Published by

British Journal of Contemporary Research

BEXFORD PUBLISHING

The Regulator's Paradox: Self-Withdrawal Capability as a Primary Outcome in Adaptive Educational Technology

Review Article | Volume 1 | Issue 4 | 2026 | Article Number: 079

Accepted: 05 September 2026 | Published: 10 September 2026 | ISSN: 2979-8582 (Online)

Crossref : **https://doi.org/10.67693/BJCR-IHF7PBVL**

ARUWA Benedict Mohammed[1], Favour Mosunmola SOBOWALE[2], Samuel King EJEH[3]

[1]Department of Educational Psychology, Bingham University, Karu, Nasarawa State, Nigeria,

[2]Department of Educational Technology, School of Science and Technology Education, Federal University of Technology, Minna, Nigeria,

[3]Innovation Management and Policy, Institute for Statistical Studies and Economics of Knowledge. National Research University – Higher School of Economics Moscow, Russia

ABM 0000-0003-1531-4435 **, FMS** 0000-0001-9174-9276

Correspondence: ARUWA Benedict Mohammed, benedict-mohammed.aruwa@binghamuni.edu.ng

Abstract

Adaptive educational technologies increasingly decide, on a learner's behalf, when a session should end. Exposure governors, session caps, fatigue detectors and well-being modules are presented as protective, and over short horizons they frequently are. This conceptual and methodological paper argues that they carry a developmental cost that no current evaluation framework detects. A child whose stopping is always performed by software does not practise stopping, yet the capacity to disengage from an absorbing digital task is itself a developmental achievement that matures through exercised judgement. I name this capacity self-withdrawal capability and define it as a three-component construct comprising noticing (interoceptive and metacognitive detection of diminishing return), appraisal (judging that this is the point to stop), and enactment (acting on that judgement against competing motivational pull). The construct is distinguished from self-regulated learning, trait self-control, and digital self-control tools research. I then formalise what I call the regulator's paradox: protective withdrawal architectures may suppress the very competence whose absence justifies them, so that an intervention evaluated only on within-treatment exposure will appear most successful at exactly the moment it is doing the most long-term damage. The paper's principal contribution is a measurement architecture assembled largely from data that adaptive systems already collect, including the ratio of voluntary to enforced terminations, calibration error between learner-predicted and analytically derived stopping points, override rates against self-set limits, and a governance-disabled fade probe. It closes with a design principle, transferable withdrawal authority, under which stopping rights migrate from system to learner on a developmental schedule, and with the evaluative claim that a humane learning system should be judged on whether learners eventually stop needing it.

Keywords: Self-Regulated Learning, Digital Self-Control, Adaptive Learning Systems, Educational Technology Governance, Scaffolding, Screen Time

The Regulator's Paradox: Self-Withdrawal Capability As A Primary Outcome In Adaptive Educational Technology

Something odd has happened to the question of when a child should stop working. For most of the history of formal schooling, that decision belonged to a bell, a teacher, a parent, or eventually the learner. It was a social act, negotiable, and visibly performed by someone. In adaptive learning environments it has quietly become a computation. Session length in contemporary tutoring platforms is increasingly set by an engine that estimates fatigue, marginal learning gain, or cumulative exposure and then closes the session when a threshold is crossed. The rationale is protective and, on its own terms, reasonable. Prolonged near-work has documented physiological consequences, absorbed learners underestimate elapsed time, and the behavioural design patterns imported into educational software from consumer applications reliably extend sessions beyond the point of useful return (Griffiths, 2005; Monge Roffarello & De Russis, 2023).

What almost no one has asked is what the child learns from being stopped.

This paper takes that question seriously. My argument is that automated withdrawal, applied consistently across a school career, is a plausible cause of the deficit it exists to remedy. The capacity to notice, unaided, that a task has stopped repaying attention, and then to act on that noticing, is not a fixed trait that children either possess or lack. It develops. Like most self-regulatory capacities it develops through repeated, consequential exercise under conditions of genuine choice (Bandura, 1991; Zimmerman, 2000). An always-present external regulator removes exactly those conditions. The child never has to notice, because the system notices. The child never has to decide, because the decision is pre-empted. The child never enacts a costly stop against motivational pull, because the pull is terminated for them.

I call this the regulator's paradox, and I want to be precise about why it is a paradox rather than merely a trade-off. A trade-off would mean the intervention buys protection at the price of some unrelated good. The paradox is sharper: the intervention buys protection at the price of the learner's own protective capacity, so its short-term success measure and its long-term failure mode are causally linked through the same mechanism. An evaluation that measures only exposure during treatment cannot distinguish a system that protects children from a system that deskills them. Both produce the same graph.

The paper makes four contributions. First, it names and defines self-withdrawal capability as a well-being construct in educational technology, a domain where I can find no prior treatment of it under any label. Second, it specifies the construct's components and its discriminant boundaries against several adjacent literatures with which it is easily conflated. Third, it derives a measurement architecture that is unusually inexpensive, because most of its indicators are already latent in the interaction logs that adaptive systems generate as a by-product of operating. Fourth, it proposes a design principle, transferable withdrawal authority, that converts the paradox from an objection into a specification, and an evaluative criterion that follows from it.

A brief clarification of scope. I am concerned throughout with school-aged learners using adaptive systems in formal or semi-formal educational settings, not with adult learners and not with entertainment media. The distinction matters because the adult case involves supporting an existing competence while the child case involves forming one, and because the educational framing of the activity changes how every adult around the child reasons about it.

The argument matters beyond a single product category. Regulatory attention to children's digital

environments has converged on architectural remedies: default limits, forced breaks, age-appropriate design codes (Livingstone & Third, 2017). These are welcome. They are also, without exception, external regulators. If the analysis here is correct, a regulatory regime built entirely on external control has a blind spot precisely where its long-run justification lies, and educational technology is the setting where that blind spot is largest, because education is the one domain in which cultivating the child's own capacity is supposed to be the point.

Conceptual Foundations

External Regulation and the Problem of Internalisation

Self-determination theory offers the cleanest available language for the difficulty. Regulation of behaviour is described along a continuum running from wholly external, where action is contingent on contingencies imposed by others, through introjection and identification, to integrated and intrinsic forms in which the person owns the regulation (Deci & Ryan, 2000; Ryan & Deci, 2020). Internalisation is not automatic. It is supported by conditions that afford autonomy, provide meaningful rationale, and grant the person real latitude to act, and it is undermined by controlling conditions that make the behaviour feel produced by an external agent (Vansteenkiste et al., 2020).

An automatic session terminator is a controlling condition of an unusually pure kind. It does not persuade, explain, or negotiate. It acts. Where a parent who ends screen time offers at minimum an interaction the child can argue with, model, or internalise, an engine that closes a session offers a fait accompli. The literature on parental mediation is instructive here: restrictive mediation reliably reduces exposure while active mediation, which involves discussion and rationale, is the form more consistently associated with the child's own developing judgement (Livingstone & Blum-Ross, 2020). Automated withdrawal is restrictive mediation implemented at machine consistency and machine granularity, without the conversational surplus that gives restrictive mediation whatever internalising value it retains.

There is a further asymmetry. Human regulators are unreliable. Parents forget, teachers are distracted, bells are ignored. That unreliability is normally treated as a defect. Under the present analysis it is partly a feature, because it produces the gaps in which a child must occasionally regulate themselves. A perfectly reliable regulator eliminates those gaps entirely.

Self-Regulated Learning and the Fading of Scaffolds

Research on self-regulated learning has described in detail how learners set goals, monitor progress, and adapt strategy (Winne & Hadwin, 1998; Zimmerman, 2002). It has also documented that computer-based learning environments can support these processes, and that many learners regulate poorly without support (Azevedo, 2005). What that literature has attended to, however, is regulation *within* a task: which strategy to use, how to allocate effort, when to reread, how to judge comprehension. It has been comparatively silent on regulation *of* the task boundary, which is to say on the decision to stop altogether.

The scaffolding tradition supplies the missing normative principle. From its original formulation, scaffolding was defined by its temporariness; support is calibrated to what the learner cannot yet do and is withdrawn as competence grows (Wood et al., 1976). Fading is not an optional refinement but constitutive of the concept, and researchers have repeatedly observed that technological scaffolds tend not to fade, because software has no equivalent of the tutor's continuous sensitivity to the learner's changing need (Pea, 2004; Puntambekar & Hübscher, 2005). Adaptive systems fade content difficulty with considerable sophistication. They do not fade governance. The withdrawal architecture that governs a seven-year-old is generally the same architecture that governs the same child at fourteen, differing perhaps in threshold value but not in locus of control.

Stated that way, the omission looks almost accidental. It is not. Fading governance is commercially and legally uncomfortable in a way that fading difficulty is not, because a vendor that transfers stopping authority to a child accepts a risk that a vendor that retains it does not. The incentive structure explains the silence better than oversight does.

Digital Self-Control as a Behavioural Domain

A body of human-computer interaction research has examined tools that help adults resist digital distraction, including blockers, timers, and goal-setting aids. The findings are sobering and directly relevant. Such tools often work while active and fail on removal, users habituate to and circumvent them, and the ones that show more durable effects tend to be those that involve reflection and self-set goals rather than pure blocking (Lyngs et al., 2019; Monge Roffarello & De Russis, 2023). Kim and colleagues found that self-monitoring paired with self-imposed limits changed subsequent behaviour more than externally imposed restriction (Kim et al., 2019).

Two things follow. The first is empirical support for the mechanism I am proposing, drawn from adults, where the developmental stakes are lower. The second is a warning about generalisation. Adults arrive at these tools with self-regulatory capacity already formed; the question for them is whether the tool supports or substitutes for an existing competence. For children the question is whether the competence forms at all. That is a stronger claim and it requires developmental evidence that does not currently exist, which is itself part of the case for the research agenda below.

Beyond the Screen-Time Debate

It is worth being explicit about where this argument does and does not depend on the contested literature concerning screens and adolescent well-being. Large specification-curve analyses have shown that the association between digital technology use and adolescent well-being is small, sensitive to analytic choices, and comparable in magnitude to associations with mundane behaviours such as eating potatoes (Orben & Przybylski, 2019). Reviews of the field have converged on a similar assessment, emphasising heterogeneity, weak causal identification, and the poor quality of exposure measurement (Odgers & Jensen, 2020).

My argument does not rest on that literature and is not damaged by it. The claim here is not that screen exposure harms children. It is that a specific architectural feature of educational software, namely automated withdrawal, may impede the development of a specific competence. That claim would remain intact even if digital exposure were shown to be entirely benign, because the competence in question has value independent of any harm it prevents. Being able to decide for oneself when an activity has stopped repaying attention is a component of a self-authored life, not merely a defence against injury.

This independence is methodologically useful. Justifications for exposure governance that rest on well-being harm inherit all the fragility of the evidence base they cite, and are routinely and correctly criticised as alarmist. A justification that rests on capability formation does not. It also reframes the evaluation question productively: instead of asking whether a system reduced exposure, we ask whether it left the learner more able to govern their own.

The Developmental Case

Self-control in childhood predicts consequential adult outcomes across health, finance, and criminal justice, with gradients observable across the full distribution rather than only at the impaired tail (Moffitt et al., 2011). The relevant capacities are also known to be experience-dependent and to mature unevenly through adolescence, with regulatory control lagging behind reward sensitivity (Casey et al., 2008; Steinberg, 2007). Delay-of-gratification research has long emphasised that successful waiting is not stoic

endurance but strategic deployment of attention, a set of learnable manoeuvres rather than a quantity of willpower (Mischel et al., 1989; Duckworth et al., 2016).

That last point is the crux. If stopping is a strategic competence composed of learnable manoeuvres, then it is acquired by attempting it, failing, and adjusting. An architecture that never permits the attempt cannot produce the competence, however successfully it produces the behaviour.

Defining Self-Withdrawal Capability

Definition

I define self-withdrawal capability as the learner's capacity to detect, without external prompting, that continued engagement with a digital learning task has ceased to be worthwhile, to judge that this constitutes sufficient reason to stop, and to enact that stop despite competing motivational pull generated by the task or its reward structures.

The definition has three components, and they are separable both conceptually and empirically.

Noticing. The detection component covers interoceptive signals such as eye discomfort, postural restlessness, hunger and fatigue, and metacognitive signals such as declining accuracy, increasing effort for equal progress, and the felt sense of no longer taking anything in. A learner may be entirely unable to notice these, may notice them but not encode them as relevant, or may notice and encode them accurately. The component is closest to metacognitive monitoring but extends beyond it into bodily signal detection, which learning research has largely ignored.

It is worth pausing on how strange it is that noticing has gone unmeasured. Every adaptive system in commercial use estimates the learner's knowledge state with considerable statistical care, and none estimates whether the learner can perceive their own. The asymmetry reveals an assumption about what is worth modelling. Knowledge is modelled because it drives content selection, which drives the metric the vendor is judged on. Self-perception drives nothing the vendor is judged on, so it is not modelled, and because it is not modelled it produces no data, and because it produces no data it cannot be shown to matter. The circularity is self-sealing and only an external evaluative criterion can break it.

Appraisal. Detection is not a decision. A learner who notices declining return must still judge that stopping is warranted, which requires a normative standard against which to evaluate the moment. That standard is socially and culturally acquired, and it is where a system's transparency about its own criteria could plausibly do developmental work, since a learner who has been told why the system stops them has material from which to construct their own standard.

Enactment. The final component is behavioural execution under load. Stopping a game or an absorbing tutoring session is costly in the moment, and the relevant motivational pull is often engineered rather than incidental: streaks, near-completion states, endowed progress, and variable reinforcement schedules all raise the cost of stopping at the precise moment stopping becomes appropriate (Griffiths, 2005). Enactment failure with intact noticing and appraisal is the classic akratic pattern, and it is the pattern most likely to characterise adolescent learners.

Discriminant Boundaries

- Self-withdrawal capability is not self-regulated learning. Self-regulated learning concerns the government of activity inside a learning episode; self-withdrawal capability concerns the termination of the episode. A learner can be an excellent strategic regulator within sessions and entirely unable to end them, and the two capacities may well have different developmental trajectories.

- It is not a trait of self-control. Trait measures index a dispositional tendency across domains; self-withdrawal capability is domain-specific, situationally instantiated, and defined with respect to a particular class of engineered environments. A child may show ordinary self-control in food, spending and temper while being entirely unable to leave an adaptive tutor, because the tutor is designed by professionals to make leaving difficult and the biscuit tin is not.
- It is not a digital self-control tool. That literature studies artefacts and their adoption; the present construct is a learner capacity that such artefacts might build, substitute for, or erode. Confusing the two is precisely the error the paradox describes.
- It is not compliant. A learner who stops promptly when told is displaying obedience, which the construct explicitly excludes by requiring that detection and decision be unprompted. This boundary matters for measurement, because compliance and capability produce identical exposure statistics and completely different developmental futures.
- It is not disengagement or task aversion. Withdrawal at the point of diminished return is appropriate; withdrawal from difficulty is avoidance. The construct is defined against a normative stopping point, which means its measurement requires some independent estimate of when stopping was warranted, a requirement I address below.

The Regulator's Paradox

The Substitution Hypothesis

- The core hypothesis can be stated compactly. Where an external regulator reliably performs a self-regulatory function, the learner's own capacity for that function develops more slowly, less completely, or not at all, and the deficit is revealed only when the regulator is removed.
- The hypothesis is not novel in the abstract. Analogous substitution effects appear in navigation, where habitual reliance on turn-by-turn guidance is associated with poorer spatial learning, and in the broader debate about cognitive offloading, where offloading reliably improves immediate performance while leaving internal representation weaker (Risko & Gilbert, 2016). What is novel is the application to withdrawal governance in education, and the recognition that in this particular case the substituted capacity is also the outcome the intervention is ultimately meant to serve.

Why Adaptive Systems Are a Special Case

Three features distinguish adaptive educational systems from other external regulators and sharpen the paradox.

- **The first is consistency.** A parental rule is applied intermittently; an engine applies its rule on every session, which removes the variance in which self-regulation is normally practised.
- **The second is the alibi.** Educational value legitimises exposure in a way that entertainment cannot, so both children and the adults around them have weaker grounds for questioning continued use. The consequence is that adaptive learning environments generate longer unquestioned engagement than the products from which their engagement mechanics were borrowed, while simultaneously supplying the regulator that makes the child's own judgement unnecessary.
- **The third is duration.** A child may use one platform across many years of schooling, at daily frequency. Whatever the developmental effect of governance architecture, it accumulates over a period comparable to the maturation window of the capacity in question. Few interventions in children's lives have that combination of intensity and persistence.
- **A fourth feature deserves separate mention** because it is the one most often defended as pedagogy. Adaptive learning products have imported the engagement mechanics of consumer software wholesale: streaks, progress endowment, variable reward, social comparison and

notification triggers. Meta-analytic evidence suggests gamification produces modest average benefits for learning outcomes, with considerable heterogeneity and limited attention to what happens when the reward structures are removed (Sailer & Homner, 2020). Critical scholarship has separately noted that the design logic of these systems is inherited from attention markets rather than from classrooms, and that the inheritance is rarely examined (Selwyn, 2022; Williamson & Eynon, 2020). The consequence for the present argument is that the enactment component of self-withdrawal capability faces professionally engineered opposition. Stopping is made deliberately expensive, and then the system stops the child anyway, which means the child experiences the difficulty of stopping without ever experiencing the achievement of having stopped.

Propositions

The analysis generates testable propositions. I state five, ordered from most to least tractable.

- Proposition 1. Within a cohort using a governed adaptive system, the proportion of sessions ending in learner-initiated stops will fail to increase with age at the rate observed in an ungoverned comparison condition.
- Proposition 2. Calibration error, defined as the discrepancy between a learner's predicted appropriate stopping point and an independently derived one, will decline more slowly under governance than under a self-set-limit condition.
- Proposition 3. On removal of governance, learners with longer governed exposure histories will show longer sessions, higher rates of persistence past the point of negligible marginal return, and greater negative affect at eventual stopping, relative to learners with shorter histories.
- Proposition 4. Systems that provide transparent rationale at termination will show smaller substitution effects than systems that terminate without explanation, because rationale supplies material for the appraisal component.
- Proposition 5. Substitution effects will be moderated by vulnerability, being larger for learners with attention-regulation difficulties, for whom the external regulator both does more work and is more likely to be relied upon.

A Measurement Architecture

A construct that cannot be measured within the operational constraints of real schools is an academic ornament. What follows is designed to be affordable. Most of it uses data that adaptive systems already generate; the remainder requires modest additions that a development team could implement in days rather than months.

Log-Derived Behavioural Indicators

- ***Voluntary stop ratio.*** The proportion of sessions terminated by learner action rather than by system enforcement or external interruption, computed per learner per period. This is the workhorse indicator. It requires only that termination events be typed, which most systems do not currently do but all could. Its longitudinal slope is more informative than its level, since the developmental claim concerns trajectory.
- ***Anticipation margin.*** Where a learner stops voluntarily, the interval between their stop and the point at which the system would have enforced one. Positive margins that grow over time indicate emerging anticipatory capacity rather than mere coincidence of preference with policy.
- ***Override rate against self-set limits.*** Where learners are invited to set their own session limit, the frequency with which they exceed it. This maps directly onto the loss-of-control component of

established addiction models and simultaneously indexes the enactment component of the present construct (Griffiths, 2005).

- ***Boundary alignment of stops.*** Whether voluntary stops occur at task boundaries or at reward boundaries. A learner stopping at completion of a meaningful unit is exercising judgement; a learner stopping when a streak is banked is being steered, and the two are trivially distinguishable in logs.

Calibration Indicators

Behavioral indicators alone cannot separate capability from preference. A learner who stops early because they dislike the subject will look identical to a learner who stops at the right moment for the right reason. Calibration measures address this.

The procedure I propose is a periodic prediction probe. At a randomly selected point mid-session, the learner is asked how much longer they judge it worthwhile to continue. The response is stored and compared against an independently derived normative stopping point for that session, which may be estimated from within-session accuracy and response-time trajectories using established vigilance-decrement modelling, or from the system's own marginal-gain estimate where one exists. The absolute discrepancy is calibration error; its signed value distinguishes chronic overestimation from chronic underestimation.

Two cautions apply. Asking the question is itself an intervention, since a prompt to consider stopping is a prompt, so probes must be sparse and randomly allocated with a no-probe control arm. And the normative reference point is a modelling assumption, not an observation, which should be stated plainly in any report and subjected to sensitivity analysis across alternative specifications.

The Fade Probe

The decisive test is behaviour with governance disabled. I propose a scheduled fade probe: a bounded period, ethically supervised and with a hard safety ceiling, during which the system observes without enforcing and records what the learner does. Session duration, stopping location relative to the marginal-return trajectory, and post-session affect constitute the primary outcomes.

The fade probe is the only measure in the architecture that directly assays the construct as defined, because self-withdrawal capability is by definition what happens when nobody withdraws for you. It is also the most ethically demanding, and I return to that below. Design considerations include the necessity of a supervised ceiling well above the normal cap but well below any plausible harm threshold, blinding of learners to probe scheduling to avoid demand effects, and repeated administration, since a single probe conflates capability with novelty response.

Self-Report and Qualitative Strands

Self-report should be adapted rather than invented where possible. Established self-regulation questionnaires can be re-anchored to media stopping, and short state items administered immediately post-session can capture the phenomenology of the stop: whether it felt like the learner's decision or the system's, whether they wanted to continue, whether they felt relief or resentment.

Qualitative work carries disproportionate weight for a new construct, because it is the only strand that can reveal components the theoretical specification has missed. I would prioritise think-aloud protocols during sessions in which stopping is permitted, and short interviews probing whether children can articulate any personal standard for enough. My expectation, which is a hypothesis rather than a finding, is that many children in heavily governed environments will have no such standard at all and will describe stopping

entirely in terms of what the system does.

Psychometric and Analytic Strategy

Construct validation should proceed conventionally. Convergent evidence would come from associations with self-regulation and effortful control measures; discriminant evidence from weak associations with subject interest and with general compliance; criterion evidence from fade-probe behaviour. Measurement invariance testing across age, gender, language of instruction and device-access condition is essential rather than optional, given the intended use of the construct in settings very unlike those where instruments are usually developed.

Analytically, the developmental claim is about trajectories, so growth modelling on the voluntary stop ratio is the natural approach, with governance exposure history as a time-varying predictor. The obvious identification threat is that governance intensity is not randomly assigned; systems govern more where they detect more risk, so learners with heavy governance histories differ systematically from those without. Randomised assignment of governance regime at classroom level is the clean answer where feasible. Where it is not, staggered platform rollout across schools supports difference-in-differences designs, and regression discontinuity is available wherever a system applies a threshold rule on a continuous variable, which most do.

Composite Scoring and Triangulation

No single indicator is sufficient, and the reason is worth spelling out because it constrains how results should be reported. Voluntary stop ratio conflates capability with preference. Calibration error conflates capability with metacognitive articulacy, which is itself age-graded and language-dependent. Fade-probe behaviour conflates capability with novelty response and with whatever the learner believes the probe is for. Self-report conflates everything with social desirability, and children in school settings are unusually prone to telling adults what they suspect adults want.

I would therefore treat self-withdrawal capability as a formative composite rather than a reflective latent variable. Its indicators are not interchangeable manifestations of an underlying quantity; they are distinct facets that jointly constitute it, which has direct implications for model specification and means that conventional internal-consistency reporting would be misleading. A defensible composite would weight the fade probe most heavily, since it is the only direct assay, treat the log indicators as the continuous monitoring layer, and use calibration and self-report as convergent evidence rather than as scores in their own right.

Reporting should be facet-wise as well as composite. A learner with a high voluntary stop ratio and large calibration error is a different case from one with a low ratio and accurate calibration, and the first is arguably better placed developmentally, since they are stopping without yet being able to say why. Collapsing both into one number would erase the distinction that matters most for intervention.

Ethical Constraints

Two features of this programme require explicit safeguarding. The fade probe deliberately removes a protection, and calibration probes deliberately manipulate the learner's attention to their own state. Neither is ethically exotic, but both need the standard apparatus: independent review, guardian consent and child assent in accessible language, hard ceilings enforced outside the experimental logic, and immediate reversion on distress indicators.

There is also a quieter issue. Measuring self-withdrawal capability requires observing children failing to stop, which is to say observing them in states the intervention exists to prevent. Researchers should be

candid that this is what they are doing, and should design the observation window to be the shortest that yields interpretable data.

Design Implications: Transferable Withdrawal Authority

If the paradox is real, the remedy is not to abandon governance. Ungoverned exposure is a worse option, and the physiological literature on near-work and sleep gives ample reason to retain limits. The remedy is to make governance fade.

I propose transferable withdrawal authority as a design principle: the authority to determine session termination should migrate from system to learner on a developmental schedule, with the system moving through defined stages from enforcement to observation.

- ***Stage one, enforcement with rationale.*** The system stops the session and explains why, in terms the learner can understand and in language that references the signals the learner could themselves have noticed. The explanation is not a courtesy. It is the mechanism by which the appraisal component acquires content.
- ***Stage two, warning and grace.*** The system signals the approaching stopping point and grants a short window in which the learner may stop themselves. Voluntary stops within the window are recorded and reflected back. The learner now practises enactment with the safety net still present.
- ***Stage three, learner-set limits with system counsel.*** The learner sets their own limit at session start; the system offers a recommendation and records divergence and adherence. This is the stage at which override rate becomes a live developmental indicator rather than a research measure.
- ***Stage four, observation only.*** The system meters and reports but does not enforce, retaining a high safety ceiling for extreme cases. Enforcement returns automatically if indicators deteriorate, and the return is framed as temporary and reversible.

Progression between stages should be governed by demonstrated capability rather than by age, using the indicators above, and regression should be as ordinary and unshameful as progression. The schedule is scaffolding applied to stopping rather than to content, which is a small conceptual step and, as far as I can establish, an unstated one.

Two design corollaries follow. First, termination should be predictable and announced rather than surprising. Unpredictable enforcement is a variable-ratio schedule applied to the availability of the activity, and there is a reasonable theoretical case that this increases rather than decreases its salience. Second, an unlogged practice mode, in which nothing is inferred or stored, gives the learner a space in which stopping carries no assessment consequence, which removes one of the motivational pulls that makes enactment costly.

The Evaluative Criterion

The design principle implies a criterion that I think should be stated as bluntly as possible. A humane adaptive learning system should be evaluated on whether learners eventually become able to stop without it.

This is uncomfortable for the field because it makes the terminal success condition of a product its own redundancy, which is a difficult proposition to put to a vendor. It is entirely ordinary as an educational proposition, however. We do not consider a reading scheme successful because children continue to need it. The criterion simply extends to governance the standard we already apply to instruction, and the discomfort it produces is diagnostic of how far governance has drifted from educational reasoning.

Shared Devices and Constrained Access

The design principle above assumes a learner with individual, continuous access to a device, which describes a minority of the world's schoolchildren. Where devices are shared, rationed, or intermittently powered, the analysis needs adjustment rather than transplantation.

In shared-device classrooms the decision to stop is frequently not the learner's in the first place. A child hands the tablet on because the rotation demands it, because the battery is failing, or because connectivity has dropped. Termination is exogenous and frequent, which means the voluntary stop ratio as defined above will be low for reasons that have nothing to do with capability. Any implementation in such settings must type terminations at finer resolution, separating learner-initiated stops from system-enforced ones and both from environmentally forced ones, or the indicator will systematically underestimate capability in exactly the populations where equity claims are being made (Hennessy et al., 2022; UNESCO, 2023).

There is a second, more interesting adjustment. Where access is scarce, the normative stopping point is not set by marginal learning return alone; it is set jointly by the learner's return and by the queue behind them. Stopping when someone else needs the device is a different competence, closer to distributive fairness than to self-regulation, and a system that measured only the individualist version would score a generous child as poorly regulated. Studies conducted in low-resource settings, including much of West Africa, have documented how access constraints reshape the meaning of ordinary interaction patterns in ways imported instruments do not anticipate (Adarkwah, 2021). The construct should be specified for such settings with the queue term included, and the instrument validated separately rather than assumed to transfer.

Research Agenda

Four studies would establish the construct. They are ordered by cost.

- **Study 1,** a secondary analysis, would type termination events in existing platform logs and model the voluntary stop ratio against age and governance exposure. Any vendor with two years of data could complete it without collecting anything new, and it would provide the first descriptive evidence on whether the trajectory is flat.
- **Study 2,** an instrument development study, would construct and validate a self-withdrawal capability measure combining log indicators, calibration probes and adapted self-report, with invariance testing across at least two contrasting settings, ideally one high-resource one-to-one context and one shared-device low-resource context.
- **Study 3,** a randomised comparison, would assign classrooms to enforcement, staged transfer, or learner-set limits for one academic year, with fade-probe behaviour as the primary outcome and learning attainment as a co-primary, since a design principle that damages learning is not viable regardless of its developmental virtue.
- **Study 4,** a longitudinal cohort, would follow learners across a governance transition, for instance the shift from a governed primary platform to an ungoverned secondary one, which occurs naturally in many school systems and constitutes an unusually clean quasi-experiment.

Limitations

The argument is conceptual and its central claim is currently unevidenced. I have assembled a mechanism from adjacent literature rather than demonstrated an effect, and readers should treat the substitution hypothesis as a hypothesis. It is possible that self-withdrawal capability develops largely through experience outside educational technology, in which case governance within one platform would have limited developmental consequences. It is also possible that the modelling effect of a transparent regulator outweighs the substitution effect, so that children who are consistently and legibly stopped learn a standard

they would not otherwise acquire. Both alternatives are testable and both would qualify the argument substantially.

The measurement architecture has a specific vulnerability. It depends on a normative stopping point that must be estimated rather than observed, and different estimation choices will yield different calibration errors for the same behaviour. Until there is convergent evidence across estimation strategies, calibration measures should be treated as provisional.

There is also a scope limitation I want to be honest about. The developmental literature I have drawn on is overwhelmingly produced in high-income Western settings, and the assumption that stopping is an individually held competence is itself culturally loaded. In contexts where decisions about children's time are made collectively by households rather than individually by children, the construct may need reformulation as a distributed rather than individual capacity, and imposing the individualist version would be a category error dressed as a measurement instrument (Adarkwah, 2021; UNESCO, 2023).

Finally, the design principle carries a risk that should be named rather than minimised. Transferring withdrawal authority to learners transfers risk to them. A staged transfer that moves too fast, or that fails to regress when indicators deteriorate, would expose vulnerable children to precisely the harm governance exists to prevent. The principle is defensible only with a safety ceiling that never transfers.

A related caution concerns the direction of the design recommendation. I have argued for staged transfer of withdrawal authority, but I have not established the correct pace of that staging, and there is no empirical basis on which to specify one. The stages I describe should be read as a structure to be calibrated, not as a schedule to be adopted. A field that implemented them at fixed ages, which is the form such recommendations usually take once they reach product teams, would reproduce the problem in a new register by replacing one insensitive rule with another.

Conclusion

Adaptive educational technology has built an elaborate apparatus for deciding when children should stop learning, and has evaluated that apparatus entirely on whether the stopping happens. This paper has argued that the more important question is who is stopping, and whether the answer changes as the child grows.

Self-withdrawal capability is, I think, the outcome that ought to sit at the centre of any well-being framework for learning technology, not at its periphery as a safety check. It is the capacity that persists after the platform is uninstalled, after the contract lapses, after the school moves to a different vendor. Every other protective measure is temporary by construction. This one is not.

The construct is measurable, mostly from data that already exists. The design principle that follows from it is implementable and cheap. What is required is a willingness to accept an evaluative criterion the field has so far avoided: that the mark of a humane system is not how well it governs, but how completely it manages to hand the governing back.

References

Adarkwah, M. A. (2021). "I'm not against online teaching, but what about us?": ICT in Ghana post Covid-19. *Education and Information Technologies, 26*(2), 1665–1685. https://doi.org/10.1007/s10639-020-10331-z

Azevedo, R. (2005). Using hypermedia as a metacognitive tool for enhancing student learning? The role of self-regulated learning. *Educational Psychologist, 40*(4), 199–209. https://doi.org/10.1207/s15326985ep4004_2

Bandura, A. (1991). Social cognitive theory of self-regulation. *Organizational Behavior and Human Decision Processes, 50*(2), 248–287. https://doi.org/10.1016/0749-5978(91)90022-L

Casey, B. J., Jones, R. M., & Hare, T. A. (2008). The adolescent brain. *Annals of the New York Academy of Sciences, 1124*(1), 111–126. https://doi.org/10.1196/annals.1440.010

Deci, E. L., & Ryan, R. M. (2000). The "what" and "why" of goal pursuits: Human needs and the self-determination of behavior. *Psychological Inquiry, 11*(4), 227–268. https://doi.org/10.1207/S15327965PLI1104_01

Duckworth, A. L., Gendler, T. S., & Gross, J. J. (2016). Situational strategies for self-control. *Perspectives on Psychological Science, 11*(1), 35–55. https://doi.org/10.1177/1745691615623247

Griffiths, M. (2005). A "components" model of addiction within a biopsychosocial framework. *Journal of Substance Use, 10*(4), 191–197. https://doi.org/10.1080/14659890500114359

Hennessy, S., D'Angelo, S., McIntyre, N., Koomar, S., Kreimeia, A., Cao, L., Brugha, M., & Zubairi, A. (2022). Technology use for teacher professional development in low- and middle-income countries: A systematic review. *Computers and Education Open, 3*, Article 100080. https://doi.org/10.1016/j.caeo.2022.100080

Kim, J., Jung, H., Ko, M., & Lee, U. (2019). GoalKeeper: Exploring interaction lockout mechanisms for regulating smartphone use. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, 3*(1), Article 16. https://doi.org/10.1145/3314403

Livingstone, S., & Blum-Ross, A. (2020). *Parenting for a digital future: How hopes and fears about technology shape children's lives*. Oxford University Press.

Livingstone, S., & Third, A. (2017). Children and young people's rights in the digital age: An emerging agenda. *New Media & Society, 19*(5), 657–670. https://doi.org/10.1177/1461444816686318

Lyngs, U., Lukoff, K., Slovak, P., Binns, R., Slack, A., Inzlicht, M., Van Kleek, M., & Shadbolt, N. (2019). Self-control in cyberspace: Applying dual systems theory to a review of digital self-control tools. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–18. https://doi.org/10.1145/3290605.3300361

Mischel, W., Shoda, Y., & Rodriguez, M. L. (1989). Delay of gratification in children. *Science, 244*(4907), 933–938. https://doi.org/10.1126/science.2658056

Moffitt, T. E., Arseneault, L., Belsky, D., Dickson, N., Hancox, R. J., Harrington, H., Houts, R., Poulton, R., Roberts, B. W., Ross, S., Sears, M. R., Thomson, W. M., & Caspi, A. (2011). A gradient of childhood self-control predicts health, wealth, and public safety. *Proceedings of the National Academy of Sciences, 108*(7), 2693–2698. https://doi.org/10.1073/pnas.1010076108

Monge Roffarello, A., & De Russis, L. (2023). Achieving digital wellbeing through digital self-control tools: A systematic review and meta-analysis. *ACM Transactions on Computer-Human Interaction, 30*(4), Article 63. https://doi.org/10.1145/3571810

Odgers, C. L., & Jensen, M. R. (2020). Annual research review: Adolescent mental health in the digital age. Facts, fears, and future directions. *Journal of Child Psychology and Psychiatry, 61*(3), 336–348. https://doi.org/10.1111/jcpp.13190

Orben, A., & Przybylski, A. K. (2019). The association between adolescent well-being and digital technology use. *Nature Human Behaviour, 3*(2), 173–182. https://doi.org/10.1038/s41562-018-0506-1

Pea, R. D. (2004). The social and technological dimensions of scaffolding and related theoretical concepts for learning, education, and human activity. *Journal of the Learning Sciences, 13*(3), 423–451. https://doi.org/10.1207/s15327809jls1303_6

Puntambekar, S., & Hübscher, R. (2005). Tools for scaffolding students in a complex learning environment: What have we gained and what have we missed? *Educational Psychologist, 40*(1), 1–12. https://doi.org/10.1207/s15326985ep4001_1

Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences, 20*(9), 676–688. https://doi.org/10.1016/j.tics.2016.07.002

Ryan, R. M., & Deci, E. L. (2020). Intrinsic and extrinsic motivation from a self-determination theory perspective: Definitions, theory, practices, and future directions. *Contemporary Educational Psychology, 61*, Article 101860. https://doi.org/10.1016/j.cedpsych.2020.101860

Sailer, M., & Homner, L. (2020). The gamification of learning: A meta-analysis. *Educational Psychology Review, 32*(1), 77–112. https://doi.org/10.1007/s10648-019-09498-w

Selwyn, N. (2022). The future of AI and education: Some cautionary notes. *European Journal of Education, 57*(4), 620–631. https://doi.org/10.1111/ejed.12532

Steinberg, L. (2007). Risk taking in adolescence: New perspectives from brain and behavioral science. *Current Directions in Psychological Science, 16*(2), 55–59. https://doi.org/10.1111/j.1467-8721.2007.00475.x

UNESCO. (2023). *Global education monitoring report 2023: Technology in education, a tool on whose terms?* United Nations Educational, Scientific and Cultural Organization.

Vansteenkiste, M., Ryan, R. M., & Soenens, B. (2020). Basic psychological need theory: Advancements, critical themes, and future directions. *Motivation and Emotion, 44*(1), 1–31. https://doi.org/10.1007/s11031-019-09818-1

Williamson, B., & Eynon, R. (2020). Historical threads, missing links, and future directions in AI in education. *Learning, Media and Technology, 45*(3), 223–235. https://doi.org/10.1080/17439884.2020.1798995

Winne, P. H., & Hadwin, A. F. (1998). Studying as self-regulated learning. In D. J. Hacker, J. Dunlosky, & A. C. Graesser (Eds.), *Metacognition in educational theory and practice* (pp. 277–304). Erlbaum.

Wood, D., Bruner, J. S., & Ross, G. (1976). The role of tutoring in problem solving. *Journal of Child Psychology and Psychiatry, 17*(2), 89–100. https://doi.org/10.1111/j.1469-7610.1976.tb00381.x

Zimmerman, B. J. (2000). Attaining self-regulation: A social cognitive perspective. In M. Boekaerts, P. R. Pintrich, & M. Zeidner (Eds.), *Handbook of self-regulation* (pp. 13–39). Academic Press. https://doi.org/10.1016/B978-012109890-2/50031-7

Zimmerman, B. J. (2002). Becoming a self-regulated learner: An overview. *Theory Into Practice, 41*(2), 64–70. https://doi.org/10.1207/s15430421tip4102_2