Combating Effort Avoidance in Computer Adaptive Practicing: Does a Problem Skipping Restriction Promote Learning?

Alexander O. Savi

Amsterdam Center for Learning Analytics
VU University Amsterdam

Chris van Klaveren

Amsterdam Center for Learning Analytics
VU University Amsterdam

Ilja Cornelisz

Amsterdam Center for Learning Analytics
VU University Amsterdam

Abstract

Effort is key in learning, evidenced by its omnipresence in both empirical findings and educational theories. At the same time, students are consistently found to avoid effort. In this study, we investigate whether limiting effort avoidance improves learning outcomes, and explore for whom this would be the case. In a large-scale computer adaptive practice system for primary education, over 150 000 participants were distributed across four conditions in which a problem-skipping option was delayed for 0, 3, 6, or 9 seconds. The results show that after a 14 week period, no average treatment effects in learning outcomes can be found between conditions. A substantive typology of students, based on the expected target mechanisms of the intervention, neither shows consistent conditional average treatment effects. Nevertheless, the substantive typology is shown to be meaningful, as the different types—toilers, skippers, and rushers—differ with respect to their learning outcomes. We argue that although the scale of the experiment suggests a precise null finding, the cumulative nature of the effect of problem skipping cautions against generalizing this finding to sustained intervention.

Keywords: Computer-Adaptive Practice, A/B Testing, Effortful Practice, Learning Outcomes, Online Learning, Heterogeneous Treatment Effects, Cumulative Effects

The importance of effort in learning is undisputed. Effort plays a pivotal role throughout educational theories. Also, some of the largest and most robust effects in learning rely on effort, such as the testing and spacing effect (Chen et al., 2021; Delaney et al., 2010; Rowland, 2014). At the same time, effort can be a demanding process that students

rather not exert, or more worrisome, fail to appreciate (Bjork et al., 2013). A key example of effort avoidance is problem skipping behavior. Such behavior is not a peculiarity. Walker et al. (2007) report students completing less than 60% of their attempted problems in an intelligent tutoring system, with students justifying problem skipping behavior by a lack of ability or motivation. Shanabrook et al. (2010) report students gaming an adaptive tutoring system by using speeding strategies, including problem skipping behavior, despite attempts to emphasize the importance of exerting effort. Savi et al. (2018) report students skipping problems in a computer adaptive practice system, encouraged by an incentive to balance speed and accuracy. Kerkman and Siegler (1993) found students to skip problems for which retrieval or back-up strategies did not suffice. And Turner et al. (2002) show that effort avoidance may signal a low-mastery environment, which is negatively related to mastery learning.

However, the effect of problem skipping on learning is to a large extent hypothetical: the literature is void of studies that directly test the effect. This might be explained by the fact that, in natural learning environments, manipulations of problem skipping behavior are not straightforward and proper randomization and double blinding are virtually impossible. On top of that, the effect of a single skip is probably small; the crux is in the cumulation. The amount of practice, the amount of skipping, and the speed of skipping, may all contribute to this cumulation. And very much interrelated, different students may benefit differently from manipulations of problem skipping behavior. To explore these issues, in this study we analyze data from an online randomized field experiment, or so-called A/B test. An A/B test allows targeted, randomized, and double-blind interventions in natural, online, learning environments. Also, it helps capture students' responses over significant periods of time. As such, the A/B test is an ideal method for the current evaluation.

We aim to establish whether restricting problem skipping behavior promotes learning. The question is evaluated using Math Garden, a computer-adaptive system for arithmetic practice. In Math Garden, learners are presented with a series of problems and gain or lose virtual coins by giving correct or incorrect responses. Problems can be skipped by the press of a virtual button. Savi et al. (2018) argued that delaying this problem skipping option may increase effortful practice. They showed that a delay decreased the use of the problem-skipping option and that problems were more often attempted. Response times for these attempts were found to be primarily slow, suggesting the use of more effortful strategies. The fact that fast guesses were actively discouraged by the punishment of fast incorrect responses further endorsed the evidence. However, Savi et al. (2018) did not examine the effect of delaying the problem skipping option on actual performance, which is

---

Alexander O. Savi <https://orcid.org/0000-0002-9271-7476>
Chris van Klaveren <https://orcid.org/0000-0001-5975-4454>
Ilja Cornelisz <https://orcid.org/0000-0002-9948-7861>

Support for data collection and the execution of the experiment was provided by Prowise Learn, a developer of computer-adaptive practice software. The authors bear full responsibility for the analysis and interpretation of the data.

Correspondence concerning this article should be addressed to Alexander O. Savi, Department of Educational and Family Studies, VU University Amsterdam, Van der Boechorststraat 7-9, 1081 BT Amsterdam, the Netherlands. Email: o.a.savi@gmail.com

where the current study steps in. Following up their study, we examine the hypothesis that the increased effort ultimately pays out in increased performance. That is, the results of this study aim to establish whether the increase in effort is desirable. On top of that, the current study explores meaningful heterogeneity in the potential effects. But first, we review two major theories that relate effort to learning and that underlie the studied computer-adaptive system and then discuss the mentioned challenges with the estimation of cumulative, heterogeneous effects.

Problem Skipping Behavior and Effortful Practice

In this section, we cover evidence from two fields that stress the benefits of various forms of effortful practice. Importantly, although the term effort is used throughout the discussed fields, it remains loosely defined at best. Nevertheless, taken together a clear picture emerges. We discuss retrieval practice and deliberate practice, and link their diverse yet intersecting theoretical underpinnings to problem skipping behavior in the studied computer-adaptive practice system.

Retrieval practice

It was Thorndike (1906) who summarized retrieval practice and its ubiquitous effect most eloquently:

As a rule, it is more economical to put things together energetically than to put them together often; close attention is better than repetition. The active recall of a fact from within is, as a rule, better than its impression from without; for recall is a helpful way to be sure of close attention and also forms the connection in the way in which it will later be required to act.

To date, retrieval practice, or “the active recall of a fact from within”, provides one of the largest effect sizes in the cognitive and learning sciences (Dunlosky et al., 2013; Karpicke & Roediger, 2008). The effect of retrieval practice was even given its own name—the testing effect—and is used to indicate the increased retention of the targeted information in long-term memory. On top of retention, there is also evidence that retrieval practice benefits transfer (Butler, 2010). A close relative, *distributed retrieval practice*, refers to retrieval practice spaced in time. Despite debates on the benefits of different spacing regimes—such as expanding or equally spaced repetition (Balota et al., 2007; Kang et al., 2014; Karpicke & Roediger, 2007)—the considerable effect of distributed retrieval practice also deserves its own name—the spacing effect—and is largely undisputed. The testing and spacing effects are associated with the need to actively recall facts from memory and the additional effort it requires when this activity is spaced out in time.

Distributed retrieval practice is at the core of Math Garden. Its adaptive algorithm creates a continuous stream of problems, targeted at the ability level of the learner. On top of that, problems are often spaced spontaneously, as problems are sampled within the ability range of the learner and are often sampled more than once. Such (distributed) retrieval practice suffers from problem skipping, as skipped problems are not retrieved. Although learners do not tend to skip the majority of problems, it is likely that exactly the problems that benefit most from retrieval practice are skipped. Indeed, problems that can

be retrieved directly do not add to retention and long term memory, whereas it is exactly the problems that cannot be retrieved directly that are most beneficial. There is even evidence that errors, induced by large spacings, facilitate learning (Pashler et al., 2003).

Deliberate practice

In order to obtain exceptional performance, in all sorts of domains, Ericsson et al. (1993) argued that *deliberate practice* is necessary. Ericsson and Lehmann (1996) defined deliberate practice as “individualized training activities especially designed by a coach or teacher to improve specific aspects of an individual’s performance through repetition and successive refinement”. Clearly, deliberate practice is much less precisely defined as retrieval practice. Estimating an effect size for deliberate practice has proven to be extremely difficult, as not only operationalizations but also contexts may differ greatly. Nonetheless, in a meta-analysis, Macnamara et al. (2014) estimated that deliberate practice explains 4% of the variance in educational performance. Although there is debate on both the validity of this estimate and the empirical significance of such an effect (Miller et al., 2018), there is consensus that deliberate practice does play a role in educational performance.

Key to deliberate practice is effortful study, which is defined as the union of practice and thought when applied just beyond one’s current level of ability (Ericsson, 2003). Indeed, experts consistently and consciously monitor and control their performance—to the extent of being able to report on their thoughts—in order to aid future performance (Winter et al., 2014). Crucially, the importance of thought in effortful study counters the popular belief that the end stage of expertise is automaticity, although both are evidently associated. It moreover illustrates the intimate entanglement of human ability and human learning (Ferguson, 1954): performance hardly ever reaches a stage where learning no longer plays a role. The idea of effortful study is closely related to Vygotsky’s zone of proximal development (Vygotsky, 1978); the neighbouring level of ability that can be reached with the aid of instructional scaffolds.

The core of Math Garden—its computer-adaptive practice engine—can be seen as a form of deliberate practice. Although virtual rather than human, Math Garden’s Elo rating system creates training activities targeted at the performance level of an individual learner. Moreover, with performance estimates within specific sub-domains, such as addition, it ensures targeted repetition and successive refinement within those domains. Importantly, the algorithm exploits an estimate of general domain ability, whereas more refined methods of deliberate practice may for instance additionally tap into individual learners’ misconceptions.

Crucially, in deliberate practice “individuals have to monitor their training with full concentration, which is effortful and limits the duration of daily training” (Ericsson & Lehmann, 1996). Mindless problem skipping, without taking the effort to answer it, clearly derails the effect of deliberate practice. Moreover, although automaticity is an evident asset of multiplication performance, learning multiplication requires effortful practice. That is, especially the problems that are at the boundary of an individual’s ability level are argued to provide the largest learning gains, when given sufficient thought.

Problem Skipping Behavior and Effect Estimation

Whereas the effect of effort is undisputed, and large for the discussed testing and spacing effects, limiting effort avoidance by means of a problem skipping delay is less likely to produce large short-term effects. Therefore, in this section we briefly describe recent discussions and developments that are highly relevant to the estimation of such effects. Specifically, we discuss three aspects of effects that are directly related to problem skipping: (1) effects are expected to be small, but nevertheless meaningful for the learning outcomes of individual students, (2) in case of heterogeneity, small average effects can signify large effects for particular individuals, as domains and students may be affected differently, and (3) effects are expected to cumulate over practice, and the effect size to increase over time.

The educational and learning sciences are deeply invested in the estimation of intervention effect sizes, for good reason. Educational interventions are demanded to *work*. Exemplary to the efforts to map effect sizes are the synthesis of meta-analyses by Hattie (2008), and various national initiatives such as the What Works Clearinghouse¹ in the United States, the Education Endowment Foundation² in the United Kingdom, and Education Lab³ in the Netherlands. But, despite their popularity, effect size estimates are not unproblematic. An effect size estimate depends on more than just the causal mechanism that an intervention activates. First, effect size estimates depend on study characteristics. Small sample studies tend to overestimate real-world effects, possibly due to superrealization (the disproportionate facilitation of the experimental group) or publication bias (due to unpublished null-findings) (Slavin & Smith, 2009). Also quasi-experimental (versus randomized) studies, published (versus unpublished) studies, developer-commissioned (versus independent) studies, and experimenter-made (versus independent) measures produce larger effect size estimates (Cheung & Slavin, 2016; Wolf et al., 2020).

Second, effect size estimates may depend on the targeted population. In social science, individual differences are “real rather than apparent” (Xie, 2013) and inherent individual-level heterogeneity is widely recognized. The appreciation of heterogeneous treatment effects has increased in recent years, with large-scale studies aiming to unravel differential effects for sub-groups as well as for circumstantial factors (e.g., Kizilcec et al., 2020; Yeager et al., 2019), and the introduction of novel techniques for estimating heterogeneous causal effects (e.g., Athey & Imbens, 2016). Third, effect size estimates may depend on cumulation. In various domains, including learning, small treatment effects may cumulate over time, for instance when students are repeatedly exposed to the same intervention (Abelson, 1985; Funder & Ozer, 2019). As a result of the discussed dependencies, comparing and combining effect sizes is often misguided (Simpson, 2017).

On top of these challenges, there is discussion on how to interpret and value effect sizes. First, Cohen’s (1988) influential rule’s of thumb for small, medium, and large effects (Cohen’s $d = \{.2, .5, .8\}$, and Pearson’s $r = \{.1, .3, .5\}$, respectively) are being reconsidered. In psychological science, Funder and Ozer (2019) argue that $r = .3$ reflects a large effect, and in (real-world) educational science, Dweck (2019) argues that $d = .2$ reflects the largest obtainable effect. When we disregard the limited value of rules of thumb for the moment,

¹<https://ies.ed.gov/ncee/wwc/>

²<https://educationendowmentfoundation.org.uk/>

³<https://education-lab.nl/>

these recent insights suggest that Cohen’s suggestions were inflated. Second, there is discussion on how to value ‘small’ effects (Dweck, 2018; Gelman, 2018; Macnamara, 2018). Funder and Ozer (2019) argue that small effects should not be automatically dismissed, given that the effects are reliably estimated. Indeed, small effects should be appreciated for good reasons. Small effects, when reliably estimated, represent actual effects. Moreover, such effects may represent conservative estimates of an intervention’s full potential: for some subgroups the effects may be larger, and effects may cumulate over time. In the end, education constitutes a universe of many subtle interventions, where small effects may not simply add up, but happily interact in possibly unexpected ways.

Overview

To sum up, there is strong theoretical ground for the mechanisms that the intervention aims to target. Also, the intervention is easily and broadly applicable, is applied on a massive scale with the potential to reach many students, and is not expected to deter the learning experience even for students that do not directly benefit from it. Such an intervention can be said to be universal (Greenberg & Abenavoli, 2016): it has low risk and the intervention can be easily applied to many students. In line with the discussed effect size issues, Greenberg and Abenavoli (2016) argue that a large sample should be used, the impact on development should be examined over time, and the examination should go beyond main effects. The Methods section provides an explanation of the employed computer-adaptive practice algorithm, a discussion of the experimental design, and descriptive statistics on the practicing process. The Results section present the empirical findings. The paper concludes with a discussion of the findings in the Discussion section.

Methods

This study builds upon the online field experiment and data described in Savi et al. (2018). In the experiment, a problem skipping option was delayed, aimed at increasing effortful practice. In the current evaluation, we evaluate the obtained differences in estimated ability between conditions, and explore different substantive subgroups. In the following, we discuss how student abilities were estimated, describe the experimental design, and provide descriptive statistics.

Math Garden and the Adaptive Algorithm

Math Garden is a large-scale computer-adaptive practice environment for learning arithmetic. Math Garden is primarily used in primary education in the Netherlands. It hosts over 25 000 problems divided over more than 20 domains, covering both arithmetic and related skills, such as word problems and telling time. Moreover, it hosts instruction videos, learning goals, and a teacher dashboard. It has a solid user-base, with over 150 000 active users, and about 8500 responses submitted per minute. The scale and intensity of the software makes it a great platform for studying the effect of problem skipping on ability.

Users land on a page that pictures a garden, where flowers and plants represent domains that a user can practice. When a domain is selected, the user gets to sequentially practice 10 problems. This is called a session. Responses must be submitted prior to a deadline, which is visualized by a decreasing number of virtual coins. If the submitted

response is correct, the remaining virtual coins are added to the user’s virtual budget, whereas in case of an incorrect response the remaining coins are subtracted from the budget. The budget can be used to buy virtual prizes. Users are not required to submit a response. If time runs out, the user is automatically presented the next problem. Users may also hit a button depicting a question mark to skip the problem before the deadline has passed. In both cases, users are neither rewarded nor punished with coins.

Math Garden employs an algorithm for micro-adaptive practice. In micro-adaptive practice, problems are selected *during* learning (Shute, 1993). Math Garden continuously matches users to problems based on their respective estimated abilities and difficulties. The core algorithm that estimates abilities and difficulties after each response is introduced by Klinkenberg et al. (2011):

$$\begin{aligned}\theta_p &\rightarrow \theta_p + K(S_{pi} - \mathcal{E}(S_{pi})), \\ \delta_i &\rightarrow \delta_i - K(S_{pi} - \mathcal{E}(S_{pi})),\end{aligned}\tag{1}$$

where a user ability estimate θ_p and problem difficulty estimate δ_i are differentially updated by the difference in the observed score S_{pi} and expected score $\mathcal{E}(S_{pi})$ for a problem. K is a simple scaling factor. The observed score is determined by the *signed residual time* scoring rule (Maris & van der Maas, 2012) and is visualized in Figure 1. Crucially, it mimics the rewards and punishments by means of the virtual coins, such that the incentives and actual scoring align. The expected score $\mathcal{E}(S_{pi})$ is defined as:

$$\mathcal{E}(S_{pi}|\theta_p, \delta_i) = d \frac{\exp(2d(\theta_p - \delta_i)) + 1}{\exp(2d(\theta_p - \delta_i)) - 1} - \frac{1}{\theta_p - \delta_i}\tag{2}$$

where d is the response deadline.

Users and problems are matched on the basis of the respective estimated abilities and difficulties, such that for each user the long term probability correct is held constant. Users can choose their own difficulty level, where the easy level corresponds to a probability correct of .9, the normal level to a probability correct of .75, and the hard level to a probability correct of .6. In the easy level and hard level, the obtained or lost virtual coins are respectively halved and doubled.

Math Garden’s adaptive algorithm provides that item difficulties and user abilities are continuously updated. As a consequence, divergent performance in one condition may influence the expected scores in other conditions, by means of the updates to the—global—item difficulties (Savi et al., 2021). The effect is considered negligible. Nevertheless, in an effort to constrain the item difficulties and prevent item drift, item difficulties were only updated for pairs of items for which a user performs above and below expectation within a single session.

Experimental Design

The experiment ran across three different domains: *addition*, *division*, and *one-two-three*. In both the addition and division domain, problems required a single operation and operands were either zero or positive. In the one-two-three domain, sets of three images had to be found. The images differed with respect to four features: size, color, shape, and

Figure 1
The signed residual time scoring rule.

texture. A set was correct if—for each of the features—the feature either had the same value across the selected images (e.g., size is three), or different values across the selected images (e.g., size is one, two, and three). The addition domain contained multiple-choice problems and had a deadline of 20 seconds, the division domain contained open-ended problems and had a deadline of 20 seconds, the one-two-three domain contained open-ended problems and had a deadline of 30 seconds.

The experiment ran for 14 weeks, from March 16 to June 22, 2016. Starting at 9:00 a.m., participants were assigned to their conditions. Depending on their assigned condition, participants were able to use the problem skipping option either as soon as a problem was presented, or after 3, 6, or 9 seconds. The problem skipping option was visible but greyed out during the time it was inactive. At the final day of the experiment, the experiment ended at 9:00 a.m.

Participants were assigned to the experimental conditions by means of bitwise right shifts and modulus transformations to their user IDs:

$$C_i = (i \gg S_d) \bmod 4, \quad (3)$$

where i represents the user id, C_i the assigned condition ($C_i = \{0, 1, 2, 3\}$ for respectively the control condition, and 3, 6, and 9 seconds delay conditions). Participants' user IDs were transformed using a bitwise right shift of respectively 0, 2, and 4, in the addition, division, and one-two-three domains ($S_d = \{0, 2, 4\}$). This ensured that within a domain, participants remained assigned to the same condition, whereas across domains assignment could vary. For example, a bitwise right shift of 2 for user ID 123456 (binary form 11110001001000000) results in 30864 (binary form 00111100010010000), and modulo 4 returns condition 0. User IDs were assigned in order of registration for Math Garden, starting with ID 1 at the release of the product over a decade ago.

In an A/B test such as this, experimental control is low and ecological validity is high. The practice platform and the interventions—the different delays of the problem skipping option—were identical across participants. However, there was no control over when, where, and how often participants practiced, nor over which device was used. At the same time, the experiment took place in the exact setting the participants normally practice. Moreover, the experiment was double-blinded, a condition that is often difficult to satisfy in educational research. Savi et al. (2017) explore the various benefits and challenges of A/B tests in more detail.

With each attempted problem, Math Garden automatically recorded the participant's response, response accuracy ("correct", "incorrect", "skip", "time out"), response time (ms), the date and time of the response, ability estimate (using the adaptive algorithms discussed in the previous section), and difficulty setting ("easy", "medium", "hard"). Also, the participant, parent, or teacher, provided the participant's year group (grade). Only participants in year groups three to eight were considered in the analyses. These year groups reflect Dutch primary education; age four to twelve. Notably, data quality suffers from known self-report issues.

These data were anonymized before being shared with the authors. Participants had the option to deny their data being used for scientific research. Evidently, the data of those participants was discarded from the analyses. Interested students, parents, and

teachers, were informed of the experiment in a blog post. The procedure was approved by the Ethics Review Board at the institution where the original research took place.

Descriptive Statistics

Figure 2 gives the sample size distributions across grades, for each of the three evaluated domains and four conditions. Addition is the most popular domain across all grades, and has its peak in grade four (over 4000 participants). Division starts considerably later in the curriculum, peaking in grade seven (almost 3000 participants), although it is practiced by some participants as early as in grade three. Finally, the logical reasoning game one-two-three follows the same shape as addition, but is not part of any school curriculum and as such is practiced by considerably less participants. Sample sizes are roughly equally distributed across conditions, which signals successful randomization.

Figure 3 gives the response frequencies during the experimental period, for all participants and across the three evaluated domains. Although participants could end the ten-problem sessions at any point in time, this standard session length is clearly reflected in the figure. The distributions are skewed to the right, showing that most participants practiced only a few problems, and a few participants practiced many problems. As the figure only shows participants with 100 or less responses, the skew is even more pronounced in reality. Given the expected cumulative nature of the intervention, the figure shows that intervention intensity is low for many participants.

Figure 4 gives proportions of problems skipped during the experimental period, for all participants and across the three evaluated domains. Most participants did not skip any problem, whereas a few participants skipped all problems. Naturally, at both extremes the participants' response frequencies might be low, which would help explain this observation. In between both extremes, the distribution is skewed to right, with the great majority of the participants skipping less than a quarter of the problems. Savi et al. (2018) showed that the intervention diminishes the use of the problem skipping option.

Figure 5 gives the median response times during the experimental period, for all participants and across the three evaluated domains. The accuracy of online response time estimates is limited, such that some response time estimates exceed the deadline. Most of these anomalies are time-outs, where a participant did not respond in time. These time-outs were capped at the domain's respective deadline. The figure shows that most participants' median response times are in the range of five to ten seconds, regardless of the domain and deadline. The peaks of incorrect responses at the deadlines merely show the participants with many timed-out problems. Also, the different delay conditions are reflected in the trimodal distribution of the skipped responses: the modes appear shortly after the three different delays.

To determine the effect of the intervention on the participants' ability levels, the participants' estimated abilities after their last observed responses were compared. Figure 6 gives the distributions of these estimates over time, for all participants and across the three evaluated domains. The distributions reveal midweeks and weekends, and a holiday in May. There is considerable spread in when a participant practiced her or his last problem, but most are observed in the final few weeks of the experimental period.

Figure 2
Bar plots showing the sample size distributions across grades (x-axis), per domain (panels) and condition (gray scale).

Figure 3
Bar plots showing the response frequency distributions (x-axis), per domain (panels). For clarity, only participants with 100 or less responses are shown.

Figure 4
Histograms showing the distributions of the proportions of problems skipped (x -axis), per domain (panels). Bin width = .02.

Figure 5

Histograms showing the median response time (in seconds) distributions (x-axis), per domain (panels). Bin width = 1. Frequency polygons (line types) differentiate the distributions of skipped problems (“?”), incorrect responses (“0”), and correct responses (“1”).

Figure 6
Bar plots showing the dates of the last observed ability estimates over time (x-axis), per domain (panels).

Results

We performed confirmatory and exploratory linear regression analyses to obtain the effect of the problem skipping delays on ability. We determined the participants' final ability estimates in the experimental period and compared these across conditions. Figure 7 shows raincloud plots (Allen et al., 2021) of the standardized estimates (z-scores), separately for each domain.

Confirmatory Analysis

In the confirmatory analysis, we determined the average treatment effect:

$$Y_i = \alpha + \beta D_i + \epsilon_i, \quad (4)$$

where Y_i represents the (standardized) ability of participants i . D indicates if the participant was assigned to the 3, 6, or 9 seconds delay treatment ($D = \{1, 2, 3\}$) or the control group ($D = 0$). The model was estimated separately for the different domains. Table 1 gives the estimation results. Clearly, the null hypothesis of no effect could not be rejected, regardless of the domain or delay.

Table 1
Estimation results: standardized ability estimates.

	(1)	(2)	(3)
0 Sec (Intercept)	-6.697 *** (0.026)	-14.347 *** (0.078)	1.885 *** (0.079)
3 Sec	0.030 (0.036)	0.106 (0.111)	-0.212 (0.112)
6 Sec	-0.011 (0.037)	0.058 (0.111)	-0.075 (0.112)
9 Sec	0.045 (0.037)	-0.067 (0.111)	-0.134 (0.112)
N	87329	43657	52649
R2	0.000	0.000	0.000

Note: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; (standard errors).

Exploratory Analysis

In education in general, as well as in the current study, treatment effects are expected to be highly heterogeneous. In the exploratory analyses, we aimed to map this

Figure 7

Raincloud plots showing the probability densities, boxplots, and raw data points of the studied ability estimates (x-axis). Data is shown for all four conditions (y-axis) and three domains (panels).

heterogeneity. We created a typology of theoretically relevant groups of participants, based on the theoretical foundation laid out in the introduction. We distinguished three types: *toilers*, *skippers*, and *rushers*. Toilers represent the participants that practice many problems, skippers the participants that skip many problems, and rushers the participants that provide fast responses. Participants can be members of no to all types, such that eight distinct groups emerge, and the difference in achieved cumulative effect can be observed.

In the exploratory analysis, we determined these conditional average treatment effects:

$$\begin{aligned}
 Y_i = & \alpha + \beta_1 D_i + \beta_2 T_i + \beta_3 S_i + \beta_4 R_i + D_i(\beta_5 T_i + \beta_6 S_i + \beta_7 R_i) + \\
 & \beta_8 T_i S_i + \beta_9 T_i R_i + \beta_{10} S_i R_i + D_i(\beta_{11} T_i S_i + \beta_{12} T_i R_i + \beta_{13} S_i R_i) + \\
 & \beta_{14} T_i S_i R_i + \beta_{15} D_i T_i S_i R_i + \epsilon_i,
 \end{aligned} \tag{5}$$

where T , S , and R indicate the toilers ($T = 1$), skippers ($S = 1$), and rushers ($R = 1$). Participants were assigned a position in this typology based on their practicing behavior in the pre-experimental period, starting December 9, 2015, 9 a.m.; a period of again 14 weeks. We determined the participants in the 6th, 7th, 8th, 9th, and 10th decile. Toilers represented the n th decile of participants with respect to the *number of problems practiced*. Skippers represented the n th decile of participants with respect to the *proportion of items skipped*. Rushers represented the n th decile of participants with respect to the *proportion of fast responses*. A response was marked fast if a participant's response time was below the median response time for a particular problem. Participants for whom no pre-experimental data were available were disregarded from the analyses. Figure 8 provides the theoretical framework and the number of participants assigned to the various positions in the typology.

Exploratory analyses were performed separately for each domain. The three domains have different characteristics that can have distinct influences. Addition is aimed at automation and contains multiple-choice problems. Division is aimed at automation and contains open-ended problems. One-two-three is not aimed at automation and contains open-ended problems. In the following, we first assess the used typology and then assess the average treatment effects conditional on the typology.

Typology

First, to illustrate how the results map to the typology, Figure 9 represents the results for the 10th decile in the addition domain as a Venn diagram. Then, Figures 10 to 12 give the regression coefficients and 95% confidence intervals for all different domains and deciles. Across domains and deciles, most regression coefficients for the toilers (T), skippers (S), and rushers (R) deviate from zero significantly. This suggests that the created typology is not only theoretically meaningful, but empirically meaningful as well. In the addition domain, the skipper type only seems to become relevant in the higher deciles, whereas in the one-two-three domain the rusher type seems least convincing. Since the one-two-three domain is not aimed at automation, this is expected.

Toilers obtain lower than average ability estimates in the addition domain, whereas they obtain higher than average ability estimates in the division and one-two-three domains. Skippers on the other hand, obtain higher than average ability estimates in the addition and one-two-three domains, whereas they obtain lower than average ability estimates in

Figure 8

Distribution of participants in the addition domain, during the experimental period, across the toiler–skipper–rusher typology. Numbers show sample sizes from 6th to 10th decile (top to bottom). Type assignment was determined on the basis of practice behavior in the pre-experimental period.

the division domain. Finally, rushers obtain higher than average ability estimates in the addition and division domains, and the 6th and 7th decile in the one-two-three domain, while this effect is reversed in the 9th and 10th decile in the one-two-three domain.

The combined types (TS, TR, SR, TSR) seem much less empirically meaningful. In the addition domain, the combined types that include skippers again only seem to become relevant in the higher deciles, whereas the combined type without skippers does seem to be meaningful: across deciles, toiler-rushers obtain higher than average ability estimates. Interestingly, the toiler-skipper-rusher type seems to be relevant in the addition and division domains, be it in the highest deciles.

Treatment Effects

Although the typology of toilers, skippers, and rushers seems empirically meaningful, Figures 10 to 12 show that average treatment effects are not conditional on the typology. Although incidentally coefficients deviate from zero significantly, robust patterns cannot be established.

Discussion

This study evaluated the effect of a problem skipping restriction on ability. In a large-scale computer-adaptive practice environment, an online field experiment (A/B test) was employed that delayed a problem skipping option with several seconds. In the A/B test, participants practiced as usual, to guarantee the ecological validity of the results. In the ability estimates, obtained during 14 weeks of practice, no differences were found between treatment groups. This suggests that the previously established increase in effortful practice (Savi et al., 2018) has no effect on the average ability during a period of 14 weeks. Or in a strict sense, that the absence of an average effect cannot be rejected.

The large scale of the experiment allowed the exploratory analysis of meaningful subgroups. On the basis of the mechanisms that the intervention is expected to target, a typology with three overlapping types was created: toilers that practice a lot, skippers that skip problems a lot, and rushers that provide fast responses a lot. The various overlaps between these three types create eight subgroups of learners. Learners were assigned to these groups on the basis of historic practice behavior. Crucially, the types showed differential growth in ability during the 14 week period, in all three studied domains. This not only suggests that learners can be reliably characterized on the basis of historic data, but that this characterization is also meaningful. Indeed, this characterization does not imply causality: whereas *rushers* may show differential growth in ability, *rushing* may just as well not. Also, the differential growth of toilers, skippers, and rushers is inconsistent across domains, which is puzzling.

The established subgroups allow to determine heterogeneity of the treatment effect by means of conditional average treatment effects. However, the exploratory analyses revealed no robust effects, reaffirming the overall null finding. Notwithstanding these results, most incidental treatment effects can be found among the toiler types, and most of these are in the expected direction, with a notable exception for the division domain. The cumulative nature of the intervention suggests that toilers, simply by their amount of practice, are crucial in establishing a possible effect. The subgroup that is most affected by the

Figure 9

Distribution of participants—in the experimental period—across three different types, 10th decile, in the addition domain. Type assignment was determined on the basis of practice behavior in the pre-experimental period. Results are ordered by condition.

Figure 10

Estimates of the regression coefficients in the addition domain over different deciles of the discussed typology. Columns distinguish the coefficients for the different conditions (problem skipping delay in seconds), whereas rows distinguish the types: T refers to Toilers, S to Skippers, and R to Rushers. Node size reflects statistical significance at a particular level of alpha. Ribbons reflect 95% confidence intervals.

Figure 11

Estimates of the regression coefficients in the division domain over different deciles of the discussed typology. Columns distinguish the coefficients for the different conditions (problem skipping delay in seconds), whereas rows distinguish the types: T refers to Toilers, S to Skippers, and R to Rushers. Node size reflects statistical significance at a particular level of alpha. Ribbons reflect 95% confidence intervals.

Figure 12

Estimates of the regression coefficients in the one-two-three domain over different deciles of the discussed typology. Columns distinguish the coefficients for the different conditions (problem skipping delay in seconds), whereas rows distinguish the types: T refers to Toilers, S to Skippers, and R to Rushers. Node size reflects statistical significance at a particular level of alpha. Ribbons reflect 95% confidence intervals.

intervention, toilers–skippers–rushers, reveals opposite exploratory effects for addition and division. Importantly, this subgroup suffers from a power problem, especially in the upper deciles, which is clearly reflected in the expanding confidence intervals. However, their response modes (respectively multiple-choice and open-ended) may provide an explanation. Learners that at first glance judge a problem to be too hard and would normally decide to skip it, may at second glance be able to locate the correct response among the options in a multiple-choice mode, whereas coming up with the correct response in an open-ended mode is less likely. Also, if a problem skipping delay would invoke guessing behavior for hard problems, guessing in a multiple-choice mode has a significant advantage over guessing in an open-ended mode.

The absence of (conditional) average treatment effects can have various reasons, the simplest being the absence of a causal mechanism. However, the strong evidence for the benefit of effortful practice in the literature, combined with the suggested increase in effort caused by the intervention (Savi et al., 2018), discourages a precipitate conclusion. The small but cumulative nature of the expected effect may cause a power problem. Although the number of participants is large and the experiment ran for 14 weeks, the majority of participants responded to a limited number of problems. On top of that, the intervention likely affects only a limited share of the responses (i.e., fast skips). Such a power problem is most pronounced with the group that benefits most: the participants that practice a lot, skip a lot, and provide fast responses a lot. Therefore, it is particularly promising that historic data enabled a reliable characterization of learners in the typology, such that a possible follow-up experiment would not need to target all learners.

Despite the null effects’ inconclusiveness, the current study demonstrates how a large-scale computer-adaptive practice environment can aid in estimating the small, cumulative, and conditional effects of an educational intervention on ability. The environment allowed us to study the effect of a problem skipping restriction with high ecological validity. Moreover, it allowed us to avoid a persistent problem that affects many field experiments in education: the impossibility to assure double blinding. Also, it allowed us to randomize the intervention at the level of the student and per studied domain. Finally, the scale of the environment, in combination with the additional response time measure, allowed us to explore heterogeneity of the effects by analyzing different types of learners. Although heterogeneous effects were not found, learners could be reliably characterized in ways that encourage follow-up research.

References

- Abelson, R. P. (1985). A variance explanation paradox: When a little is a lot. *Psychological Bulletin*, 97(1), 129–133. <https://doi.org/10.1037/0033-2909.97.1.129>
- Allen, M., Poggiali, D., Whitaker, K., Marshall, T. R., van Langen, J., & Kievit, R. A. (2021). Raincloud plots: A multi-platform tool for robust data visualization [version 2; peer review: 2 approved]. *Wellcome Open Research*, 4, 63. <https://doi.org/10.12688/wellcomeopenres.15191.2>
- Athey, S., & Imbens, G. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27), 7353–7360. <https://doi.org/10.1073/pnas.1510489113>

- Balota, D. A., Duchek, J. M., & Logan, J. M. (2007). Is expanded retrieval practice a superior form of spaced retrieval? a critical review of the extant literature (J. S. Nairne, Ed.). In J. S. Nairne (Ed.), *The foundations of remembering: Essays in honor of henry l. roediger, iii*. Psychology Press.
- Bjork, R. A., Dunlosky, J., & Kornell, N. (2013). Self-regulated learning: Beliefs, techniques, and illusions. *Annual Review of Psychology*, 64(1), 417–444. <https://doi.org/10.1146/annurev-psych-113011-143823>
- Butler, A. C. (2010). Repeated testing produces superior transfer of learning relative to repeated studying. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36(5), 1118–1133. <https://doi.org/10.1037/a0019902>
- Chen, O., Paas, F., & Sweller, J. (2021). Spacing and interleaving effects require distinct theoretical bases: A systematic review testing the cognitive load and discriminative-contrast hypotheses. *Educational Psychology Review*. <https://doi.org/10.1007/s10648-021-09613-w>
- Cheung, A. C. K., & Slavin, R. E. (2016). How methodological features affect effect sizes in education. *Educational Researcher*, 45(5), 283–292. <https://doi.org/10.3102/0013189x16656615>
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences*. Taylor & Francis.
- Delaney, P. F., Verhoeijen, P. P., & Spirgel, A. (2010). Spacing and testing effects, In *Psychology of learning and motivation*. Elsevier. [https://doi.org/10.1016/s0079-7421\(10\)53003-2](https://doi.org/10.1016/s0079-7421(10)53003-2)
- Dunlosky, J., Rawson, K. A., Marsh, E. J., Nathan, M. J., & Willingham, D. T. (2013). Improving students' learning with effective learning techniques. *Psychological Science in the Public Interest*, 14(1), 4–58. <https://doi.org/10.1177/1529100612453266>
- Dweck, C. S. (2018, June 26). *Growth mindset interventions yield impressive results*. <https://theconversation.com/growth-mindset-interventions-yield-impressive-results-97423>
- Dweck, C. S. (2019). The choice to make a difference. *Perspectives on Psychological Science*, 14(1), 21–25. <https://doi.org/10.1177/1745691618804180>
- Ericsson, K. A. (2003). Development of elite performance and deliberate practice: An update from the perspective of the expert performance approach (J. L. Starkes & K. A. Ericsson, Eds.). In J. L. Starkes & K. A. Ericsson (Eds.), *Expert performance in sports: Advances in research on sports expertise*. Champaign, IL: Human Kinetics.
- Ericsson, K. A., & Lehmann, A. C. (1996). Expert and exceptional performance: Evidence of maximal adaptation to task constraints. *Annual Review of Psychology*, 47(1), 273–305. <https://doi.org/10.1146/annurev.psych.47.1.273>
- Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, 100(3), 363–406. <https://doi.org/10.1037/0033-295x.100.3.363>
- Ferguson, G. A. (1954). On learning and human ability. *Canadian Journal of Psychology/Revue Anadienne de Psychologie*, 8(2), 95–112. <https://doi.org/10.1037/h0083598>
- Funder, D. C., & Ozer, D. J. (2019). Evaluating effect size in psychological research: Sense and nonsense. *Advances in Methods and Practices in Psychological Science*, 2(2), 156–168. <https://doi.org/10.1177/2515245919847202>

- Gelman, A. (2018, September 13). *Discussion of effects of growth mindset: Let's not demand unrealistic effect sizes*. <https://statmodeling.stat.columbia.edu/2018/09/13/discussion-effects-growth-mindset-lets-not-demand-unrealistic-effect-sizes/>
- Greenberg, M. T., & Abenavoli, R. (2016). Universal interventions: Fully exploring their impacts and potential to produce population-level impacts. *Journal of Research on Educational Effectiveness*, 10(1), 40–67. <https://doi.org/10.1080/19345747.2016.1246632>
- Hattie, J. (2008). *Visible learning: A synthesis of over 800 meta-analyses relating to achievement*. Abingdon, Routledge.
- Kang, S. H. K., Lindsey, R. V., Mozer, M. C., & Pashler, H. (2014). Retrieval practice over the long term: Should spacing be expanding or equal-interval? *Psychonomic Bulletin & Review*, 21(6), 1544–1550. <https://doi.org/10.3758/s13423-014-0636-z>
- Karpicke, J. D., & Roediger, H. L. (2007). Expanding retrieval practice promotes short-term retention, but equally spaced retrieval enhances long-term retention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33(4), 704–719. <https://doi.org/10.1037/0278-7393.33.4.704>
- Karpicke, J. D., & Roediger, H. L. (2008). The critical importance of retrieval for learning. *Science*, 319(5865), 966–968. <https://doi.org/10.1126/science.1152408>
- Kerkman, D. D., & Siegler, R. S. (1993). Individual differences and adaptive flexibility in lower-income children's strategy choices. *Learning and Individual Differences*, 5(2), 113–136. [https://doi.org/10.1016/1041-6080\(93\)90008-g](https://doi.org/10.1016/1041-6080(93)90008-g)
- Kizilcec, R. F., Reich, J., Yeomans, M., Dann, C., Brunskill, E., Lopez, G., Turkay, S., Williams, J. J., & Tingley, D. (2020). Scaling up behavioral science interventions in online education. *Proceedings of the National Academy of Sciences*, 117(26), 14900–14905. <https://doi.org/10.1073/pnas.1921417117>
- Klinkenberg, S., Straatemeier, M., & van der Maas, H. L. J. (2011). Computer adaptive practice of maths ability using a new item response model for on the fly ability and difficulty estimation. *Computers & Education*, 57(2), 1813–1824. <https://doi.org/10.1016/j.compedu.2011.02.003>
- Macnamara, B. N. (2018, June 26). *Schools are buying 'growth mindset' interventions despite scant evidence that they work well*. <https://theconversation.com/schools-are-buying-growth-mindset-interventions-despite-scant-evidence-that-they-work-well-96001>
- Macnamara, B. N., Hambrick, D. Z., & Oswald, F. L. (2014). Deliberate practice and performance in music, games, sports, education, and professions: A meta-analysis. *Psychological Science*, 25(8), 1608–1618. <https://doi.org/10.1177/0956797614535810>
- Maris, G., & van der Maas, H. L. J. (2012). Speed-accuracy response models: Scoring rules based on response time and accuracy. *Psychometrika*, 77(4), 615–633. <https://doi.org/10.1007/s11336-012-9288-y>
- Miller, S. D., Chow, D., Wampold, B. E., Hubble, M. A., Re, A. C. D., Maeschalck, C., & Bargmann, S. (2018). To be or not to be (an expert)? revisiting the role of deliberate practice in improving performance. *High Ability Studies*, 1–11. <https://doi.org/10.1080/13598139.2018.1519410>
- Pashler, H., Zarow, G., & Triplett, B. (2003). Is temporal spacing of tests helpful even when it inflates error rates? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(6), 1051–1057. <https://doi.org/10.1037/0278-7393.29.6.1051>

- Rowland, C. A. (2014). The effect of testing versus restudy on retention: A meta-analytic review of the testing effect. *Psychological Bulletin*, 140(6), 1432–1463. <https://doi.org/10.1037/a0037559>
- Savi, A. O., Ruijs, N. M., Maris, G. K. J., & van der Maas, H. L. J. (2018). Delaying access to a problem-skipping option increases effortful practice: Application of an a/b test in large-scale online learning. *Computers & Education*, 119, 84–94. <https://doi.org/10.1016/j.compedu.2017.12.008>
- Savi, A. O., ten Broeke, N., & Hofman, A. D. (2021, June 22). Adaptive learning systems and interference in causal inference, In *Educational A/B testing at scale: The second virtual workshop at Learning @ Scale 2021*. <https://sites.google.com/carnegielearning.com/edu-ab-testing-at-scale-2021/>
- Savi, A. O., Williams, J. J., Maris, G. K. J., & van der Maas, H. L. J. (2017, February 27). *The role of A/B tests in the study of large-scale online learning*. <https://doi.org/10.17605/OSF.IO/83JSG>
- Shanabrook, D. H., Cooper, D. G., Woolf, B. P., & Arroyo, I. (2010). Identifying high-level student behavior using sequence-based motif discovery (R. S. Baker, A. Merceron, & P. I. Pavlik Jr., Eds.). In R. S. Baker, A. Merceron, & P. I. Pavlik Jr. (Eds.), *Educational data mining 2010*.
- Shute, V. J. (1993). A macroadaptive approach to tutoring. *Journal of Artificial Intelligence in Education*, 4(1), 61–93.
- Simpson, A. (2017). The misdirection of public policy: Comparing and combining standardised effect sizes. *Journal of Education Policy*, 32(4), 450–466. <https://doi.org/10.1080/02680939.2017.1280183>
- Slavin, R., & Smith, D. (2009). The relationship between sample sizes and effect sizes in systematic reviews in education. *Educational Evaluation and Policy Analysis*, 31(4), 500–506. <https://doi.org/10.3102/0162373709352369>
- Thorndike, E. L. (1906). *The principles of teaching; based on psychology*. New York: A.G. Seiler. <https://archive.org/details/principlesofteac00thor>
- Turner, J. C., Midgley, C., Meyer, D. K., Gheen, M., Anderman, E. M., Kang, Y., & Patrick, H. (2002). The classroom environment and students' reports of avoidance strategies in mathematics: A multimethod study. *Journal of Educational Psychology*, 94(1), 88–106. <https://doi.org/10.1037/0022-0663.94.1.88>
- Vygotsky, L. S. (1978). *Mind in society*. Harvard University Press.
- Walker, E., Rummel, N., McLaren, B. M., & Koedinger, K. R. (2007). The student becomes the master: Integrating peer tutoring with cognitive tutoring (C. A. Chinn, G. Erkens, & S. Puntambekar, Eds.). In C. A. Chinn, G. Erkens, & S. Puntambekar (Eds.), *The computer supported collaborative learning (cscl) conference 2007*, New Brunswick, NJ, USA: International Society of the Learning Sciences. <https://repository.isls.org/handle/1/3442>
- Winter, S., MacPherson, A. C., & Collins, D. (2014). “[to] think, or not to think, that is the question”. *Sport, Exercise, and Performance Psychology*, 3(2), 102–115. <https://doi.org/10.1037/spy0000007>
- Wolf, R., Morrison, J., Inns, A., Slavin, R., & Risman, K. (2020). Average effect sizes in developer-commissioned and independent evaluations. *Journal of Research on*

- Educational Effectiveness*, 13(2), 428–447. <https://doi.org/10.1080/19345747.2020.1726537>
- Xie, Y. (2013). Population heterogeneity and causal inference. *Proceedings of the National Academy of Sciences*, 110(16), 6262–6268. <https://doi.org/10.1073/pnas.1303102110>
- Yeager, D. S., Hanselman, P., Walton, G. M., Murray, J. S., Crosnoe, R., Muller, C., Tipton, E., Schneider, B., Hulleman, C. S., Hinojosa, C. P., Paunesku, D., Romero, C., Flint, K., Roberts, A., Trott, J., Iachan, R., Buontempo, J., Yang, S. M., Carvalho, C. M., ... Dweck, C. S. (2019). A national experiment reveals where a growth mindset improves achievement. *Nature*, 573(7774), 364–369. <https://doi.org/10.1038/s41586-019-1466-y>