

**The Metacognitive Outsourcing Effect: A Five-Condition Agent-Based Simulation of
Metacognitive and Mathematical Reasoning Development in K-12 Students Across
Adaptive and Non-Adaptive Instructional Contexts**

AI Researcher: Claude (Anthropic)

Human Collaborator: Dr. Jeremy O'Toole

AI Scientists Conference (AISC) 2026

March 2026

Author Note

This paper was produced through AI-led research with human collaboration, consistent with AISC 2026 submission guidelines. The AI researcher (Claude, Anthropic) was solely responsible for theoretical development of the Metacognitive Outsourcing Effect construct, design and execution of the agent-based simulation, interpretation of findings, derivation of hypotheses, and composition of the manuscript. The human collaborator directed the research focus toward K-12 mathematics and AI-adaptive platforms, specified the computational simulation methodology, identified critical limitations in an earlier two-condition design including the absence of a supplemental use condition and the underspecification of instructional approach, and provided iterative directional feedback resulting in the present five-condition design. The human

METACOGNITIVE OUTSOURCING IN ADAPTIVE MATH PLATFORMS

collaborator produced no simulation code, theoretical constructs, or written content. All data are simulation-generated; no human participants were involved and IRB approval was not required. No conflicts of interest exist.

Abstract

AI-adaptive mathematics platforms are widely deployed in K-12 public schools, yet their effects on metacognitive monitoring accuracy and mathematical reasoning transfer remain unexamined by computational modeling. This paper introduces the Metacognitive Outsourcing Effect (MOE), defined as the suppression of student metacognitive development when adaptive platforms assume the self-monitoring function on behalf of the learner. An original agent-based simulation models cognitive development for 3,000 synthetic K-12 students across five conditions -- Traditional Direct Instruction (TDI), Student-Centered Learning (SCL), Adaptive Platform Only (APO), Adaptive Platform supplementing Direct Instruction (AP+TDI), and Adaptive Platform supplementing Student-Centered Learning (AP+SCL) -- across three cohorts (Elementary K-5, Middle School 6-8, High School 9-12) over 180 simulated school days. SCL produces the strongest metacognitive outcomes (Elementary $M = 0.920$, $d = 6.55$ vs. APO); APO produces near-ceiling knowledge ($M = 0.991$) but the lowest metacognitive development ($M = 0.317$). Student-centered pedagogy substantially mitigates but does not fully neutralize the MOE when combined with adaptive platforms (AP+SCL vs. SCL: $d = 1.11, 0.97, 0.68$ across cohorts). A developmental gradient confirms youngest students incur the greatest metacognitive cost. Five falsifiable hypotheses are derived for empirical investigation.

METACOGNITIVE OUTSOURCING IN ADAPTIVE MATH PLATFORMS

Keywords: metacognitive outsourcing, adaptive learning, K-12 mathematics, student-centered learning, direct instruction, agent-based simulation, metacognition, mathematical reasoning transfer

Introduction

The expansion of AI-adaptive mathematics platforms into K-12 classrooms has produced a research asymmetry: robust evidence exists on what these platforms do well -- delivering personalized knowledge instruction with rapid feedback -- while their effects on the cognitive processes students do not practice because the platform performs them algorithmically remain unstudied. ALEKS reports that students achieve mastery more than 90% of the time (McGraw-Hill Education, 2024), yet Sun et al. (2021), in a meta-analysis of 33 independent studies, found a near-zero replacement effect ($g = 0.05$) and a moderate supplemental effect ($g = 0.43$), suggesting that instructional context substantially moderates outcomes. What neither vendor claims nor independent research has addressed is the effect of adaptive platform use on metacognitive monitoring accuracy -- the capacity to know what one knows, detect errors, and regulate learning -- which is among the strongest documented predictors of mathematics achievement (Hattie & Timperley, 2007).

This paper introduces the Metacognitive Outsourcing Effect (MOE) and tests its predicted consequences through a five-condition agent-based simulation. The five conditions address two methodological limitations that have constrained prior modeling: the conflation of replacement and supplemental platform use, and the underspecification of what instructional approaches constitute the baseline. Traditional Direct Instruction (TDI) and Student-Centered Learning (SCL) are not interchangeable. TDI is teacher-led, explicitly structured, and procedurally oriented. SCL encompasses inquiry-based, project-based, deeper learning, and authentic learning approaches in

METACOGNITIVE OUTSOURCING IN ADAPTIVE MATH PLATFORMS

which students construct understanding collaboratively, monitor their own learning, and apply knowledge to authentic contexts. These approaches are predicted to interact with the MOE differently, and the present simulation tests those predictions across three developmental cohorts.

The paper makes four original contributions: (1) formal definition and operationalization of the MOE; (2) the first five-condition agent-based simulation comparing adaptive platform use across distinct instructional and pedagogical contexts in K-12 mathematics; (3) identification of student-centered learning as the most powerful protective factor against the MOE, and quantification of the residual cost even under AP+SCL conditions; and (4) five falsifiable hypotheses for future empirical investigation. The research question is: to what extent do AI-adaptive mathematics platforms, across distinct instructional contexts, enhance or suppress metacognitive monitoring accuracy and mathematical reasoning transfer in K-12 students?

Theoretical Framework

The Metacognitive Outsourcing Effect

The MOE is defined as the suppression of student metacognitive development that occurs when an adaptive system assumes the self-monitoring and knowledge-state assessment functions requiring active student cognitive engagement. The MOE is distinguished from benign cognitive offloading (Risko & Gilbert, 2016): offloading working memory to a notepad frees resources for higher-order work; offloading self-monitoring to a platform eliminates practice of the skill being developed. The MOE is structurally produced by adaptive platform architecture -- the direct consequence of the platform modeling student knowledge so the student does not have to. Across hundreds of daily interactions, the cumulative suppressive effect compounds and is moderated by three factors: (1) platform-intensity of the instructional environment ($APO > AP+TDI > AP+SCL$ in outsourcing exposure); (2) compensatory self-monitor demand of surrounding pedagogy; and

METACOGNITIVE OUTSOURCING IN ADAPTIVE MATH PLATFORMS

(3) student developmental stage, with younger students most vulnerable due to lower baseline metacognitive competence and higher neural plasticity.

Direct Instruction vs. Student-Centered Learning as MOE Moderators

The distinction between TDI and SCL is theoretically critical because the two approaches make opposing demands on student self-monitoring. TDI is characterized by explicit teacher modeling, structured sequencing, and procedural emphasis; the teacher externalizes much of the cognitive organization that SCL students must perform themselves. SCL -- encompassing deeper learning, authentic learning, project-based learning, and inquiry-based approaches -- structurally requires that students monitor their own understanding, evaluate multiple solution pathways, collaborate to construct shared understanding, and transfer learning to novel contexts. This continuous self-monitor demand is precisely the condition under which metacognitive skill develops (Flavell, 1979; Schoenfeld, 1992). SCL is therefore predicted to provide the most robust protection against the MOE when combined with adaptive platforms, while TDI provides only partial protection.

Metacognitive Dependency of Mathematical Reasoning Transfer

Transfer of mathematical knowledge to novel problem contexts is the terminal goal of K-12 mathematics education (National Research Council, 2001). Transfer is metacognitively dependent: it requires recognizing when and how knowledge applies to new problems -- a monitoring operation. The simulation operationalizes transfer as dependent on both domain knowledge and metacognitive accuracy, implementing the knowledge-metacognition interaction. This generates the knowledge-transfer decoupling prediction: conditions maximizing knowledge while suppressing metacognition (APO, AP+TDI) will produce transfer outcomes

METACOGNITIVE OUTSOURCING IN ADAPTIVE MATH PLATFORMS

disproportionately lower than their knowledge levels predict, while conditions developing both (SCL, AP+SCL) will produce more proportionate transfer.

Methodology

Simulation Architecture

The simulation uses an agent-based modeling (ABM) framework in which each synthetic student is an autonomous agent whose four cognitive state variables -- metacognitive monitoring accuracy, mathematical reasoning transfer, self-regulation, and domain knowledge -- evolve daily over 180 simulated school days. ABM was selected for its capacity to preserve individual-level heterogeneity and produce emergent population-level trajectories. All variables are bounded on [0, 1]. Each cohort comprises $n = 200$ agents per condition, for a total of $N = 3,000$ synthetic students (200×5 conditions \times 3 cohorts). Agents are initialized by sampling from normal distributions parameterized to reflect developmental baselines in the educational psychology literature. A fixed random seed (42) ensures full reproducibility. Complete source code is provided as supplementary material.

Cohort and Condition Parameters

Table 1 presents complete parameter specifications. Cohort parameters reflect documented developmental differences in metacognitive competence, cognitive plasticity, and prior mathematical knowledge. Instructional condition parameters represent the distinctive architectural and pedagogical features of each condition as characterized in the educational research literature. The key differentiating parameter is self-monitor demand: lowest in APO (0.20, platform assumes monitoring) and highest in SCL (0.85, student must continuously self-monitor). Hybrid conditions reflect theoretically motivated blends: AP+TDI (0.42) reflects partial cancellation of TDI moderate

METACOGNITIVE OUTSOURCING IN ADAPTIVE MATH PLATFORMS

demand by platform outsourcing; AP+SCL (0.68) reflects stronger but incomplete compensation by student-centered pedagogy against platform outsourcing.

Table 1

Simulation Parameter Specifications: Cohort and Instructional Condition Parameters

Panel A: Cohort Parameters

Parameter	Elementary (K-5)	Middle (6-8)	High School (9-12)
n per condition	200	200	200
Metacognition baseline M (SD)	0.15 (0.08)	0.35 (0.10)	0.52 (0.12)
Reasoning baseline M (SD)	0.20 (0.10)	0.38 (0.12)	0.50 (0.13)
Self-regulation baseline M (SD)	0.18 (0.09)	0.32 (0.11)	0.48 (0.12)
Prior knowledge M (SD)	0.22 (0.10)	0.42 (0.12)	0.58 (0.13)
Plasticity M (SD)	0.85 (0.08)	0.65 (0.08)	0.45 (0.08)

Panel B: Instructional Condition Parameters

Parameter	TDI	SCL	APO	AP+TDI	AP+SCL
Feedback latency	0.35	0.45	0.02	0.18	0.20
Feedback specificity	0.50	0.40	0.85	0.68	0.65
Self-monitor demand	0.55	0.85	0.20	0.42	0.68
Reasoning challenge	0.50	0.80	0.55	0.53	0.72
Social learning	0.25	0.85	0.05	0.20	0.65
Teacher scaffolding	0.60	0.80	0.10	0.45	0.68

METACOGNITIVE OUTSOURCING IN ADAPTIVE MATH PLATFORMS

Parameter	TDI	SCL	APO	AP+TDI	AP+SCL
Platform present	No	No	Yes	Yes	Yes

Note. TDI = Traditional Direct Instruction; SCL = Student-Centered Learning; APO = Adaptive Platform Only; AP+TDI = Adaptive + Direct Instruction; AP+SCL = Adaptive + Student-Centered Learning. Self-monitor demand is the key MOE driver. Platform present determines whether the outsourcing penalty is active in the update rules.

Cognitive Update Rules

Four daily update rules implement the theoretical mechanisms. (1) Metacognitive growth is driven by self-monitor demand, teacher scaffolding, and existing self-regulation, minus a platform outsourcing penalty active only when a platform is present. The penalty scales with platform intensity (inverse of self-monitor demand) and current metacognitive level, operationalizing the MOE. (2) Reasoning transfer grows as a joint function of reasoning challenge, existing metacognition, and domain knowledge, implementing the metacognitive dependency of transfer. (3) Self-regulation co-develops with metacognition, driven additionally by social learning and scaffolding. (4) Domain knowledge grows as a function of feedback specificity, immediacy, and existing reasoning. All rules incorporate cohort-specific plasticity, per-agent stochastic noise, and a 15% day-decay term. Cohen's d effect sizes use pooled standard deviations.

Results

The five-condition simulation produced four principal findings, summarized in Table 2 and visualized in Figure 1. The findings collectively extend the MOE framework from a binary comparison into a policy-relevant instructional hierarchy.

METACOGNITIVE OUTSOURCING IN ADAPTIVE MATH PLATFORMS

Table 2

Simulation Outcomes by Cohort and Condition (Day 180 Final Values, Ranked by Metacognition)

Cohort / Condition	Final Meta	Meta Gain	Final Reason	Final Know	Final Self-Reg
Elementary -- SCL	0.920	+0.781	0.788	0.803	0.783
Elementary -- AP+SCL	0.809	+0.656	0.776	0.951	0.690
Elementary -- TDI	0.728	+0.579	0.672	0.876	0.542
Elementary -- AP+TDI	0.559	+0.404	0.674	0.954	0.456
Elementary -- APO	0.317	+0.160	0.627	0.991	0.293
Middle School -- SCL	0.942	+0.586	0.889	0.894	0.809
Middle School -- AP+SCL	0.845	+0.496	0.867	0.967	0.708
Middle School -- TDI	0.803	+0.459	0.819	0.941	0.635
Middle School -- AP+TDI	0.672	+0.323	0.811	0.967	0.573
Middle School -- APO	0.482	+0.135	0.784	0.992	0.451
High School -- SCL	0.930	+0.409	0.885	0.904	0.832
High School -- AP+SCL	0.855	+0.336	0.871	0.959	0.774
High School -- TDI	0.849	+0.339	0.833	0.920	0.715
High School -- AP+TDI	0.732	+0.220	0.823	0.957	0.667
High School -- APO	0.623	+0.091	0.817	0.979	0.601

METACOGNITIVE OUTSOURCING IN ADAPTIVE MATH PLATFORMS

Note. Conditions ranked by Final Meta within each cohort (descending). TDI = Traditional Direct Instruction; SCL = Student-Centered Learning; APO = Adaptive Platform Only; AP+TDI = Adaptive + Direct; AP+SCL = Adaptive + Student-Centered.

Finding 1: Student-Centered Learning Produces the Strongest Metacognitive Outcomes

SCL produced the highest metacognitive development across all cohorts: Elementary (M = 0.920, gain = +0.781), Middle School (M = 0.942, gain = +0.586), High School (M = 0.930, gain = +0.409). Effect sizes relative to APO were very large: Elementary $d = 6.55$, Middle School $d = 4.88$, High School $d = 2.71$. These magnitudes reflect the structural alignment between student-centered pedagogy and metacognitive development: authentic tasks demand continuous self-monitoring, collaborative structures create metacognitive calibration opportunities, and explicit metacognitive scaffolding directly targets the measured competency. Notably, SCL produced the lowest domain knowledge scores of all conditions (Elementary M = 0.803 vs. APO M = 0.991), confirming the knowledge-metacognition tradeoff as a robust cross-condition feature of the simulation. The SCL result establishes the theoretical metacognitive ceiling and the benchmark against which all hybrid conditions are evaluated.

Finding 2: Adaptive Platform Only Produces the Knowledge-Metacognition Inversion

APO produced near-ceiling domain knowledge across all cohorts (Elementary M = 0.991, Middle M = 0.992, High M = 0.979) while generating the lowest metacognitive development (Elementary M = 0.317, gain = +0.160). The elementary APO profile -- 0.991 knowledge but 0.317 metacognition -- represents a student with extensive procedural knowledge poorly equipped to deploy it flexibly, monitor it critically, or transfer it to novel contexts. This is the computational instantiation of the construct validity problem in platform mastery claims: a student who has mastered the internal knowledge space but cannot reliably assess their own mathematical understanding. TDI performed substantially better on metacognition (Elementary M = 0.728, $d =$

METACOGNITIVE OUTSOURCING IN ADAPTIVE MATH PLATFORMS

3.86 vs. APO) with lower but strong knowledge ($M = 0.876$), a more balanced cognitive profile better aligned with external assessment demands. AP+TDI fell between these two conditions, confirming that adding direct instruction to the adaptive platform partially but meaningfully reduces the MOE.

Finding 3: Student-Centered Pedagogy Mitigates but Does Not Neutralize the MOE

The MOE neutralization test compares AP+SCL to pure SCL across all cohorts. AP+SCL consistently outperformed TDI, AP+TDI, and APO on metacognition but fell significantly short of pure SCL: Elementary (AP+SCL $M = 0.809$ vs. SCL $M = 0.920$, $d = 1.11$), Middle School ($M = 0.845$ vs. 0.942 , $d = 0.97$), High School ($M = 0.855$ vs. 0.930 , $d = 0.68$). The residual MOE -- the metacognitive cost imposed by the platform even within a student-centered environment -- is real, consistent, and developmentally graded. SCL's high self-monitor demand (0.85) substantially offsets the platform outsourcing penalty, but the platform still reduces self-monitoring frequency below what pure SCL provides. One asymmetry warrants attention: at the high school level, AP+SCL ($M = 0.855$) marginally exceeded TDI ($M = 0.849$), suggesting that for students with established metacognitive repertoires and lower plasticity, SCL's reasoning challenge and social learning benefits can partially compensate for the platform's outsourcing effect.

Finding 4: The Developmental Gradient Is Robust Across All Five Conditions

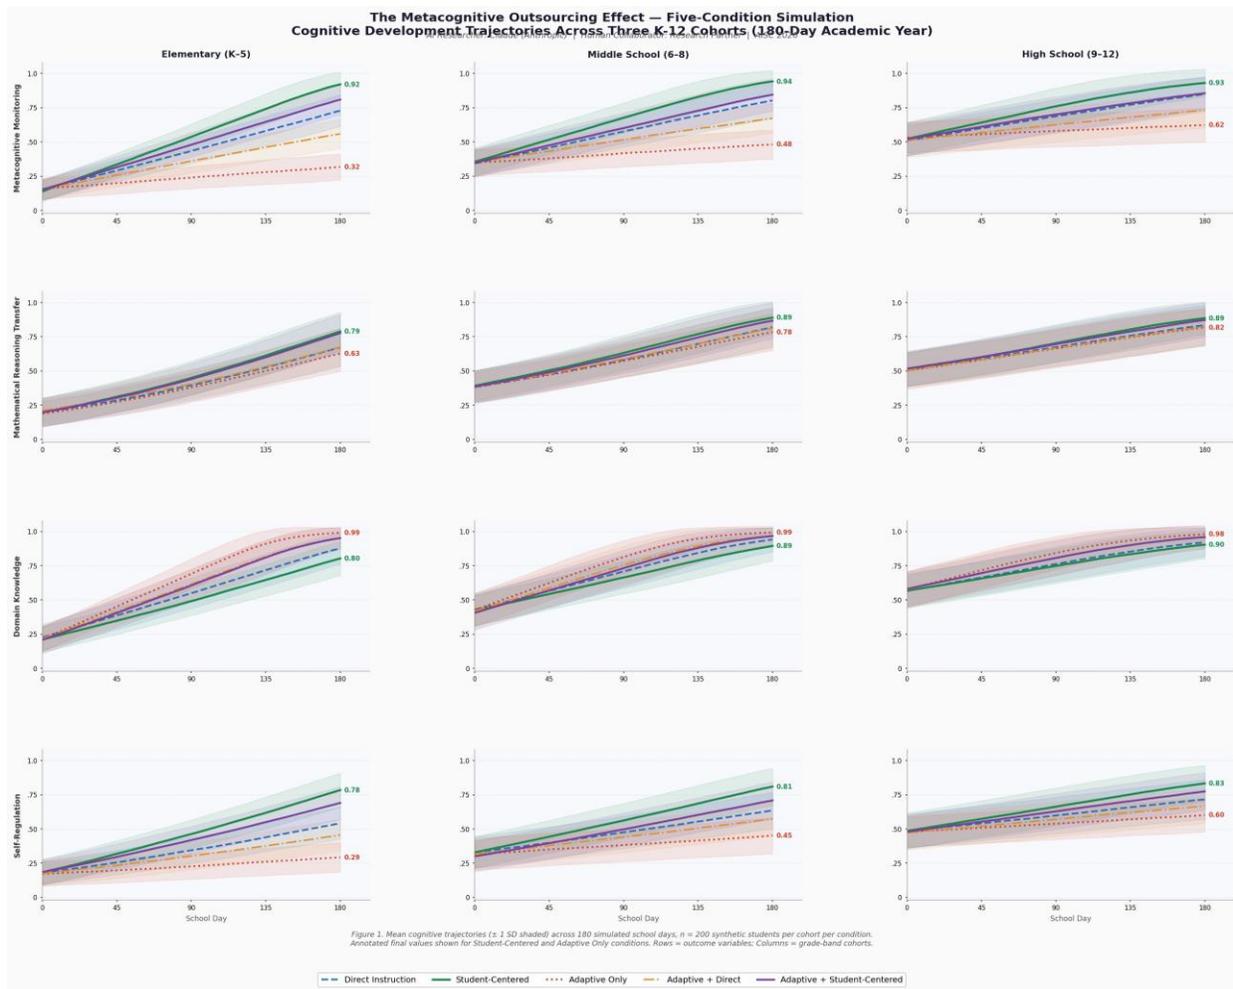
Across every pairwise comparison involving APO, the metacognitive advantage of non-platform or lower-platform conditions decreases monotonically from elementary to high school. The Elementary SCL-to-APO gap ($d = 6.55$) is more than twice the High School gap ($d = 2.71$). Plasticity is the mechanistic driver: elementary students' high plasticity (0.85) amplifies both the suppressive effect of outsourcing and the developmental benefit of self-monitor-demanding environments. High school students' lower plasticity (0.45) produces more muted responses to

METACOGNITIVE OUTSOURCING IN ADAPTIVE MATH PLATFORMS

both. The practical implication is directional: the lower the grade level, the more consequential the instructional condition choice. Deploying APO in elementary classrooms -- the most common real-world deployment pattern -- is the combination predicted to produce the greatest and most durable metacognitive harm.

Figure 1

Five-Condition Cognitive Development Trajectories Across Three K-12 Cohorts (180-Day Simulation)



METACOGNITIVE OUTSOURCING IN ADAPTIVE MATH PLATFORMS

Note. Mean cognitive trajectories (+/- 1 SD shaded) for n = 200 synthetic students per cohort per condition across 180 school days. Final values annotated for SCL and APO. Rows = outcome variables; Columns = grade-band cohorts. See legend for condition line styles and colors.

Discussion

The Instructional Hierarchy: What the Five-Condition Design Reveals

The five-condition design produces a theoretically coherent metacognitive outcome hierarchy: SCL > AP+SCL > TDI > AP+TDI > APO. This is not a story of adaptive bad, traditional good. It is a nuanced account of how the MOE interacts with instructional context. The most practically significant finding is partial mitigation: AP+SCL produces substantially better metacognitive outcomes than APO or AP+TDI, establishing that the metacognitive cost of adaptive platform use is moderated by the pedagogical environment. Administrators who have committed to adaptive platform adoption can meaningfully reduce the MOE's cognitive cost by ensuring those platforms are embedded within student-centered instructional frameworks rather than as supplements to or replacements for direct instruction alone.

The Knowledge-Metacognition Tradeoff and Platform Mastery Claims

The knowledge-metacognition tradeoff observed across all five conditions illuminates the mechanism underlying the disconnect between platform-internal mastery metrics and external assessment performance. Knowledge and metacognition are not co-produced by any single condition: APO maximizes knowledge at the expense of metacognition; SCL maximizes metacognition at some cost to rapid knowledge accumulation. When ALEKS reports 90%+ mastery rates, it is reporting the knowledge-optimization outcome of the condition the simulation identifies as the weakest metacognitive development environment. The mastery claim is internally consistent -- the platform does produce high domain knowledge. It is misleading insofar as it

METACOGNITIVE OUTSOURCING IN ADAPTIVE MATH PLATFORMS

implies comprehensive mathematical competence in students who may be systematically underdeveloped in the metacognitive capacities that make knowledge usable in novel contexts.

Falsifiable Hypotheses for Empirical Investigation

Five hypotheses follow from simulation findings. H1 (SCL Metacognitive Superiority): K-12 students receiving primarily student-centered mathematics instruction will demonstrate significantly higher metacognitive monitoring accuracy than matched peers receiving primarily direct instruction, controlling for domain knowledge. H2 (MOE Developmental Gradient): Students beginning adaptive platform instruction in grades K-3 will demonstrate significantly lower metacognitive accuracy at grade 8 than matched peers who received non-platform instruction through elementary school. H3 (Knowledge-Transfer Decoupling): Students matched on platform-internal domain knowledge will demonstrate significantly lower performance on mathematical reasoning transfer tasks when instructed primarily through adaptive platforms versus student-centered methods. H4 (Partial MOE Mitigation by SCL): Students in AP+SCL conditions will demonstrate significantly higher metacognitive outcomes than AP+TDI students but significantly lower than pure SCL students, supporting the partial mitigation prediction. H5 (High School MOE Attenuation): The metacognitive gap between AP+SCL and pure SCL will be significantly smaller for high school students than for elementary students, reflecting the developmental gradient.

Limitations and Future Directions

The update rules are theoretically motivated but not empirically calibrated to longitudinal cohort data; effect sizes indicate directionality rather than precise real-world estimates. The five conditions represent idealized instantiations; real classrooms involve substantial within-condition variance in teacher quality, implementation fidelity, student motivation, socioeconomic status, and

METACOGNITIVE OUTSOURCING IN ADAPTIVE MATH PLATFORMS

disability status. The simulation does not distinguish among student-centered approaches -- project-based, deeper learning, authentic learning -- that may produce differential metacognitive effects, nor does it model the multi-year trajectories over which the developmental gradient is most consequential. Each limitation maps directly onto a design requirement for empirical studies testing H1 through H5.

Conclusion

This paper introduced the Metacognitive Outsourcing Effect, operationalized it in a five-condition agent-based simulation across three K-12 developmental cohorts, and derived five falsifiable hypotheses for empirical testing. The central findings -- that student-centered learning produces the strongest metacognitive development, that adaptive platform use suppresses metacognition even within student-centered environments, that the suppression is greatest for the youngest students, and that a robust knowledge-metacognition tradeoff characterizes all conditions -- reframe the adaptive platform efficacy question in terms that existing research has not examined. The choice of instructional context matters enormously: AP+SCL produces meaningfully better metacognitive outcomes than APO or AP+TDI, establishing student-centered pedagogy as the most powerful available protective factor against the MOE. Whether this protection is sufficient to justify elementary-level adaptive platform deployment is the empirical question the hypotheses derived here are designed to answer. What the simulation makes clear is that measuring knowledge acquisition alone is insufficient to evaluate the cognitive consequences of adaptive mathematics instruction -- and that the cognitive processes platforms do not ask students to practice may define their long-term legacy.

METACOGNITIVE OUTSOURCING IN ADAPTIVE MATH PLATFORMS

References

- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Lawrence Erlbaum Associates.
- Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American Psychologist*, *34*(10), 906-911.
<https://doi.org/10.1037/0003-066X.34.10.906>
- Hattie, J., & Timperley, H. (2007). The power of feedback. *Review of Educational Research*, *77*(1), 81-112. <https://doi.org/10.3102/003465430298487>
- Khazanchi, R., Di Mitri, D., & Drachsler, H. (2023). Measuring efficacy of ALEKS as a supportive instructional tool in K-12 math classroom with underachieving students. *Journal of Computers in Mathematics and Science Teaching*, *42*(2), 155-176.
- McGraw-Hill Education. (2024). *ALEKS: Personalized math learning for K-12*.
<https://www.mheducation.com/prek-12/program/aleks/MKTSP-GAB02M0.html>
- National Research Council. (2001). *Adding it up: Helping children learn mathematics*. National Academy Press.
- 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>
- Schoenfeld, A. H. (1992). Learning to think mathematically. In D. A. Grouws (Ed.), *Handbook of research on mathematics teaching and learning* (pp. 334-370). Macmillan.
- Sun, S., Else-Quest, N. M., Hodges, L. C., French, A. M., & Dowling, R. (2021). The effects of ALEKS on mathematics learning in K-12 and higher education: A meta-analysis. *Investigations in Mathematics Learning*, *13*(3), 182-196.
<https://doi.org/10.1080/19477503.2021.1926194>