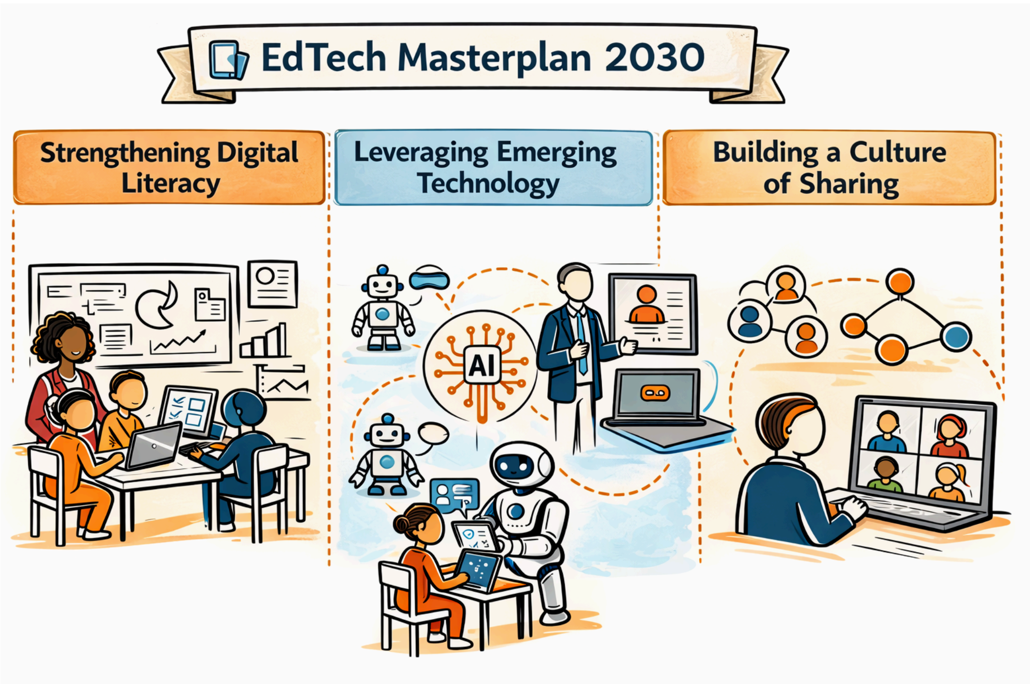

In July 2024, MOE revised the primary school ICT provision norm, from 1 computing device per 5.5 students, to 1 device per 3.5 student. The change permitted improved device access for primary school students. Most recently circa 2026, ICT

Enablement Planning was introduced to provide School Leaders with greater clarity on ICT readiness and priorities, while aligning schools with system-wide digital initiatives that served to better manage workload.

These efforts reflect strong policy intent. However, they also sharpened the concern that motivated this project: Despite clearer direction, improved access and stronger enablement, we as primary ICT KPs perceive the degree of meaningful classroom EdTech adoption to be uneven. We suspect a gap between system-level support and teachers' lived classroom realities. This provided the impetus for this project.

1.2 Assumptions About Uneven Adoption

First, we assumed that uneven EdTech adoption could not be solely explained by teacher resistance alone. We suspected that some teachers who appear reluctant may feel burdened by the effort required to learn new tools, and implement them proficiently in classroom settings.

Second, we assumed that not all use of technology passes off as meaningful EdTech adoption and were particularly guarded against one-off use by teachers seeking purely student engagement. We sought to find out if technology was being used in a purposeful and sustained way to support teaching and learning.

Third, we assumed that teachers carry out forms of hidden work that were not always recognised either in formal implementation or as part of performance appraisal. Even with greater availability of devices, stronger collaborative practices and structured ICT planning, teachers may still have to invest significant effort in lesson planning, experimentation, real-time troubleshooting and refining classroom routines. We suspect that this invisible workload may explain why EdTech adoption remains uneven despite substantial government investment.

The above assumptions situate our study around a practical leadership problem. If meaningful EdTech adoption remains uneven despite greater support and emphasis, then the issue may not lie only in availability of technology but also in how teachers experience the work of adoption in everyday practice. For ICT KPs and School Leaders, understanding the ethnographic realities may better inform us on how to create conditions to shift EdTech adoption stances from compliance towards purposeful, sustainable and meaningful use.

2 Literature Review

2.1 Definitions

Two constructs carry the most conceptual weight and require careful definition: teacher beliefs and invisible workload.

2.1.1 Teacher Beliefs

Teacher beliefs in the context of this project, refer to the deeply held convictions that teachers carry about the nature of teaching, the conditions under which learning occurs, and the role of technology within both. These are not surface-level attitudes or expressed preferences. Instead, they are durable, dispositionally-embedded orientations that function as interpretive filters for all new instructional information (Ertmer, 2005).

Teacher beliefs are not readily revised through training or instruction alone. People tend to preserve conceptual coherence in their personal theories over time, attending selectively to information that confirms existing assumptions while marginalising disconfirming evidence (Ertmer, 2005).

This matters for EdTech adoption because a teacher who believes that good teaching is fundamentally about direct instruction will interpret even the most well-designed digital tool through that lens.

2.1.2 Invisible Workload

Invisible workload is the conceptual centrepiece of this project. It refers to the unacknowledged, uncompensated labour that teachers absorb when integrating EdTech tools. This includes time spent on learning the tool, redesigning lessons, troubleshooting failures, managing unfamiliar digital workflows, and recovering cognitively and emotionally from integration that does not go as planned.

This burden is invisible not because it is small, but because it falls outside formal systems that recognise, measure, and compensate teacher effort. Workload-measurement instruments, time audits, and performance appraisal frameworks are not designed to capture lesson redesign at midnight, the cognitive cost of a failed pilot, or the emotional labour of navigating institutional ambiguity in real time.

Invisible workload in EdTech integration is not simply a matter of extra hours. When viewed through a labour theory lens, digital technologies have reconstituted education work in ways that produce both the standardisation and intensification of teachers' labour. Teachers engage in diverse forms of what the authors describe as immaterial labour: the production of information, knowledge, communication, and affect (Selwyn et al., 2016). As a result, teachers' work has come to resemble a labour process which is described as limitless, shapeless work that is never over and never done (Connell, 2009).

The emotional dimension of this workload deserves particular attention. In the context of EdTech adoption, this maps onto cognitive and emotional recovery costs: the residual anxiety of a lesson that failed because a platform crashed, the mental rehearsal of what to try differently next time, the low-grade stress of being perpetually on the wrong side of a learning curve.

2.2 Core Constructs

The theoretical architecture of this project draws on two interconnected frameworks: Ertmer's two-level barrier model and Bandura's self-efficacy theory. This section develops each framework in turn before tracing the relationship between them.

2.2.1 Ertmer's Two-Level Barrier Framework

Ertmer distinguishes between two qualitatively different types of barriers that obstruct technology use in classrooms (Ertmer, 1999, 2005).

First-order barriers are external, resource-based constraints that are largely institutional or infrastructural in nature. These include limited hardware access, insufficient technical support, inadequate time for preparation, and the absence of relevant digital resources. While significant, first-order barriers are in principle addressable through targeted investment and administrative action. While Singapore's sustained commitment to ICT infrastructure has substantially reduced first-order barriers across its school system, infrastructure alone is insufficient in driving meaningful integration (Natarajan et al., 2021).

Second-order barriers are intrinsic to the teacher and are far more resistant to change. These arise from teachers' deeply held beliefs about teaching, learning, and the role of technology within both. Literature suggests that these beliefs function as a filter through which all new instructional information is processed. Teachers selectively attend to evidence that confirms existing assumptions while setting aside disconfirming data.

Ertmer argues that it is these second-order barriers, not access or training, most powerfully predict whether teachers will adopt technology in pedagogically meaningful ways. In so doing, Ertmer postulates the primacy of second-order barriers as the explanatory variable for persistent gaps between EdTech investment and pedagogical transformation.

The primacy of second-order barriers corroborates with observations of how selected Singapore teachers navigated both big "R" reforms (system-level ICT Masterplans) and small "r" reforms (school-level innovations); they revealed that even in schools with strong infrastructure and leadership support, teachers' willingness to exercise agency in technology integration depended heavily on their beliefs about their own capacity to lead and to shape practice. In particular, there exists a "I am just a teacher" syndrome which captures a surrender of agency that maps directly onto Ertmer's second-order barriers: teachers who do not believe they have the power to effect change are unlikely to invest the effort required for meaningful integration, regardless of what first-order conditions allow (Hong, 2023).

2.2.2 Bandura's Self-Efficacy Theory

Intersecting with Ertmer's barrier framework is Bandura's theory of self-efficacy, which is defined as an individual's judgment of their own capability to organise and execute the courses of action required to achieve specific goals, particularly in diverse, unpredictable, and stressful situations (Bandura, 1997). It is critical that self-efficacy concerns not the skills a person actually possesses but their beliefs about what they can accomplish with those skills.

Literature suggests that teachers with higher computer self-efficacy use technology more often, experience less technology-related anxiety, and are more likely to view themselves as capable of handling new systems (Sun & Yan, 2025).

Conversely, teachers with low self-efficacy are more susceptible to frustration, more prone to avoidance, and more likely to abandon integration attempts when they encounter obstacles. Therefore, when self-efficacy is low, the unacknowledged cognitive and emotional labour (i.e. invisible workload) of technology troubleshooting feels disproportionately heavy.

Self-efficacy is not fixed but rather can be built or eroded (Bandura, 1997) through

- mastery experiences (succeeding at a task),
- vicarious experiences (observing a peer succeed),
- social persuasion (being told by someone credible that one can do it), and
- physiological states (managing the anxiety and fatigue that integration generates).

Each of the above modalities has direct implications for how schools design professional development and support structures around EdTech.

2.2.3 Synthesis: How Beliefs, Workload, and Self-Efficacy Interact

Our hypothesized relationship operates as follows: second-order belief barriers shape technology self-efficacy, which determines the extent to which invisible workload is experienced as prohibitive. When beliefs are misaligned and self-efficacy is low, the invisible labour of integration becomes subjectively heavy, reinforcing disengagement and performative compliance. Where beliefs are aligned and self-efficacy is high, the same workload is experienced as manageable, even meaningful, investment.

This is not a static relationship. Invisible workload, when unacknowledged by school leadership, functions as a steady accumulation of unsupported failure experiences. That is the condition under which Bandura's theory predicts erosion which chips away at the mastery-experience foundation on which self-efficacy is built.

Over time, this erosion can deepen second-order barriers. Teachers whose confidence has been quietly drained by invisible workload are more likely to retreat to safe, transmissive practice. They come to interpret any new EdTech mandate as further evidence that the system does not understand their reality. Hong's documentation of the "baby elephant syndrome" in Singapore schools illustrates the end-state of this cycle: through accumulated experience, teachers learnt that it is futile to resist the constraints they face, and consequently surrender their agency altogether.

We needed to surface the beliefs and self-efficacy perceptions that drive these dynamics. Therefore, we approached data collection with the aim of eliciting the affective and experiential dimensions of teacher knowledge. This necessity influenced our choice of ethnographic methodology described in the following chapter.

3 Impetus for the Project

3.1 The Amplifying Effect of Invisible Workload

Literature suggests that second-order barriers and low self-efficacy make EdTech adoption subjectively harder. Invisible workload is the mechanism through which that difficulty compounds.

The increased digitisation of schools has received surprisingly little critical attention, despite growing evidence that digital technologies can make education work more complex rather than less (Selwyn et al., 2016). Teachers' labour is already characterised by competing demands: the production of learning, affect, and bureaucratic artefacts simultaneously (Connell, 2009). Technology integration adds a further layer of invisible demands to this already intensive process.

This invisible workload is not pedagogically neutral. When teachers with low self-efficacy or misaligned beliefs encounter the real costs of integration, those costs are subjectively experienced as evidence that EdTech is not worth the effort. A teacher who spends two hours preparing a Student Learning Space ("SLS") lesson only to have half the class go off-task does not simply experience a failed lesson. That teacher accumulates a data point against future investment.

This may create a self-reinforcing cycle. High invisible workload reduces self-efficacy. Reduced self-efficacy increases perceived workload burden. Both together reinforce existing belief barriers to adoption. The cycle is difficult to interrupt precisely because the workload that drives it is invisible to the systems designed to support teachers: it does not appear in timetables, workload agreements, or performance conversations.

Understanding and interrupting this cycle is the practical motivation for this project.

3.2 Why Primary Schools

Primary schools occupy a foundational but underexamined position in Singapore's EdTech landscape. Unlike secondary school students, the vast majority of primary school students do not have PLDs. The teacher is therefore the primary, and often sole, mediator of all technology experiences in the classroom. This places a disproportionate burden of design, curation, and adaptation on primary school teachers. It also means that the quality of teacher belief and workload management is especially consequential: if the individual teacher chooses not to initiate EdTech, it does not happen.

Additionally, primary school teachers operate within a context of institutional mixed messaging. MOE’s Educational Technology Division (“ETD”) actively promotes AI-assisted learning and adaptive platforms. Yet teachers simultaneously navigate a broader social climate that is increasingly cautious about children's digital exposure. Public health guidance from the Ministry of Health (“MOH”) recommends limiting screen time for younger learners, and this caution is reinforced by growing public discourse around the effects of excessive screen use on children's attention, sleep, and social development.

These contradictory pressures create what amounts to a compliance bind. Teachers are expected to integrate EdTech meaningfully, yet they also absorb ambient signals that technology exposure for young children should be carefully managed. Without clear guidance on how to reconcile these expectations, teachers must navigate the tension largely on their own. This discourages authentic risk-taking and favours safe, minimal-compliance approaches to technology use.

4 Research Methodology

We sought to understand how teachers experience the time, effort, opportunities, and challenges associated with integrating EdTech into classroom practice. We aimed to surface human factors such as beliefs, confidence, motivation, pedagogical orientation, and school context, which play an important role in shaping adoption,

4.1 Qualitative and Ethnographic

We employed an interpretivist, qualitative exploratory methodology, supported by score-coding of the resulting data.

Teachers' epistemic beliefs, emotional experience, and the invisible workload are not captured well by surveys or structured instruments. They require an approach that preserves the interpretive and affective dimensions that quantitative instruments tend to flatten.

We therefore drew on ethnographic principles. Ethnography aims to capture what people genuinely believe, not what they perform for an audience. Asked directly whether they support EdTech adoption, most teachers will give the expected answer. What we sought lay underneath: the reservations, the workarounds, the quiet compromises that shape actual classroom practice.

Ethnographic work has a known challenge: the researcher can never be fully absent. Body language, tone, and the social dynamic between interviewer and teacher all introduce noise. Experienced ethnographers manage this through reflexivity, typically by keeping a research log. Even so, managing bias across dozens of interviews is demanding and imperfect.

We do not claim to have conducted ethnography in its full sense. True ethnography requires sustained immersion in a setting, and our study did not do that. What we did was apply ethnographic sensibilities to the design of structured, chatbot-mediated interviews: prioritising lived experience over opinion, probing for classroom reality over abstraction, and creating conditions under which teachers could speak candidly. The methodological choice that made this possible is described in the next section.

4.2 Chatbot-Mediated Ethnographic Interviews

We deployed a structured conversational chatbot on the EdCafe platform as our primary data collection instrument. Teachers responded under self-chosen pseudonyms, either by typing or by speaking via voice capture. The voice option also accommodated participants more conversant in their Mother Tongues. The chatbot recorded transcription and, where needed, translation.

We chose this design for three reasons.

1. **Consistency**

   Every teacher encountered an identical interviewer. There was no variation in tone, no visible reaction to a controversial answer, and no unconscious probing that pursued one teacher's thread more than another's. This consistency is a prerequisite for defensible comparison across 61 respondents.

2. **Candour**

   The anonymity afforded by pseudonyms reduced concerns about how a supervisor might view any given response. The asynchronous format also gave teachers time to reflect, and allowed them to pause and return to an interrupted session. We judged that these conditions would produce more candid accounts than a scheduled face-to-face interview with a known researcher.

3. **Scale**

   We completed 61 interviews within a week. A conventional qualitative study of this size would require months of scheduling and transcription. The chatbot compressed that cycle, allowing us to direct effort towards analysis rather than logistics.

There are limitations we must acknowledge:

1. **Emotional Attunement**

   A chatbot cannot read the room, sense hesitation, notice what is being avoided, or gently redirect a teacher who has misunderstood a question. A skilled human interviewer brings the emotional attunement needed to draw out a reluctant or evasive participant.

2. **Sustaining Duration of Interview**

   Some sessions in our study were notably short, and the chatbot's follow-up prompts were not always sufficient. The chatbot also lacks the relational warmth that can open a participant who is cautious about disclosure.

We accepted this trade-off deliberately. The chatbot bought scale and consistency at the cost of depth and flexibility. For a study aimed at surfacing broad dispositional patterns across 61 teachers, we deemed the trade-off acceptable.

The fact that teachers were willing to share candidly with a machine, under a pseudonym, at a time of their own choosing, may itself say something about how rarely these experiences are asked about at all.

4.2.1 Chatbot Parameters

The interview chatbot was prompt-engineered, and was not a general-purpose chatbot given a loose brief. Instead, it was given a precise role, a fixed question sequence across five modules, and strict behavioural rules:

- one question at a time
- a maximum of one probe per question
- no invented follow-ups,
- paraphrasing only at module boundaries, and
- a permanent state rule preventing any return to earlier modules once closed.

The rules regulating “emotion” were also explicit:

- acknowledge in one sentence, and
- return to the question flow.

These constraints served two purposes. They ensured consistency across all 61 interviews, which is a prerequisite for defensible comparison during analysis.

They also prevented the chatbot from drifting into invented lines of inquiry that would introduce noise the research team could not trace or justify.

The five modules covered

- Profile,
- Current ICT practice,
- Workload experience,
- Beliefs and meaning, and
- System and support.

The full question set and probe logic are documented in Appendix A.

4.2.2 Strengths and Limitations of the Interview Chatbot

The chatbot completed 61 interviews across five schools in 8 days spanning a weekend (27 March to 3 April 2026). This pace would not have been possible with 5 human interviewers.

The asynchronous format allowed teachers to respond at their convenience. Several respondents completed sessions across multiple sittings. The text-based default might have lowered the threshold for candour: teachers could reveal sentiments that they might not say aloud, particularly about workload frustration and institutional pressure.

The transcripts yielded concrete, experience-grounded data. Teachers described specific lessons, named specific tools, and shared lesson-preparation steps. The chatbot's consistent question structure ensured comparable coverage across all 61 respondents, reducing interviewer variability.

Three limitations surfaced during analysis.

Question Repetition

The chatbot occasionally re-asked questions the respondent had already answered. Respondent "Beloved" replied "already answered this earlier" 7 consecutive times before the chatbot moved on. Respondent "mh" ended the session after noting: "Seeing the repeat of questions." This caused visible frustration and, in at least 2 cases, premature exit.

Participant fatigue

Several respondents signalled exhaustion toward the end of their sessions. Both respondents "Breadtalk" an “UserA” asked, "When will the questions end?".

The chatbot's fixed question sequence did not adapt to response depth. A teacher who gave a rich, multi-paragraph answer received the same number of follow-up probes as one who answered in two words.

Conversational Rigidity.

The chatbot sometimes failed to follow the respondent's lead. Respondent "Jes" raised device logistics as a primary concern. The chatbot redirected toward teaching and learning. Jes responded: "You have deviated from my problem." The chatbot did not recover. This was the clearest case yet to illustrate where a human interviewer would have produced richer data by staying with the respondent's framing.

Dataset Remains Valid

These limitations did not invalidate the dataset. They constrained its depth in specific cases. The richest transcripts came from respondents who engaged at length despite the chatbot's structural constraints. The thinnest transcripts came from respondents who disengaged early, often because the chatbot's repetition signalled that their input was not being heard.

For future iterations, adaptive sequencing (skipping questions already answered), response-length sensitivity (fewer probes for detailed answers, more for sparse ones), and stronger topic-following would address the three main weaknesses without sacrificing the format's strengths in scale and consistency.

4.3 Responsible Use of AI for Qualitative Data Analysis in a Verifiable Manner

To safeguard against AI hallucination and confident assertion in the absence of evidence, we designed the analysis process to be human-auditable at every step.

Our approach is consistent with Wang's (2025) characterisation of LLMs in qualitative research: that they transform qualitative analysis from a solitary craft into a dialogical coding process, and that the core competence required is a "literacy for human-AI collaboration" rather than a delegation of judgment to the machine. The passes described below operationalise that principle. The AI proposed, we interrogated, and final judgments rested with the research team.

4.3.1 Use of AI for Data Analysis

Coding the 61 transcripts was carried out against 12 spectra:

S/N	Spectrum	Range
S1	Years of Service	Beginning Teacher to experienced teacher
S2	Work Burden	Work creator to work saver
S3	Facilitating Conditions	Structurally failing to fully enabling
S4	Volition / Agency	Compliance-driven to intrinsically-motivated
S5	Resilience	Deterred by failure to adaptive for continuation
S6	Psychological Safety	Isolated and unsafe to safe to fail
S7	Troubleshooting Confidence	Gives up easily to expert problem-solver
S8	Collegiality	Isolated to community organizer
S9	Pedagogy	Instructivist to constructivist
S10	Blended Efficacy	No integration to systemic mutual influence of offline and online learning
S11	Perceived Student Impact	No benefit to transformative
S12	Perceived Student Digital Readiness	Major barrier to lesson flow to fully self-managing

The exact descriptors for the 6 ranks in each spectra can be found in Appendix.

Each spectrum was scored on a six-point rank with detailed descriptors for each rank. AI was used to propose draft scores, but its outputs were never accepted at face value. The process involved two human-in-the-loop passes.

In the first pass, an AI language model was provided with the 12 spectra and the six rank descriptors for each. It proposed draft scores for every respondent, together with evidence quotes drawn from the transcript. We reviewed each score against the transcript. Whenever we updated a score, the AI was instructed to characterise the change and to check it against all existing scores, surfacing any inconsistencies for human confirmation.

A second calibration pass focused on spectra where borderline cases clustered, particularly S4, S5, S6, and S8, where distinctions between adjacent scores required closer judgment. Across both passes, 53 scoring changes were made through human intervention. Each change is traceable to a specific respondent and spectrum.

To make the analysis reproducible and accessible, the full dataset was built into an interactive dashboard. The dashboard displays every respondent's coded score on each spectrum, the tone of voice registered across the interview, and the evidence quotes that support each score. Any reader can open a teacher's card, read the transcript, and check whether the scores assigned match what the teacher said. The correlation analysis across all 45 spectrum pairs is also visible, colour-coded by effect size, and exportable. Nothing is hidden behind a summary.

The result is an AI-assisted study where the AI's role is specific, bounded, and checkable. Humans produced the instrument, refined the scoring, and produced the dashboard. This is not a study that left AI to its own device. It is a study that used AI in ways a sceptical reader can interrogate.

5 Findings and Discussion

5.1 Profile of Respondents

61 primary school teachers completed the chatbot interview between 27 March and 3 April 2026. Participation was voluntary. Teachers chose their own pseudonyms and could use either text or voice to respond.

The sample covered a broad spread of experiences. Using the six-band years-of-service coding (S1), the distribution was as follows:

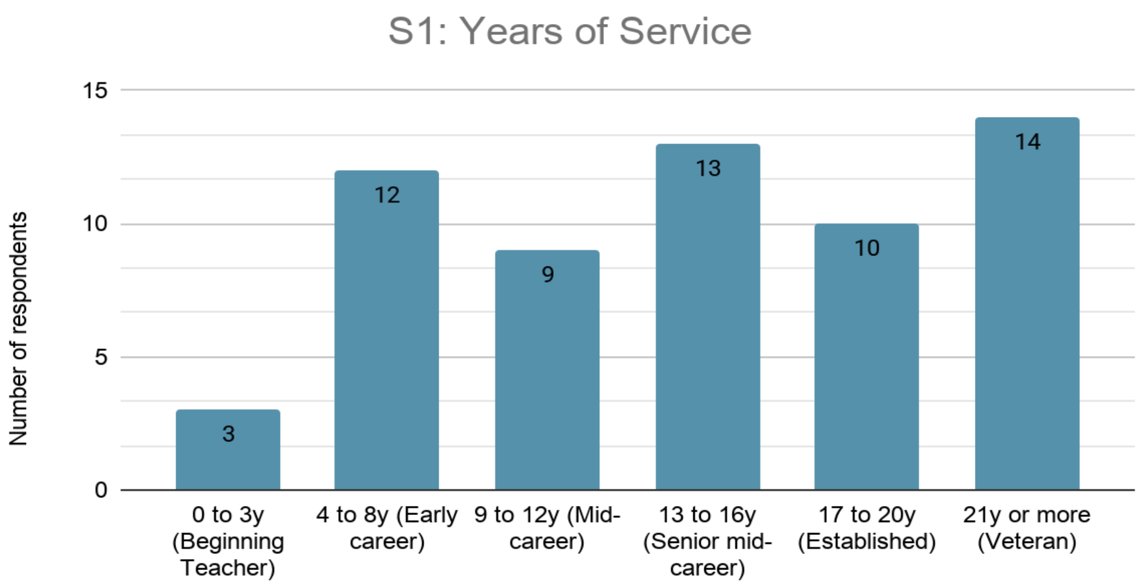

No single career stage dominated the sample. Veterans (21+ years) were the largest single group at 23%, but early career, mid-career, and late career teachers were all represented in meaningful numbers. The number of Beginning Teachers is comparatively small but not under-represented; the distribution reflects the low pace of hiring mainly to make up for systemic attrition and retirement.

This spread matters for the findings that follow: any pattern that holds across the sample is unlikely to be an artefact of one career stage alone.

Teachers came from a range of subject areas, including English Language, Mathematics, Mother Tongue languages, Science, Social Studies, Art, Music, and Physical Education. Some held formal appointments (Senior Teacher, Subject Head, Head of Department) while others were classroom teachers without formal leadership titles. Four pseudonyms appeared twice in the dataset by coincidence; these respondents were disambiguated by subject and years of service.

5.2 Descriptive Statistics Across the 12 Spectra

Each of the 61 transcripts was scored across 12 spectra on a 1-to-6 rank scale. Fig. 5.2 summarises the central tendency and spread for each spectrum. The rank descriptors for each band can be found in Appendix A.

S/N	Spectrum	Mean	Mode	Median	SD
S2	Work Burden	3.10	3	4.0	0.74
S3	Facilitating Conditions	2.97	3	3.0	0.81
S4	Volition / Agency	3.49	3	3.0	0.90
S5	Resilience	3.20	3	3.0	0.79
S6	Psychological Safety	3.20	3	3.0	0.79
S7	Troubleshooting Confidence	3.07	3	3.0	0.81
S8	Collegiality	3.15	3	3.0	0.74
S9	Pedagogy	3.16	3	3.0	0.61
S10	Blended Efficacy	3.00	3	3.0	0.98
S11	Perceived Student Impact	3.20	3	3.0	0.83
S12	Perceived Student Digital Readiness	2.98	3	3.0	0.71

5.2.1 Modal Rank of 3

For every non-administrative spectrum (S2 to S12), rank 3 is the modal score at the population level. At the individual level, however, most respondents do not sit at rank 3 across all spectra. Teachers show peaks and troughs across the 12 dimensions, but these peaks and troughs do not align across teachers into a sharper collective pattern. What emerges at the aggregate is not uniformity but a distribution. Together, the aggregate settles at the functional midpoint; individual variation persists within it. As a system, it is stable but merely functionally adequate.

5.2.2 Tight Clustering of the Average

Excluding S1 (Years of Service), the means across S2 to S12 occupy a narrow range: from 2.97 (S3 Facilitating Conditions) to 3.49 (S4 Volition / Agency). On a 6-point scale, that is a spread of roughly half a rank across all 11 adoption-related spectra. No spectrum sits clearly above or below the midpoint.

The spread of individual scores around each mean is also modest. Standard deviations across S2 to S12 range from 0.61 to 0.98, meaning that for most spectra, the typical teacher sits within one rank of the spectrum's mean. The distribution is tight both between spectra and within spectra.

Read together, the pattern is consistent. The midpoint does not describe routine, high-quality adoption. It describes conditional adoption: teachers who see value in EdTech but whose use depends on a set of conditions that are only partially in place.

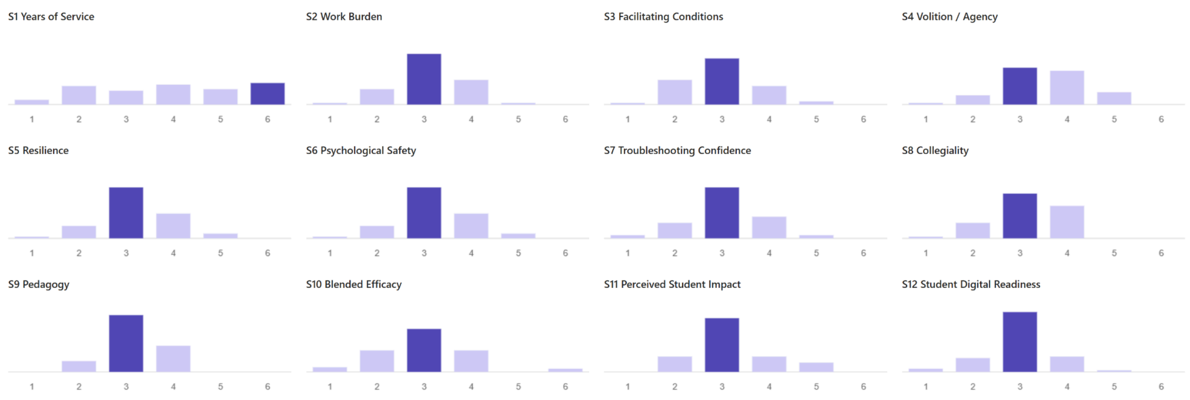

5.2.3 Tone of Voice Across Spectra

We also performed a tone of voice ("ToV") analysis to describe how teachers feel when responding in accordance to each spectra. ToV was coded on each transcript across every non-administrative dimension as Positive (P), Unresolved (U), or Negative (N). S1 was excluded as it carries no affective content.

The overall ToV at the respondent level was P = 25 (41%), U = 27 (44%), and N = 9 (15%). Taken together, 52 of 61 teachers (85.2%) were coded Positive or Unresolved. Broadly negative dispositions toward EdTech are rare in this sample.

The per-spectrum breakdown sharpens the picture. Positive tone concentrates in a small number of spectra. Negative tone concentrates in one.

Spectrum	**Positive (%)**	**Unresolved (%)**	**Negative (%)**	**Dominant**
S2 Work Burden	28	56	16	U
S3 Facilitating Conditions	30	39	**31**	U
S4 Volition / Agency	**49**	41	10	**P**
S5 Resilience	28	56	16	U
S6 Psychological Safety	33	49	18	U
S7 Troubleshooting Confidence	26	57	16	U
S8 Collegiality	33	54	13	U
S9 Pedagogy	**48**	44	8	**P**
S10 Blended Efficacy	33	51	16	U
S11 Perceived Student Impact	**46**	43	11	**P**
S12 Student Digital Readiness	18	**64**	18	U

Three patterns emerge.

Ertmer’s 2nd-Order Barriers Are Not the Binding Constraint

Positive tone concentrates on S4 (49%), S9 (48%), and S11 (46%). These are the spectra of conviction, pedagogy, and perceived student benefit. When teachers talk about why they use EdTech, how they teach with it, and what they see it doing for students, the affective register lifts. This is the affirmative core of the dataset.

Impact of Ertmer’s 1st-Order Barriers Initially Underestimated

Second, S3 (Facilitating Conditions) carries the highest negative proportion of any spectrum at 31%. The nearest others sit below 20%. S3 is the only dimension where N approaches P in magnitude (30% P vs 31% N). Infrastructure friction is the one part of the adoption experience teachers are willing to name in affectively negative terms.

Student Digital Readiness: Reported but not Evaluated

Third, S12 (Student Digital Readiness) records the highest unresolved proportion at 64%. Teachers describe student readiness factually rather than with affect. This reads as descriptive reporting rather than evaluation.

Read alongside the rank data in sections 5.2.1 and 5.2.2, the tone picture completes the picture. Teachers are engaged about conviction, pedagogy, and student outcomes. What they resist affectively is the infrastructure that carries the work, and what they report flatly is the readiness of the learners. This sets up the discussion that follows where the gap between positive disposition and midpoint practice becomes the central puzzle.

5.3 Correlation Overview

To examine how the 12 spectra relate to one another, Pearson correlations were computed for all 66 pairwise combinations (S1 to S12). Strength of association was interpreted using Cohen's R^2 thresholds (Cohen, 1988):

Cohen's R^2 thresholds	**Strength of association**
<0.01	negligible
0.01–0.09	weak
0.09–0.25	moderate
0.25–0.49	strong
≥ 0.49	Very strong

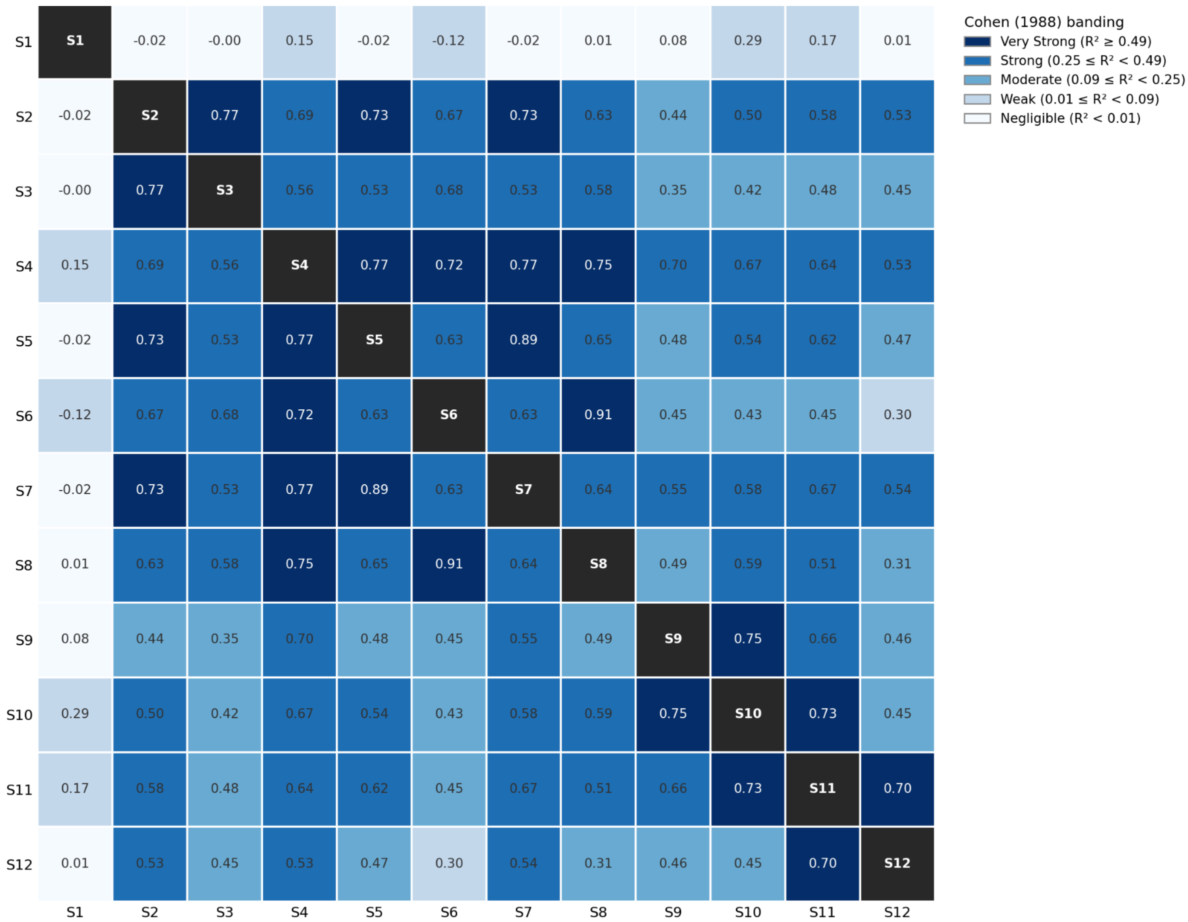

Cohen (1988) banding
Very Strong ($R^2 \geq 0.49$)
Strong ($0.25 \leq R^2 < 0.49$)
Moderate ($0.09 \leq R^2 < 0.25$)
Weak ($0.01 \leq R^2 < 0.09$)
Negligible ($R^2 < 0.01$)
S1 S1 -0.02 -0.00 0.15 -0.02 -0.12 -0.02 0.01 0.08 0.29 0.17 0.01
S2 -0.02 S2 0.77 0.69 0.73 0.67 0.73 0.63 0.44 0.50 0.58 0.53
S3 -0.00 0.77 S3 0.56 0.53 0.68 0.53 0.58 0.35 0.42 0.48 0.45
S4 0.15 0.69 0.56 S4 0.77 0.72 0.77 0.75 0.70 0.67 0.64 0.53
S5 -0.02 0.73 0.53 0.77 S5 0.63 0.89 0.65 0.48 0.54 0.62 0.47
S6 -0.12 0.67 0.68 0.72 0.63 S6 0.63 0.91 0.45 0.43 0.45 0.30
S7 -0.02 0.73 0.53 0.77 0.89 0.63 S7 0.64 0.55 0.58 0.67 0.54
S8 0.01 0.63 0.58 0.75 0.65 0.91 0.64 S8 0.49 0.59 0.51 0.31
S9 0.08 0.44 0.35 0.70 0.48 0.45 0.55 0.49 S9 0.75 0.66 0.46
S10 0.29 0.50 0.42 0.67 0.54 0.43 0.58 0.59 0.75 S10 0.73 0.45
S11 0.17 0.58 0.48 0.64 0.62 0.45 0.67 0.51 0.66 0.73 S11 0.70
S12 0.01 0.53 0.45 0.53 0.47 0.30 0.54 0.31 0.46 0.45 0.70 S12
S1 S2 S3 S4 S5 S6 S7 S8 S9 S10 S11 S12

We then extract the ten strongest correlations, grouped by construct:

Pair	r	R^2	Cohen band
Dispositional / Cultural			
S6 Psychological Safety x S8 Collegiality	+0.905	0.820	Very strong
S5 Resilience x S7 Troubleshooting Confidence	+0.885	0.783	Very strong
Ertmer's First-Order Conditions			
S2 Work Burden x S3 Facilitating Conditions	+0.772	0.596	Very strong
S2 Work Burden x S7 Troubleshooting Confidence	+0.730	0.533	Very strong
Ertmer's Second-Order Conditions			
S4 Volition x S7 Troubleshooting Confidence	+0.769	0.592	Very strong
S4 Volition x S5 Resilience	+0.769	0.591	Very strong
S4 Volition x S8 Collegiality	+0.751	0.564	Very strong
Pedagogical Practice and Impact			
S9 Pedagogy x S10 Blended Efficacy	+0.750	0.562	Very strong
S10 Blended Efficacy x S11 Perceived Student Impact	+0.733	0.537	Very strong
S11 Perceived Student Impact x S12 Digital Readiness	+0.701	0.492	Very strong

5.3.1 Initial Observations

S1 (Years of service) is not a meaningful predictor of any adoption-related spectrum in this dataset. Across all 11 of its pairings, 7 are negligible and 4 are weak.

S4 (Volition / Agency) appears in both cluster groupings. It correlates strongly with the dispositional and cultural cluster (S5, S6, S7, S8) and also with the pedagogical practice and impact cluster (S9, S10, S11). No other spectrum behaves this way. S4 functions as a bridge variable between how teachers feel about their working conditions and how they translate that into classroom practice. This bridging role will be further explored.

5.4 Headline Finding: Conditional Adoption

The defining finding of this study is not that teachers resist EdTech. It is that they adopt it conditionally. Teachers speak about conviction, pedagogy, and student outcomes in positive affective terms. They are not hostile, but instead see value.

Yet every non-administrative spectrum has rank 3 as its modal score. This implies that our system sits at the functional midpoint. This is the central puzzle:

Why does positive belief produce only adequate practice?

While functional adequacy is not failure, it is also not flourishing. This happens when teachers who believe in the tool deliver EdTech-integrated lessons under conditions that are only partially in place. They adopt EdTech to the extent that conditions permit. When conditions are partial, adoption is partial.

5.4.1 Evidence 1: Teachers are not resentful

Teachers speak about EdTech with conviction more often than with resentment. Across the 61 transcripts' ToV,

- 52 (85.2%) were coded as overall Positive or Unresolved overall.
- 49% of teachers were coded Positive on S4 (Volition / Agency), the highest proportion of any spectrum.

5.4.2 Evidence 2: Teachers perceive value

Teachers perceive real value for students. On S11 (Perceived Student Impact), 51 of 61 teachers (83.6%) scored between 3 and 5. They see EdTech producing moderate-to-strong benefits in their classrooms. Quotes across the dataset return repeatedly to the same affordances: immediate feedback, visibility of every student's response, data for reteaching, differentiated pacing.

5.5 Invisible Workload as the Hidden Chain

The gap between conviction and practice is where the invisible workload sits.

5.5.1 Adopting EdTech Perceived as Incurring Personal Cost

Several teachers describe the same pre-delivery lesson-preparation sequence.

- Decide on the lesson objective.
- Evaluate if EdTech adds value.
    - If yes, select the app.
        - Write the prompts.
        - Book the iPads or laptop trolley.
        - Test the platform.
        - Troubleshoot when it fails.

Respondent “Breadtalk” calls this preparation "unseen." Respondent “UserA” distinguishes between EdTech making learning better, and EdTech making teaching easier. His conclusion is that good use of technology makes learning better, but not necessarily easier for the teacher.

This preparation happens outside visible teaching time. Respondent “Milktea” surmises that most of that extra prep comes from (his) personal time. Respondent “Red bean” prepares during free periods and calls the work "unseen." Respondent “Lola”, a veteran with 37 years of teaching experiences, prepares for lessons during rest time at home. Her S10 (Blended Efficacy) score is 4, above what her facilitating conditions (S3 Facilitating Conditions = 2) would predict. “Lola” achieves that score at a measurable personal cost.

The invisible workload is the mechanism that converts conviction into practice. Where the workload is absorbed, adoption progresses. Where it is not, adoption stalls at the midpoint regardless of willingness.

5.5.2 Invisible Burden is Perceptually Real

This is the Invisible Burden. It is not the absence of belief, nor the absence of skill. It is the uncounted hours between the decision to use EdTech and the moment students begin the task. The burden is invisible because it sits outside the curriculum, outside the timetable, and outside the measures the system uses to judge good teaching.

5.5.3 Correlations Involving Invisible Burden

The statistical correlates follow this logic. S2 (Work Burden) shows Very Strong correlations with

- first-order conditions S3 (r=0.772), S7 (r=0.730),

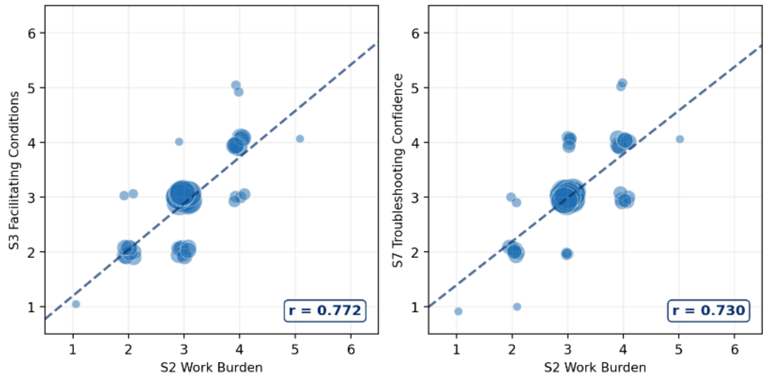

- second-order dispositions S5 (r=0.728), S4 (r=0.692), and S6 (r=0.672).

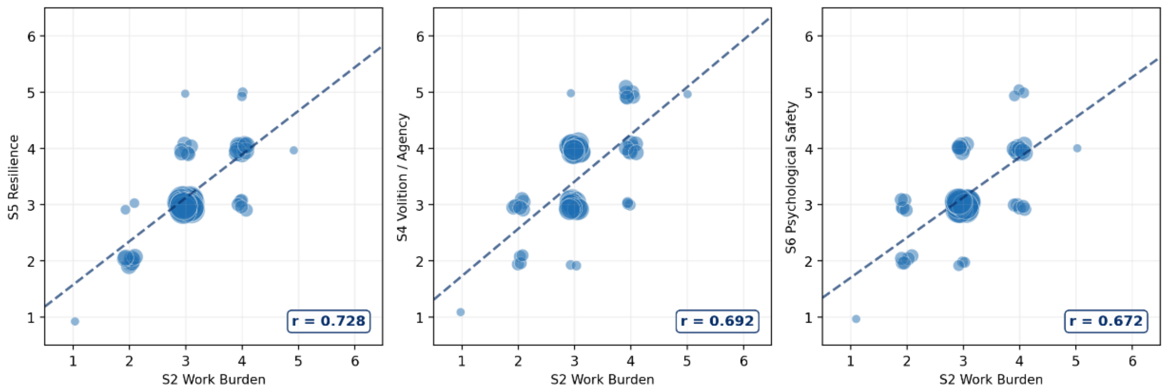

When teachers report higher workload, they also report weaker facilitating conditions, lower troubleshooting confidence, lower resilience, lower volition, and lower psychological safety. Burden does not sit alone. It propagates across the full dispositional landscape.

5.5.4 Implications of Findings

Conditional adoption cannot be solved by training alone

Training addresses competence. It does not reduce the invisible workload that competent teachers still carry.

System depends on teachers who absorb the burden privately

The teachers who reach the upper bands of S10 (Blended Efficacy) and S11 (Perceived Student Impact) in this sample population are often the teachers who work from home during rest periods. This is a finding about sustainability, not just adoption.

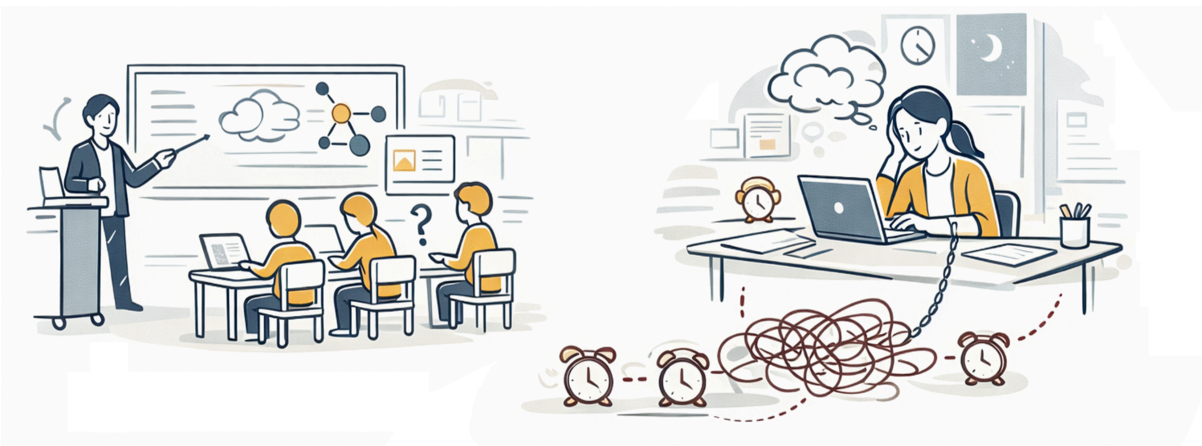

5.6 First-Order and Second-Order Conditions

Ertmer (1999) distinguishes between first-order barriers (infrastructure, access, time, training) and second-order barriers (beliefs, pedagogical orientation, willingness to change). The correlation structure of this dataset suggests that both categories are active, but they do not carry equal weight.

5.6.1 First-Order Conditions: Necessary but Not Sufficient

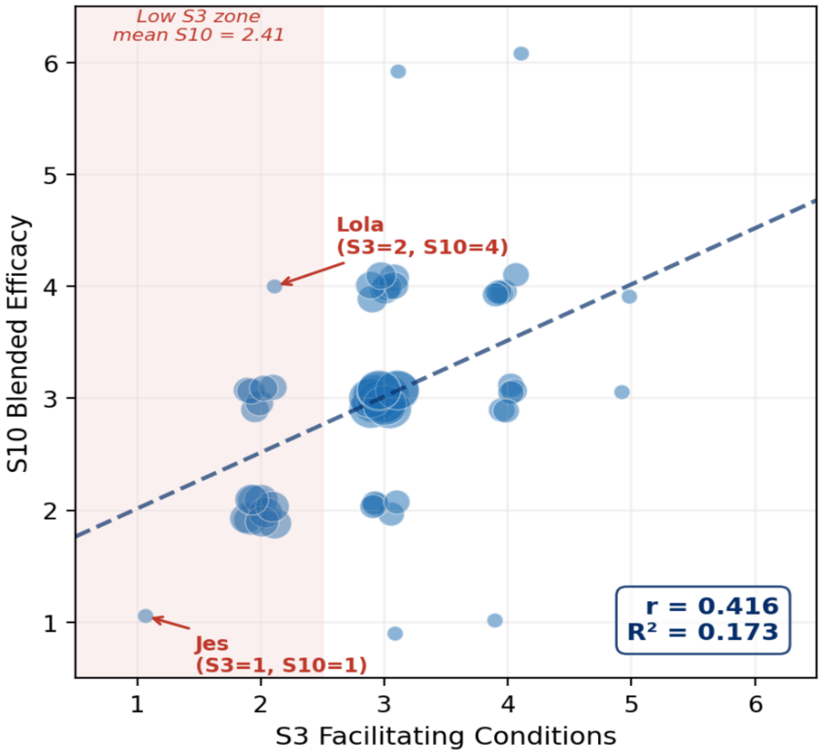

S3 (Facilitating Conditions) sets a floor on practice. Better infrastructure is correlated with better blended practice. Among the 17 teachers who scored S3 at 2 or below, the mean S10 (Blended Efficacy) score was 2.41. Among the 14 teachers who scored S3 at 4 or above, the mean S10 was 3.50.

But S3 (Facilitating Conditions) explains only 17.3% of the variance in S10 (Blended Efficacy) ($r=0.416$, $R^2=0.173$). Infrastructure sets a ceiling. It does not determine where under that ceiling a teacher lands.

Two respondents illustrate the boundaries. Respondent "Jes" (S3=1, S10=1) is the floor case. Her entire interview centres on device logistics. She cannot get past the first-order barrier. There is no pedagogical content to code for because she perceived that the infrastructure had never permitted a lesson to begin with.

Respondent "Lola" (S3=2, S10=4) is the exception that proves the rule. She achieves blended practice well above what her facilitating conditions would predict. But this comes at personal cost. She prepares during rest time at home. Her resilience (S5=2) and psychological safety (S6=2) scores are both low. She has compensated for structural failure with private effort.

Respondent "1960s" captures the first-order experience in operational terms. In one lesson, about half the students failed to log in and were locked out of the platform before the task could begin. When infrastructure collapses, no amount of conviction or pedagogical skill can rescue the lesson.

5.6.2 Second-Order Conditions: The Tighter Construct

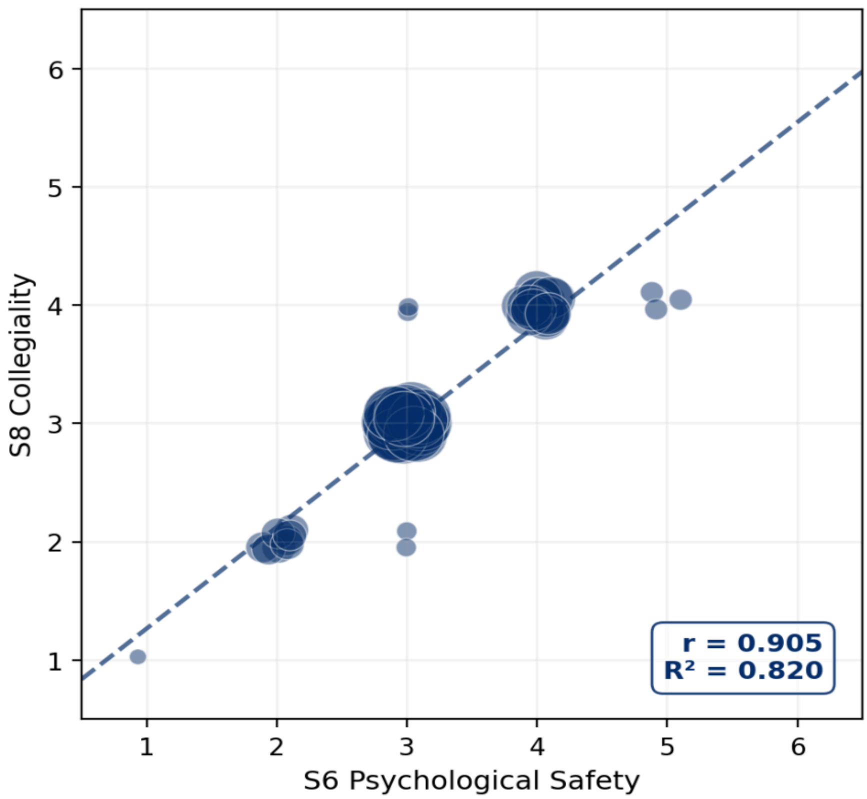

The strongest correlation in the entire dataset is not between infrastructure and practice. It is between two second-order spectra: S6 (Psychological Safety) and S8 (Collegiality), at r=0.905, R^2=0.820.

This is not a coincidence of measurement. Psychological safety and collegiality are conceptually distinct. S6 captures whether teachers feel safe to experiment, fail, and ask for help. S8 captures whether they actually share practices, resources, and feedback with colleagues. That these two move almost in lockstep says something about the schools in this sample: where safety exists, sharing follows. Where it does not, teachers work alone.

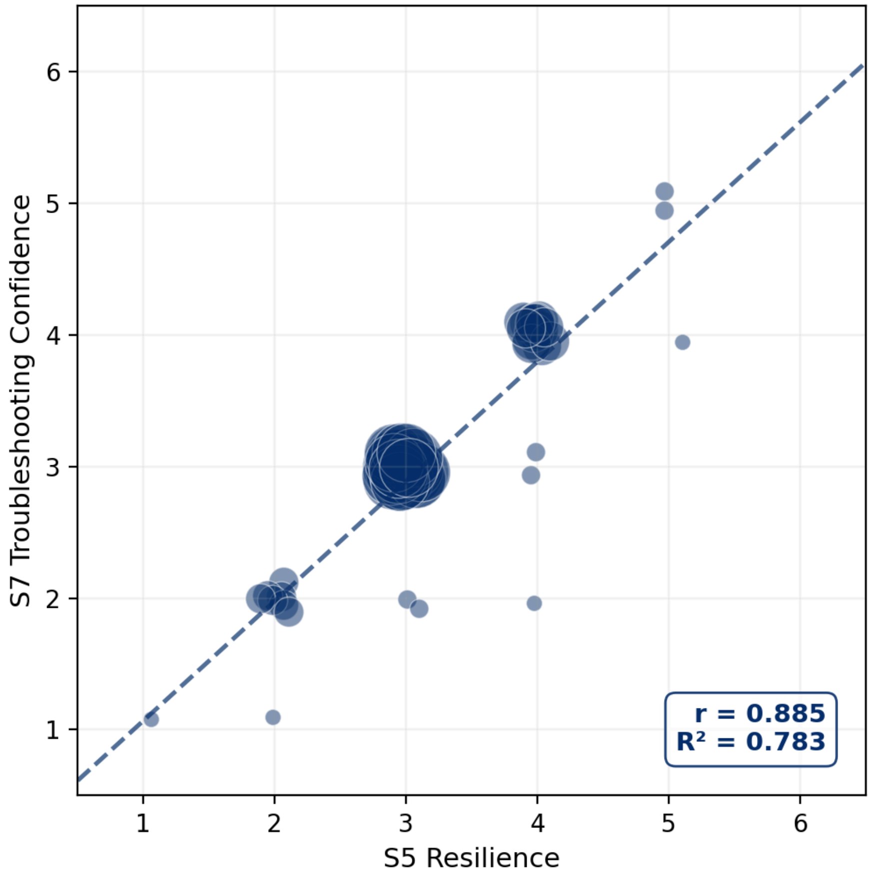

The second-tightest pair is also second-order: S5 (Resilience) and S7 (Troubleshooting Confidence), at r=0.885, R^2=0.783. Teachers who persist when things go wrong are also the teachers who can fix what goes wrong. These are not the same skill, but they rise and fall together.

Taken together, the second-order spectra form a coherent dispositional cluster. S6×S8 and S5×S7 are its anchors. S4 (Volition / Agency) bridges this cluster to the practice spectra.

5.6.3 Conditioning a Systemic Response

First-order conditions are prerequisites. Without easy access and availability of functioning devices, reliable networks, and bookable equipment, EdTech lessons do not happen. The data confirms this.

But the data also shows that once first-order conditions reach a functional threshold, it is the second-order environment that determines how far practice develops. School culture, not infrastructure spending, is the stronger predictor of adoption quality.

This has a practical consequence. Schools that invest heavily in devices and platforms but neglect psychological safety and collegial sharing are likely to see adoption plateau at the midpoint. The infrastructure will be in place. The culture will not carry it further.

5.6.4 Shifting the Needle Under Constraints

When working within constraints, the data suggest staged sequencing. First-order conditions need to reach a functional minimum, not a gold standard. Teachers do not need the newest devices. They need devices that are available, bookable, and connected when the lesson begins.

Once that threshold is met, further spending on hardware produces diminishing returns.

The higher-leverage investment is in the second-order environment. It requires time and structure. Three possible moves sit within most school leaders' reach:

Protected time for collegial sharing

Collegial sharing does not happen reliably in informal settings. It requires scheduled, low-stakes opportunities for teachers to show each other what they have tried, including what did not work. Respondent "mh" describes this in practice: "Professional dialogues with colleagues during TTT also help to broaden my perspectives." The key word is "during." The sharing happens because the time is protected.

Online communities such as the Singapore Learning Designers Circle (“SgLDC”) serve a complementary but different function. They provide visibility into what other educators are doing, and at their best they circulate ideas and tools widely.

However, the design of large showcase communities can inadvertently produce comparison rather than collaboration. A teacher at rank 3 who sees a polished example from a teacher at rank 5 may feel behind rather than supported. The S6×S8 finding suggests that what moves practice is not exposure to excellence but safety to experiment. That safety is built in small teams, face to face, with permission to fail. It is not built by scrolling or inducing “fear of missing out” in teachers.

Acknowledging and reducing the invisible workload

Section 5.5 identified the metaphorical “chain” that sits between the decision to use EdTech and the moment students begin the task. Much of this chain is operational, not pedagogical: booking devices, configuring accounts, testing platforms, troubleshooting failures.

We can shorten this chain through practical steps: pre-configured device sets, shared lesson templates, centrally managed platform accounts, and streamlined booking systems. Each of these removes a link from the chain that currently falls on individual teachers. The goal is not to eliminate preparation. It is to reduce the proportion of preparation that is logistical rather than instructional.

Respondent "Kay" names the bottleneck plainly: "Planning the resources as well as getting the iPads." If getting the iPads were not part of the teacher's job, the planning could focus on learning.

Building volition

The S4 bridge finding (Section 5.3) shows that conviction connects disposition to practice. S4 correlates with both the dispositional cluster (S5, S6, S7, S8) and the practice cluster (S9, S10, S11). Conviction is the hinge.

If specific tools or platforms are mandated into usage without teacher input, there is a systemic risk that S4 (Volition / Agency) will degrade from being an asset into a compliance exercise.

Respondent "Happy" captures this precisely: "Will use only if need to and only if it helps to make lessons more effective." Her conviction is high (S4=5) because it is self-directed. Respondent "1960s" names the risk from the other direction: "High expectations lead to 'forced' adoption of technology which is not sustainable."

The practical implication is to mandate the outcome, not the tool. Set the expectation that lessons use technology where it adds value, then let teachers choose how. This preserves the agency that S4 measures and that the data shows is necessary for practice to follow.

5.7 Pedagogy × Blended Efficacy for Impact

The previous sections examined what holds EdTech adoption at the midpoint. This section examines what moves it forward.

The data identifies a directional chain: S9 (Pedagogy) shapes S10 (Blended Efficacy), and S10 (Blended Efficacy) is the strongest single predictor of S11 (Perceived Student Impact). This is the pathway through which EdTech reaches students.

5.7.1 Pedagogical Orientation Shapes Blending

S9 (Pedagogy) and S10 (Blended Efficacy) correlate at r=0.750, R^2=0.562. This correlation suggests that teachers with a more constructivist pedagogical orientation are more likely to blend technology into their teaching in sustained and purposeful ways.

This is not a claim about ideology because S9 (Pedagogy) measures what teachers do: whether they design lessons where students interact with content, produce responses, and receive feedback through technology, or whether technology serves primarily as a presentation tool.

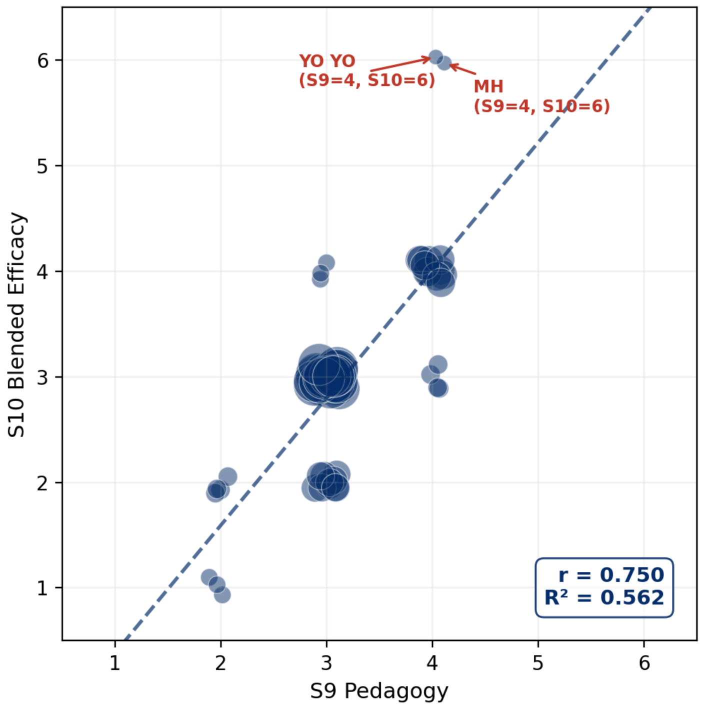

Respondent "Jellyfish" (S9=4 and S10=4.) articulates the design-led approach, "Be clear of lesson objectives, know my success criteria, plan logical flow in students' use of tech, provide anchor points." Her blended practice follows from the pedagogical design.

Respondent "mh" (S9=4 and S10=6) shows what happens when pedagogical orientation meets sustained execution. She describes a systematic cycle, SLS pre-assessment, ClassPoint during the lesson, Data Assistant for analysis, reteaching based on results, then Google Docs for student collaboration. She is one of only two teachers at the S10 ceiling. The other is Respondent "YO YO" (S9=4 and S10=6), who built a blended writing loop in Chinese using SchoolAI.

Both ceiling cases share a feature. Their blended practice is not a single tool applied to a single lesson. It is a workflow: multiple platforms connected in sequence, each serving a specific pedagogical function. This level of integration does not emerge from tool training. It emerges from a pedagogical framework that treats technology as a means, not an end.

5.7.2 Blended Efficacy Drives Perceived Impact

Fig. 5.7.2: Top three predictors of S11 Perceived Student Impact (n = 61)

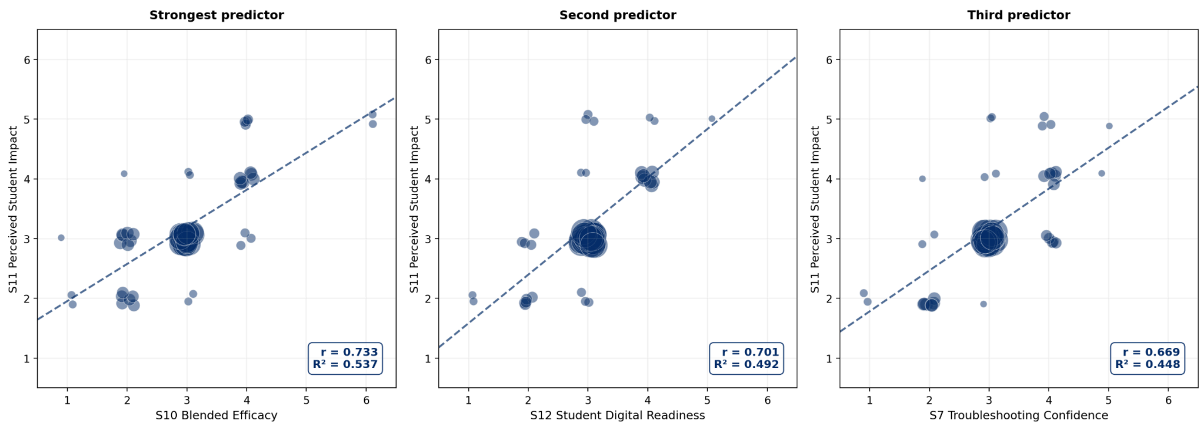

S10 and S11 correlate at r=0.733, R^2=0.537. The next strongest predictor of S11 is S12 (Student Digital Readiness) at R^2=0.492, followed by S7 (Troubleshooting Confidence) at R^2=0.448.

This means that how well a teacher blends technology into instruction matters more for perceived student outcomes than any other factor the framework measures.

The affordances teachers cite are consistent. Respondent "Michael" values technology because it lets him "get students' responses throughout the lesson as a form of feedback" and "elicit misconceptions and alternative solutions for class discussion." Respondent "doremi" describes students receiving "immediate feedback" that allows them to "reflect on their understanding and make improvements." Respondent "Jojo" names the visibility affordance: "Using ICT allows me to hear from all students, which without ICT I can't."

These are not descriptions of novelty. They are descriptions of formative assessment made scalable. The perceived student impact is not about engagement or motivation. It is about feedback loops that would not exist without the technology.

5.7.3 Two Anomalies

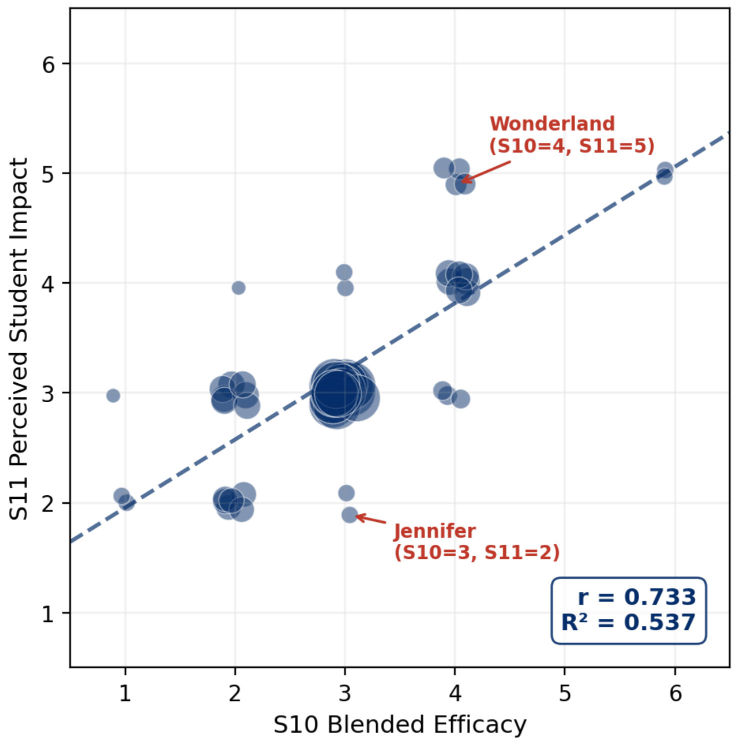

Respondent "Wonderland" scores S11=5 with S9=3. She achieves high perceived student impact without a strongly constructivist orientation. Her method is data-driven reteaching: she uses student response data to adjust instruction, but within a largely instructivist frame. This suggests that the S9→S10→S11 chain is the primary pathway, but not the only one. A teacher who uses technology narrowly but effectively for data-informed decisions can still perceive strong impact.

Respondent "Jennifer" presents the inverse. She scores S4=4 (high conviction) but S11=2 (low perceived impact). Her interview describes a chatbot lesson that confused students. Conviction without pedagogical fit produces frustration, not outcomes. Jennifer confirms that S4 is necessary but not sufficient. The pathway requires S9 and S10 to carry conviction into practice.

5.7.4 Implications

The practical chain is: conviction (S4) shapes pedagogical orientation (S9), which shapes blended practice (S10), which produces student impact (S11).

Each link must hold. Conviction alone does not produce impact (respondent “Jennifer”). Blended practice without constructivist framing can still work if the teacher uses data well (respondent “Wonderland”). But the dominant pathway runs through all four.

For school leaders, this means that professional development aimed at improving student outcomes through EdTech should focus on pedagogical design, not tool proficiency. The question is not "can you use ClassPoint?" It is "what learning goal does ClassPoint serve in this lesson, and how will you use the data it produces?"

5.8 Student Digital Readiness

S12 (Student Digital Readiness) is the second-strongest predictor of S11 (Perceived Student Impact) at $r=0.701$, $R^2=0.492$. It sits just behind S10 (Blended Efficacy) in predictive weight. This is a finding about the receiving end of the adoption chain. What teachers perceive as impact depends partly on whether students can use the technology placed in front of them.

5.8.1 The Tightest Distribution

S12 (Student Digital Readiness) shows the least spread of any spectrum. 39 of 61 teachers (63.9%) sit at rank 3. Only one teacher scores S12 at rank 5. Only two score S12 at rank 1.

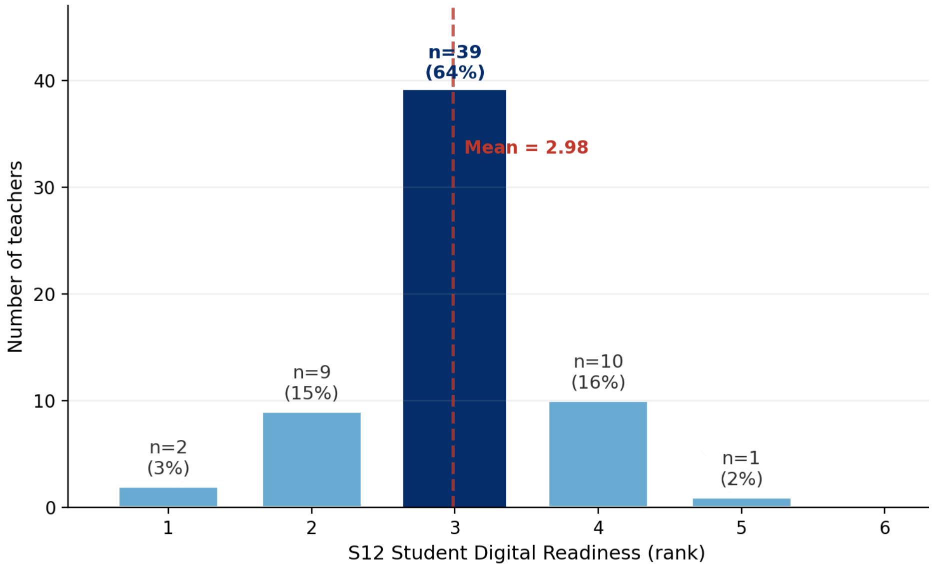

This distribution shows that teachers broadly agree that their students are moderately ready for digital learning. The 64% Unresolved tone coding on S12, the highest of any spectrum, reinforces this reading. Teachers describe student readiness as a condition they observe, not as something they feel strongly about. It is reported, not evaluated.

5.8.2 Where Readiness Constrains Impact

Respondent "DEY" is the clearest case. She teaches Primary 1 and scores S12=1, the joint-lowest in the sample. Her description is concrete: "P1s have difficulty going onto SLS platform due to their typing speed and readiness... all these take time." The barrier is not pedagogical. It is developmental. Six-year-olds cannot type fast enough to complete tasks within lesson time. DEY scores S11=2. She sees limited student impact because her students cannot yet operate the tools.

DEY also calls EdTech adoption "a must-do." Her S4 (Volition / Agency) is 3. She is not resistant. She is structurally blocked by the age of her learners. This is not a second-order barrier. It is a student-side constraint that sits outside the teacher's control entirely.

Across the 11 teachers who score S12 at 2 or below, the pattern is consistent. Their mean S11 is 2.27. Low student readiness is associated with low perceived impact regardless of what the teacher does.

5.8.3 Where Readiness Enables Impact

Respondent "Jojo" is the only teacher at S12=5. She scores S11=5. Her framing is specific: "Using ICT allows me to hear from all students, which without ICT I can't." The affordance she names, visibility of every student's response, only works if students can produce those responses. Her students can and EdTech delivers.

Among the 11 teachers who score S12 at 4 or above, the mean S11 is 4.27. High student readiness does not guarantee high perceived impact, but it removes the floor that low readiness imposes.

5.8.4 Implications

Student digital readiness is not a teacher-side variable. It cannot be addressed through professional development or school culture alone. It is shaped by student age, prior exposure, home access, and curriculum sequencing.

MOE's Guide to the Development of Digital Literacy and Technological Skills (2025) provides a structured progression. Basic digital skills are mapped to key junctures: device operations and login at P1, file management at P3/4, keyboarding at P5/6, and system troubleshooting at Secondary 1. The 9 Digital Competencies under the FTAC frame (Find, Think, Apply, Create) set developmental milestones across primary, secondary, and pre-university levels.

The guide is clear that these milestones are reference points for the design of learning experiences, not criteria for student assessment. Schools are expected to customise success criteria to suit their own profiles and needs. This design is deliberate. It allows for the variation in student readiness that this dataset documents.

However, the gap between policy intent and classroom experience remains visible in the data. Respondent "DEY" teaches P1. Her students struggle with login speed and platform navigation. The guide places basic device operations and login at the P1 juncture, but DEY's experience suggests that the time required for these foundational skills competes directly with curriculum time for EdTech-integrated lessons. At the P1 level, building digital readiness is itself the lesson. Expecting technology integration outcomes comparable to upper primary at this stage misreads where students sit on the progression.

For school leaders, this means that EdTech adoption expectations should be differentiated by level. The DLTS guide provides the learning outcomes. What schools need to define are the success criteria that reflect realistic baselines at each juncture. A P1 class where every student logs in independently by Term 3 is a different achievement from a P5 class running a collaborative Google Docs task. Both are valid. They should not be measured on the same scale.

5.9 What Years of Service Does Not Tell Us

A common assumption in EdTech policy is that experience predicts adoption. Veteran teachers may be assumed to be less willing or less able to integrate technology. Beginning teachers are assumed to be more ready.

Our data does not support either claim.

5.9.1 S1 Is a Null Predictor

S1 (Years of Service) has 11 pairwise correlations with the adoption-related spectra (S2 to S12). 7 are negligible ($R^2 < 0.01$). 4 are weak (R^2 between 0.01 and 0.09).

This means that experience does not predict

- workload perception (S1×S2: r=-0.023)
- resilience (S1×S5: r=-0.016)
- troubleshooting confidence (S1×S7: r=-0.023), or
- predict collegiality (S1×S8: r=0.008).

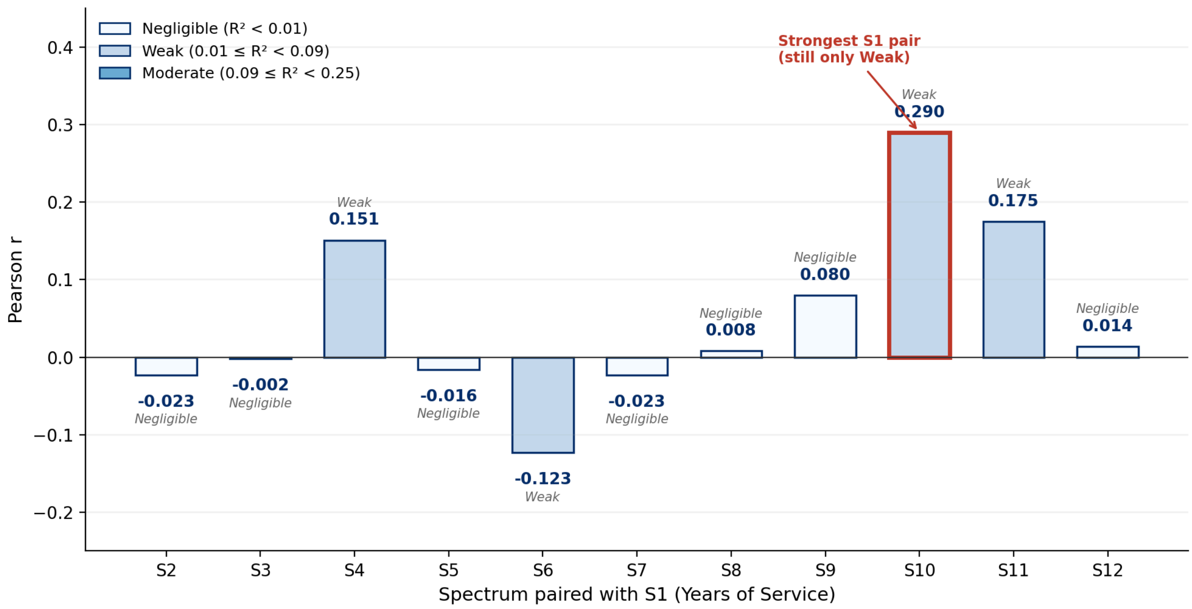

In the heatmap, the S1 row and column are visibly pale against the dense blue of the remaining spectra. Experience is disconnected from every dimension the framework measures.

5.9.2 The Veteran Burden Paradox

If experience does not predict adoption, what does the veteran group actually look like?

Among the 14 teachers coded S1=6 (21+ years), S10 scores range from 2 to 6. The mean is 3.36. The spread is enormous. Respondent "mh" (S1=6, S10=6) runs a systematic blended workflow across four platforms. Respondent "BOT" (S1=6, S10=2) says: "So many more new software/apps in the market, no time to learn." Respondent "Techtired" (S1=6, S10=2) has the only S7=1 in the dataset. These are teachers with the same career length and completely different adoption profiles.

Seven veterans (S1≥5) score S2 at 2 or below: “BOT”, “Breadtalk”, “Lola”, “Techtired”, “red bean”, “Joyce (EL/SS)”, and “DEY”. These are teachers who report high workload burdens despite long careers. Their experience has not insulated them from the invisible workload described in Section 5.5. In some cases it has deepened it. Respondent “BOT” has been teaching for 25 years. Her response to EdTech is exhaustion, not resistance.

This is the Veteran Burden Paradox. Long service does not protect against workload burden. It does not build resilience or troubleshooting confidence. What it does, in a small number of cases, is provide the time to develop deep pedagogical integration (respondents “mh”, “YO YO”). But these are outliers. Most veterans sit at the same midpoint as everyone else.

5.9.3 Implications

The practical consequence is direct. Schools should not use years of service as a proxy for EdTech readiness. Segmenting teachers into "digital natives" and "digital immigrants" by career stage misreads the data. A beginning teacher with low S4 (Volition) may be less ready than a veteran with high S4.

The segmentation that matters is by conviction (S4), not by career stage (S1). Section 5.3 showed that S4 bridges the dispositional and practice clusters. Section 5.6 showed that second-order conditions outweigh first-order conditions. S1 sits outside both clusters. It is an administrative fact, not an adoption variable.

Professional development should not be differentiated by age cohort. It should be differentiated by where teachers sit on the spectra that actually predict practice: S4, S9, and S10.

5.9.4 The Neutral Majority: Compliance Without Conviction

At 44% of the sample, respondents with overall Unresolved tone constitute the single largest affective category. The dataset does not divide into enthusiasts and resistors. Most teachers occupy a more equivocal position: functionally compliant, neither burdened nor energised, using technology because it is expected without either strong aversion or strong investment.

This affective configuration is the direct correlate of the modal rank-3 finding. A teacher who scores 3 across most spectra and codes Unresolved on overall tone is not disengaged. She is operating within the system as the system currently holds her.

The neutral majority is not indifferent. Many score Positive tone on S4 or S11 when describing specific moments of pedagogical conviction or observed student benefit respectively. Respondent "Michael" is coded Unresolved overall, but his tone lifts when he describes using technology to "get students' responses throughout the lesson as a form of feedback" and "elicit misconceptions and alternative solutions for class discussion." The conviction exists. It is situational, not sustained.

Their neutral tone is holistic rather than dimensional. It is an absence of activated investment at the level of the whole EdTech experience, not an absence of belief in its potential value.

The intervention question that follows is therefore not how to persuade reluctant teachers to use more technology. It is how to convert the latent conviction already visible in S4 (Volition / Agency) and S9 (Pedagogy) tone into the sustained, purposeful practice that rank-3 teachers are not yet producing.

The barrier is not apathy. It is context.

5.10 Layered Adoption EdTech Stack

Our data suggest that EdTech adoption in this sample operates as a layered system. Four conditions stack on top of one another . Each layer is prerequisite for the one above it. Each sets a ceiling on what the next can achieve.

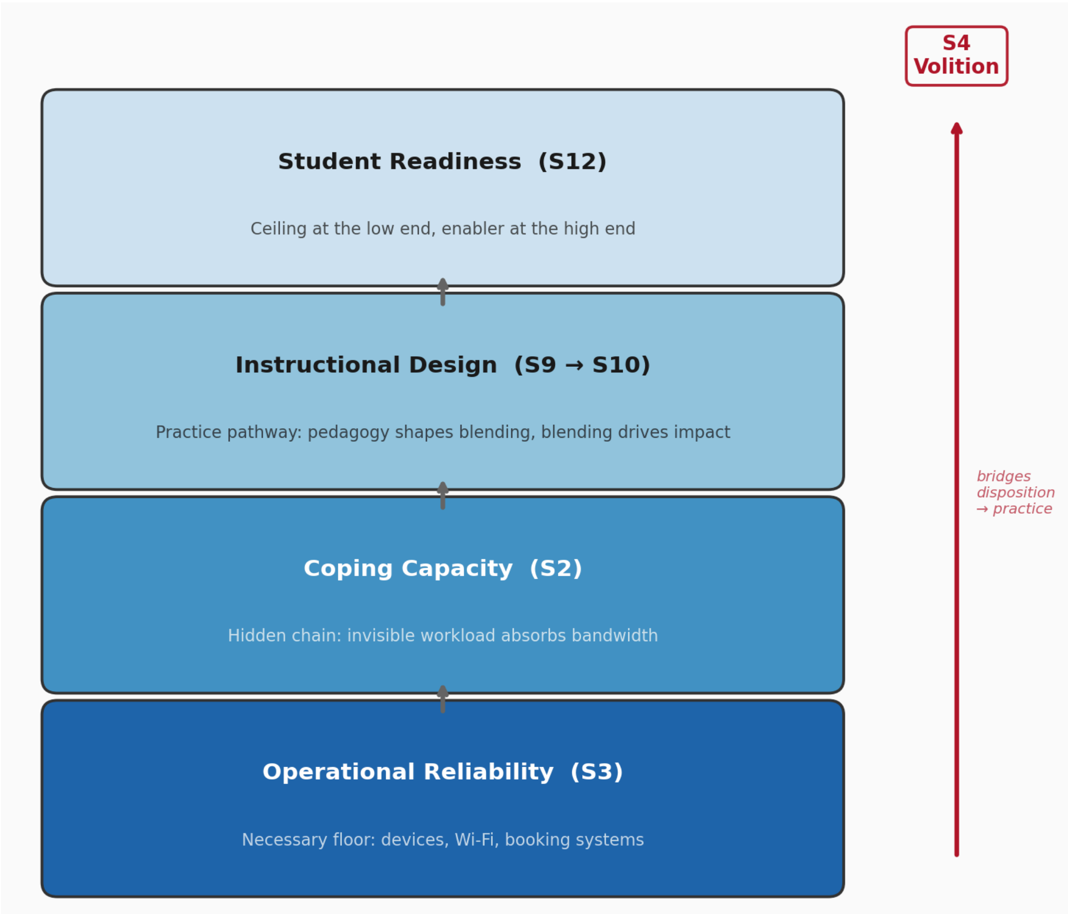

Like Bloom's taxonomy of cognitive demand, the model is cumulative: higher-order outcomes depend on lower-order foundations being in place.

Unlike Bloom's taxonomy, the layers describe systemic conditions rather than cognitive skills.

5.10.1 Foundational Layer: Operational Reliability

The foundation. Without this, nothing above it activates.

When S3 (Facilitating Conditions) falls to rank 2 or below (n=17), average S10 (blended efficacy) drops to 2.41. When S3 rises to 4 or above (n=14), it reaches 3.50.

While reliable infrastructure does not produce good practice, unreliable infrastructure prevents it. This is the prerequisite that makes all subsequent layers possible.

5.10.2 Coping Capacity

The absorption layer. Do teachers have bandwidth beyond survival?

S2 (Work burden) correlates with every adoption-related spectrum.

When invisible workload consumes coping capacity, teachers default to what is safe. Fixing operational reliability clears the floor. Layer 2 determines whether teachers have room to stand on it.

5.10.3 Instructional Design

The practice layer: where adoption becomes visible in the classroom.

The chain S9 (Pedagogy) → S10 (Blended Efficacy) → S11 (Perceived Student Impact) is the strongest sequential pathway in the dataset. But it cannot activate without the two layers below it. Respondent "Lola" achieves S10=4 despite S3=2, but only by using personal rest time. Conversely, Respondent "Wonderland" produces S11=5 from S9=3 when infrastructure holds and blended tools provide data-driven re-teaching.

This layer rewards any teacher who reaches it with working conditions intact, but punishes effectively EdTech-enabled teachers by charging them the invisible workload if the lower layers are not well-provided.

5.10.4 Layer 4: Student Readiness (S12)

The ceiling. A condition the teacher cannot fully control.

S12 (Student Digital Readiness) has the tightest distribution in the dataset (64% at rank 3). At the low end, Respondent "DEY" (S12=1) spends lesson time on login support. The planned lesson cannot run. At the high end, Respondent "Jojo" (S12=5) reports zero operational friction.

Where schools invest in foundational digital skills early, this constraint lifts by the time students reach upper primary levels. Where they do not, it persists.

5.10.5 The Bridging Variable: S4 Volition / Agency

S4 (Volition / Agency) does not sit within any single layer. It operates across all four. It is what converts adequate conditions into purposeful action. Two teachers with the same S3 (Facilitating Conditions) and S2 (Work Burden) can produce different S10 (Blended Efficacy) outcomes depending on whether they believe the effort is worth it.

5.10.6 Reading the Model

The distilled results do not claim all 4 layers must be optimised simultaneously. It claims each layer sets a ceiling on what the layer above can achieve. Progress within specific schools requires diagnosis of which layer is currently the binding constraint, then targeted action at that layer.

The question for school leadership is not "are our teachers using enough EdTech?" It is: "which layer is currently holding them back?"

6 Conclusion and Recommendations

6.1 Conclusion

This study examined how 61 Singapore primary school teachers experience EdTech adoption. We found that teachers adopt EdTech conditionally, not reluctantly. They are not resistant. They see value. But their practice stalls at functional adequacy because the conditions for sustained adoption are only partially in place.

The data, interpreted through Ertmer's (1999) framework of first-order and second-order barriers, revealed a consistent pattern. First-order barriers (infrastructure, devices, booking systems) set a necessary floor but do not drive adoption on their own. Second-order barriers (beliefs, confidence, pedagogical conviction) form a tighter construct and explain more of the variance in practice. The strongest correlations in the dataset sit within the second-order clusters:

- S6 (Psychological Safety) × S8 (Collegiality) at r=0.905, and
- S5 (Resilience) × S7 (Troubleshooting Confidence at r=0.885

These are the conditions that predict whether a teacher with working devices actually uses them well.

What we deem as “invisible workload” is situated in-between these two correlations. Teachers described a chain of preparation, testing, booking, troubleshooting, and recovery that falls outside timetables and workload measures. This burden is absorbed privately, often during rest periods and weekends. Work burden (S2) correlates with every adoption-related spectrum. It is not an isolated complaint. It is a systemic feature of the EdTech adoption experience.

S1 (Years of service) does not predict adoption. S1 has no moderate or strong correlation with any other spectrum. The segmentation that matters is by conviction (S4), not career stage.

We synthesised these findings into a four-layer model. Operational reliability (S3) forms the foundation. Coping capacity (S2) sits above it. Instructional design (S9 → S10) is the practice layer. Student readiness (S12) sets the ceiling. S4 (Volition) bridges all four.

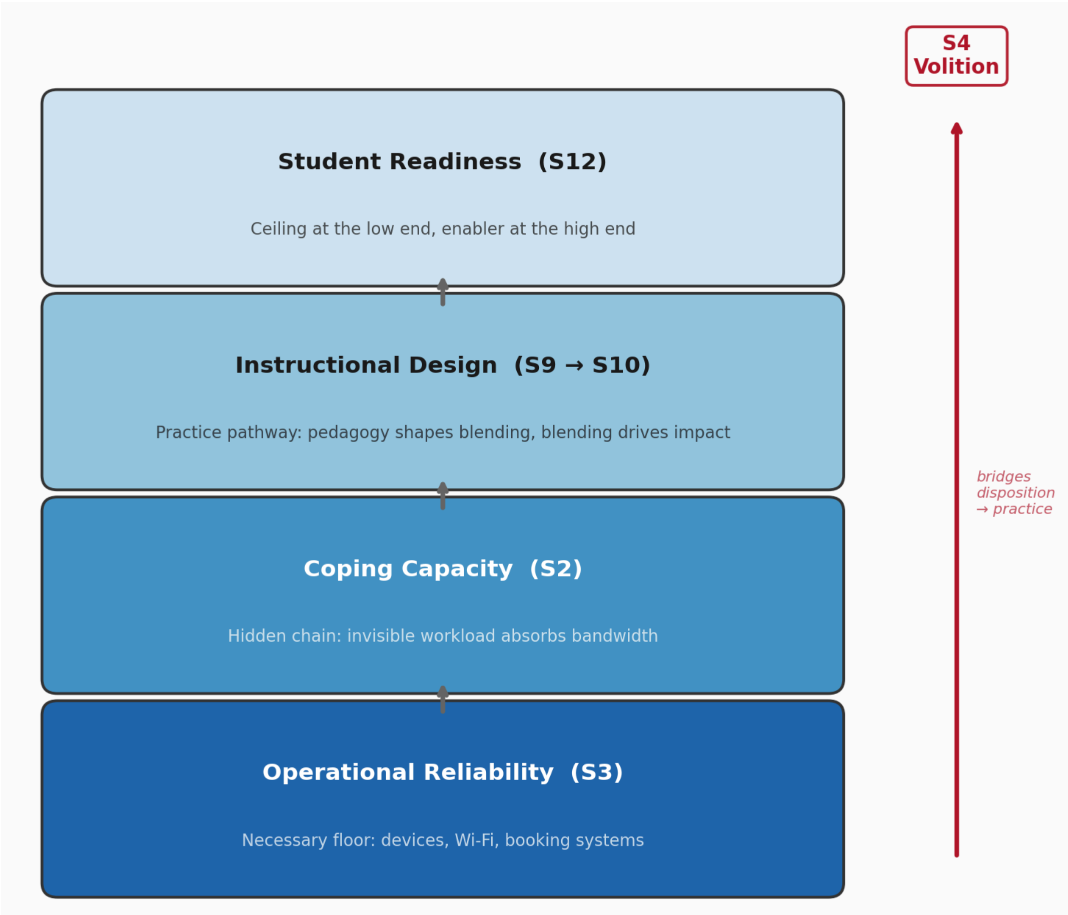

Each layer is prerequisite for the next. Each sets a ceiling on what the layer above can achieve. The question for school leaders is not whether teachers are using enough EdTech. It is which layer is currently the binding constraint.

6.2 Recommendations

The following recommendations are addressed to ICT Key Personnel (KPs) and school leaders. Each corresponds to a layer in the adoption model proposed in Section 5.10.

6.2.1 Secure the Foundation: Operational Reliability

The ICT KP roles require more than pedagogical coordination. It requires technical credibility. The KPs overseeing EdTech implementation should have sufficient technical knowledge to diagnose common failures and make informed investments into choice hardware: Wi-Fi load issues, device charging cycles, platform compatibility, and login architecture. Teachers' need for troubleshooting during lessons should be minimized proactively and not be reliant on just-in-time ("JIT") interventions by support staff. Where the KP lacks this background, the recommendation is not replacement but development. MOQ HQ and schools can invest in

- design of resilient and robust school-wide provisions with a keen consideration from the users' (teachers' in class) point of view, and
- building the KP's technical capability rather than relying on teachers to work around systemic gaps.

Systemic Pain Points beyond School-level Control

Practically, these examples remain unsolved systemically:

- School-wide daily device monitoring for battery charge levels, OS updates, and geolocation
- Flexible device-booking and device-collection processes and systems which should allow same-day, JIT access where possible.
- Wi-Fi AP capacity should be stress-tested against concurrent class usage, not single-device benchmarks. The current hard-limit of 40 devices per AP causes consistently-intermitted dropped connections due to primary schools frequently having to service more than 40 devices in a class (e.g. 42 students per class, and teachers' multiple devices)
- Reliable login password systems that adhere to the need for personalized, non-standard passwords. An example could be a system that allows each school to specify and refresh at will fixed conventions of MIMS passwords that includes passphrases and known values such as "iL0veFrontierDDMMYYY" where DDMMYYYY refers to the students' birthdate.)

6.2.2 Acknowledge and Reduce the Invisible Workload

Coping capacity cannot improve if the workload remains invisible. School leaders should posture recognition of the preparation overhead that EdTech requires. This means recognising that an EdTech lesson is not equivalent in preparation time to a non-EdTech lesson, minimally within internal staff-wide communications.

Protected time for lesson design, tool experimentation, and collegial sharing should be built into the timetable rather than left to teachers' personal time.Where structured PD time is already allocated, its focus should shift from tool demonstration to supported lesson building. Teachers learn tools by using them for real lessons, not by watching presentations about them.

6.2.3 Invest in Instructional Design, Not Tool Proliferation

The strongest pathway to student impact in the data is S9 (Pedagogy) → S10 (Blended Efficacy) → S11 (Perceived Student Impact). Pedagogical orientation shapes blended practice, which drives impact. Schools should resist the temptation to introduce new tools each term. Instead, they should deepen teachers' capacity with selected, reliable platforms.

The recommendation is to mandate outcomes, not tools. A school that requires "one formative assessment cycle per fortnight using any platform the teacher is confident with" will produce better results than one that requires "use of [specific new tool] by Term 2." Teachers who feel ownership over their tool choice score higher on S4 (Volition). Teachers who feel mandated score lower.

6.2.4 Build Student Digital Readiness Progressively

Student readiness (S12) constrains adoption at the low end and enables it at the high end. The constraint is most acute at lower primary, where typing speed, login procedures, and platform navigation consume instructional time.

The recommendation is to treat digital readiness as a taught skill, not an assumed one. Schools that invest in structured digital literacy at P1 and P2 (login routines, basic navigation, typing practice) will find the S12 constraint lifts by P4. Schools that defer this investment will face the same friction at every level, every year.

6.2.5 The Invisible Workload of ICT KPs

MOE's Digital Literacy and Technological Skills guide (2025) provides a progressive framework across P1 to P6 junctures. It specifies learning outcomes but not success criteria. Schools must define their own.

This work cannot sit with the ICT KP alone. Defining what digital readiness looks like in a P2 English composition lesson requires expert input from English teachers. But initiating that conversation, securing buy-in from department KPs, obtaining permission to "take up their lesson time" and co-creating assessment rubrics that are both construct-valid and practical falls to the ICT KP.

This is the invisible workload of the ICT KP. The study surfaced invisible workload for classroom teachers. The same dynamic operates one level up. The ICT KP absorbs the coordination cost of making digital readiness a shared responsibility rather than a siloed mandate.

Where that cost is unacknowledged, the KP either burns out or stops trying. Where it is recognised and supported with protected time and leadership backing, cross-departmental ownership becomes possible.

6.2.6 Build Volition, Not Compliance

S4 (Volition/Agency) bridges every layer in the model. It is the variable that converts adequate conditions into purposeful practice. Schools cannot mandate volition. They can create the conditions that nurture it.

Three practices support this.

1. give teachers genuine choice over which tools they use.
2. create low-stakes opportunities to experiment and fail without consequence to appraisal.
3. make successful practice visible through peer sharing rather than top-down showcases.

Teachers in this study who described positive EdTech experiences consistently referenced collegial exchange, not formal training, as the catalyst.

The goal is not universal enthusiasm. It is informed professional judgment: teachers who choose technology when it serves learning and decline it when it does not.

6.3 Limitations and Future Studies

Every study has boundaries. This section names them honestly, notes what this study already did to mitigate each one, and identifies where the next study could push further.

6.3.1 Researcher Subjectivity

Qualitative research carries interpretive risk. The researcher's own position as a practitioner-researcher in the Singapore primary school system shaped the lens through which transcripts were read and scored.

This study mitigated the risk in two ways.

1. The interview chatbot standardised the question sequence across all 61 respondents, removing interviewer-specific probing bias.
2. Scores were assigned against explicit rank descriptors documented in the analysis parameters (Appendix [X]), not against impressionistic judgment.

However, the scoring remained a single-rater process. Future studies should incorporate inter-rater reliability measures. An example of such a measure could be to study if a teacher or middle manager with a different background, code the same ranks for the same interview transcript.

A multi-rater design using the same 12-spectrum framework would test the robustness of the scoring and strengthen the findings.

6.3.2 Sample Scope

The sample comprised 61 primary school teachers across minimally 5 local government primary schools. This is substantial for qualitative research but narrow in two respects.

1. Convenience sampling means the respondents may over-represent teachers who were willing to engage with a chatbot about EdTech.
2. The study captured only the primary school context, where students do not have personal learning devices (PLDs). Secondary schools, where PLDs are standard, would alleviate the S3 (Facilitating Conditions) constraint. Replicating this study with secondary teachers would test whether the four-layer model holds when the foundational layer is largely resolved.

6.3.3 Data Collection Window

The chatbot collected data over eight days (27 March to 3 April 2026). Teacher attitudes toward EdTech are not static. They shift with term pressures, school events, and policy changes. A one-week window captures a snapshot, not a trajectory.

This study gained consistency from the tight window: all 61 respondents answered during the same period under the same institutional conditions. A longitudinal design tracking the same teachers across two or three terms would reveal whether the score-3 cluster is stable or whether it shifts in response to specific interventions.

6.3.4 Chatbot as Interview Medium

The chatbot format enabled scale (61 interviews in one week) and consistency (standardised question flow). The next methodological contribution would be to test whether adaptive sequencing (skipping questions already answered), response-length sensitivity (fewer probes for detailed answers), and stronger topic-following improve both data quality and participant experience. The chatbot-mediated ethnographic interview is a young method. This study demonstrated its viability. The next study could refine its execution.

6.3.5 From Perception to Practice

This study captured what teachers say about their EdTech experience. It did not observe what they did. Self-reported data is subject to social desirability and inaccurate self-perception. A teacher who describes meaningful blended practice may, under observation, reveal something different.

Future studies could pair chatbot interviews with classroom observation to triangulate reported experience against observed practice. The 12-spectrum framework provides a ready-made observation lens. A study that scores the same teacher from both transcript and observation would test the alignment between perception and practice, and reveal where the two diverge.

Appendix

The HTML dashboard (containing the individual descriptors for rank coding for all spectra, all pairwise correlations and all verbatim interview transcripts) as well as the prompt-engineering behind the interview chatbot can be accessed via:

http://tiny.cc/2026mlsCPT31apdx

References

Bandura, A. (1997). *Self-Efficacy: The Exercise of Control*. Worth Publishers.

Cohen, J. (1988). *Statistical Power Analysis for the Behavioral Sciences* (2nd ed.). Routledge. https://doi.org/10.4324/9780203771587

Connell, R. (2009). Good teachers on dangerous ground: towards a new view of teacher quality and professionalism. *Critical Studies in Education*, *50*(3), 213-229. https://doi.org/10.1080/17508480902998421

Ertmer, P. A. (1999). Addressing First- and Second-Order Barriers to Change: Strategies for Technology Integration. *Educational Technology Research and Development*, *47*(4), 47-61. https://doi.org/10.1007/bf02299597

Ertmer, P. A. (2005). Teacher Pedagogical Beliefs: The Final Frontier in Our Quest for Technology Integration? *Educational Technology Research and Development*, *53*(4), 25-39. https://doi.org/10.1007/BF02504683

Hong, H. (2023). *How Teacher Leadership Develops and its Effects on Information Communication Technology Enabled Reforms in Singapore Schools* [Doctoral dissertation, National Institute of Education, Nanyang Technological University]. https://hdl.handle.net/10497/26186

Natarajan, U., Lim, K. Y.T., & Laxman, K. (2021). A national vision for information and communication technologies in education: reflections on Singapore's ICT technologies Masterplans. *International Journal of Educational Management*, *35*(5), 943. Emerald Insight. https://doi.org/10.1108/IJEM-11-2020-0532

Selwyn, N., Nemorin, S., & Johnson, N. (2016). High-tech, hard work: an investigation of teachers' work in the digital age. *Learning, Media and Technology*, *42*(4), 390-405. https://doi.org/10.1080/17439884.2016.1252770

Sun, Y., & Yan, Z. (2025). Factors influencing teachers' technology adoption in technology-rich classrooms: model development and test. *Educational technology research and development*, *73*(1), 1481-1496. https://doi.org/10.1007/s11423-025-10454-5

Tan, S. H., & Ng, D. (2021, June 23). S'pore govt to spend $3.8b on ICT, including school auto-marking system. *The Straits Times*. Retrieved March 25, 2026, from https://www.straitstimes.com/tech/tech-news/spore-govt-to-spend-38b-on-ict-including-school-auto-marking-system

Wang, S. (2025). AI for social science: A sociology PhD candidate's autoethnography on how LLMs are changing research work. *AI Magazine*, *46*(e70046), na. https://doi.org/10.1002/aaai.70046