Nudging students to reduce procrastination in office hours and forums

Eric Fouh, Wellington Lee, Ryan S. Baker
University of Pennsylvania
Philadelphia, PA
Email: efouh@cis.upenn.edu, wtnlee@seas.upenn.edu, ryanshaunbaker@gmail.com

***Abstract*—In this article, we present the results of a study aiming to understand the impact of email nudge notification on students' procrastination in office hours, and Piazza (QA forum) in a CS1 course at a large research university. With this study, we sought to understand if email nudges can be a useful tool in improving student's learning behaviors, especially procrastination. After the first two homeworks, we randomly split students into two groups; the treatment group received the email, and the control group did not. The treatment group was further divided into two groups: one for the students who performed above the median (of the combined grades of homework 1 and 2) and those who performed below the median. Each sub-group received a slightly different version of the email. We found that students in the treatment group did not change their office hours' attendance and Piazza interactions. We also found no difference in homework grades. However, students in the treatment group used more free late days on the following (third) homework. However, the change was short-lived, and they reverted to the pre-email level of late days usage on the fourth homework.**

***Keywords*-Nudge; Learning Behavior; Procrastination**

I. Introduction

Many students who take an introductory programming course come to the course with a lack of familiarity with fundamental computer science concepts and skills [1]. Some but not all of these students struggle. One of the foremost reasons students struggle is a lack of self-regulated learning (SRL) skill [2], [3].

SRL has been defined as an active, constructive process where learners set learning goals and then attempt to monitor, regulate, and control their cognitive, metacognitive, motivational, and affective processes to achieve those goals [4]. The lack of effective SRL skills can result in less than optimal learning behaviors like procrastination and failure to seek help (for example, from course staff) when needed [5], [6].

In this article, we present the results of a study aiming to understand the effect of an email notification (nudge) on students' learning behavior in a CS1 course at a large research university. After the first two homeworks (roughly a month into the semester), we randomly split students into two groups; one group received the email, and the other did not. The treatment group was further divided into two groups: one for the students who performed above the median (of the combined grades of homework 1 and 2), and those who performed below the median. Each sub-group received a slightly different version of the email. Both versions of the email reminded the students of the negative relation between procrastination and grades and between office hours' wait time and homework due date. The emails encouraged them to start the assignment early, engage with the course staff during office hours, or post questions on Piazza.

We found that students in the treatment group did not interact more with the course staff (office hours attendance) or participate more on Piazza after receiving the email. However, students in the treatment group who performed below the median used more (free) late days on the following homework (#3), and their performance improved. We also found that the change was short-lived, and they reverted to the pre-email level of late days usage on homework 4, and their performance followed a similar trend.

II. Research questions

We hypothesized that an email notification reminding students to start homework and seek help early would lead to learning behavior changes that will impact students' performance.

The specific learning behaviors that we are monitoring are **late participation (procrastination), interactions on Piazza, and office hours attendance**; and the performance is measured as the **grade** earned on the homeworks.

The following research questions guided our investigation:

- RQ1: What is the impact of the email notification on students' learning behaviors?
- RQ2:What is the impact of the email notification on students' performance?
- RQ3: How long does the email notification impact on students' behavior lasts?

III. Related work

SRL involves cognition, meta-cognition, and motivational aspects [7]; and there is considerable evidence that the use of SRL processes and strategies are correlated with positive learning outcomes in STEM domains [8]–[10]. In their taxonomy of intervention types in computer science education, Szabo et al. noted that "metacognitive" interventions *"can include opportunities for students to reflect on their learning,* *but also updates that aim to improve time management, organization, or remove or reduce plagiarism"* [11]. Nudge is defined as *"ways of influencing choice without limiting the choice set or making alternatives appreciably more costly in terms of time, trouble, social sanctions, and so forth."* [12]. Furthermore, a nudge that makes its intent and means clear to the person being nudged and that prompt the user to reflect and choose differently is classified as a *transparent type 2 nudge* [13].

The impact of nudges in education is mixed. In their review of nudges in education, Damgaard and Nielsen [14] found the followings: informational and reminders nudges do not always improve student outcomes; interventions targeting skills like grit, forward-looking behavior generally have a positive impact; and that social comparison nudges have provided mixed results.

Although many CS educators have reported many metacognitive interventions, very few nudge email intervention results have been published. An example of transparent type 2 nudge email was used to reduce procrastination by Edwards et al. [15]. Procrastination is a big problem for many college students [16], and is associated with worse learning outcomes [17]. In a junior-level advanced data structures course, the authors sent regular emails to students reminding them how their progress compares to others (social comparison), as well as noting the course expectations. They found that nudging students with emails reduced the number of late submissions by one third. Another transparent type 2 nudge email intervention is described in [18]. The authors gathered data from the IDE and used the university's registrar system to send nudge emails to students enrolled in an introductory CS course, giving individualized advice on how to succeed better. The treatment group had a higher average grade, although the overall pass rate was not different between groups.

IV. Participants and Methods

Participants were students enrolled in the CS1 course at the University of Pennsylvania. Students were assigned to this course if they did not have previous programming experience. Students have four late days to use as they wish throughout the semester. Students can use at most two late days on any given assignment, and there is a total of 9 programming assignments in the course. No submission is allowed two days after the due date of the homework. Late days are deducted automatically for late submissions, and students do not need to ask permission or give any reason for using them. Due to COVID-19, the course instructor changed the late days after homework 4, and students were allowed to submit all the remaining homeworks up to 48 hours after the due date without penalty. We used late days as a proxy to estimate how much additional time students spent on homework. If there is no difference in late days usage between the groups before the intervention, and we see a difference after the intervention, the intervention likely impacted the behavior change.

Our experiment used the first four programming assignments of the course and involved a random assignment of the students ($n = 218$) into two roughly equally sized groups, both with 109 students. The group assignment happened at the beginning of the semester (before the course drop deadline), and at the end of the experiment, 192 students were enrolled in the course. Within the treatment group, we will refer to the group of students in the treatment group who performed below the median as $T_L(n = 35)$, and the ones who performed at or above the median as $T_H(n = 58)$.

After the first two homework assignments, we sent an email to the students within the treatment group, with the contents varying a little between T_L and T_H. The email was sent after homework 2 grades were released (three days after the due date), and four days before the next homework (3) due date. The emails were sent early on the first day of office hours. The course staff holds office hours Monday through Thursday, and homeworks are due on Thursday evening. Students in the control group did not receive the email.

A copy of the email (T_L) is included in section IV-A. Note that the only difference between the two versions was the first paragraph. T_H participants' email had *"Good job so far in [Class Name]! You've done really well on the homework."* instead of *"We noticed that you lost some points on ..."*.

The emails were designed to convey the following information to the students:

- Procrastination is correlated with lower grades.
- Office hours get significantly more crowded closer to the due date.
- Many resources available to them in case they needed help.

To measure this intervention's effect, we continually tracked students' behaviors throughout the next two homework (two weeks of classes). We collected daily interaction logs from Piazza and our office hours queue software. We tracked daily whether or not a student interacted with Piazza or went to office hours. The office hour queue software recorded attendance to office hours. Students had to log in and submit their question(s) before meeting with a TA; they cannot just show up.

A. Sample Email

Here is the email that was sent to students in the T_L group:

Dear xx,

We noticed that you lost some points on the first few homework assignments for various reasons and we would be happy to help ensure you understand everything- xxxx is fast moving and it's important to understand early concepts well.

We just wanted to reach out to reiterate that Piazza and Office Hours are an excellent resource for getting help, *especially as homework assignments become more complex and more difficult to debug.*

Starting early and coming to Office Hours earlier is extremely important as many scientific studies have shown that procrastination and grades are negatively correlated. That is, the later a student starts, usually the less likely they are to do well and learn fully from the assignment (and the same holds for exams).

If that's not enough, the following are average wait-time for Office Hours in the days leading up to a homework due date:

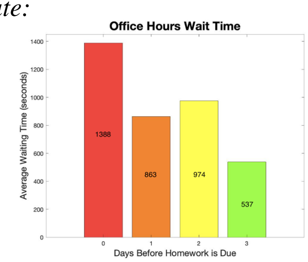

Hope to see you soon,
Prof xxx and your xxxx Course Teaching Staff

V. Interactions on Piazza

A. Data

In its raw format, the data we collected every day was the sum of the cumulative number of posts made and viewed over the account's lifetime. We put the data in an indicator variable format, where a student would have a 1 in a specific day if they had interacted with the question/answer platform and a 0 otherwise. Thus, each student had a vector of length 28 (7 days per homework) for the four homeworks we used in our study.

We then further split each student's vectors into pre-intervention and post-intervention regions.

B. Analysis

To analyze this data, we performed a k-means clustering analysis. We clustered the data before and after the intervention, and we found that the k-means algorithms centroids represent the following student profiles:

1) **Early participation**: These students were active on Piazza as soon as the homework was released and stopped being active before the due date. $n = 65$ pre-intervention, and $n = 33$ post-intervention.
2) **No participation**: These students did not use Piazza much. $n = 63$ pre-intervention, and $n = 68$ post-intervention.
3) **Late participation**: These students were mainly active on Piazza close to (2 days before) the due date. $n = 28$ pre-intervention, and $n = 59$ post-intervention.
4) **Constant participation**: These students were active on Piazza from the moment the homework was released until the due date. $n = 46$ pre-intervention, and $n = 42$ post-intervention.

Figures 1, 2, 3, and 4 show examples of each profile.

Splitting up each student's interaction vector into pre-intervention/ post-intervention allowed us to see how students moved between the different clusters, and potentially make claims about how our treatment potentially impacted the participants.

We speculated that:

1) Profiles 1 and 4 display positive behavior. Students in those groups are unlikely to be procrastinators.
2) Some students with profile 3 are likely to be procrastinators.

C. Results

We define a positive movement as a student moving from profile 2 or 3 to profile 1 or 4 (negative to positive profiles). We define a negative movement as the opposite (moving from cluster 1 and 4 to 2 or 3).

Movements between profiles 2 and 3 are considered "consistent negative", and movements between 1 and 4 are "consistent positive", as this means that students are not changing behaviors and are staying within their own types of groups.

Table I
Movements between Piazza interactions profiles

Type of Movement	Control	Treatment
Negative change (1,4 $\longrightarrow$ 2,3)	0.21	0.23
Positive change (2,3 $\longrightarrow$ 1,4)	0.06	0.09
Consistent positive (1 $\longleftrightarrow$ 4)	0.25	0.27
Consistent negative (2 $\longleftrightarrow$ 3)	0.47	0.40

Table II
*

Each cell represents the proportion of participants who changed their Piazza interactions after the intervention.

Table I shows how participants moved between profiles during the experiment. Overall there were similar movements between the control and the treatment group. The largest apparent difference was a 7% difference between the control and treatment group in the consistent negative category, but this was not statistically significant according to a two sample test of the significance of the difference between independent proportions, $Z = 1.0424$, two-tailed $p = 0.298$.

VI. Office Hours Attendance

A. Data

The course staff collects granular data on when students come to office hours. Every single interaction must be logged per course policy. As a result, we generated a very similar profile as the Piazza data above, with indicator variables depicting whether students come to office hours or not.

The course staff only hosts office hours Monday through Thursday, so as opposed to 14 data points (the forum is

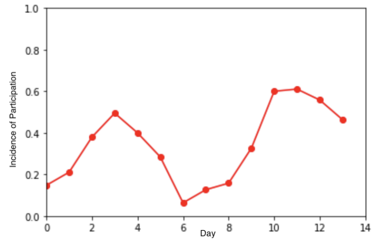

Figure 1. Profile 1: Early participation

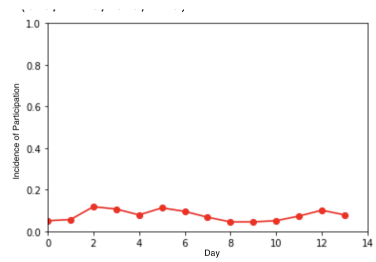

Figure 2. Profile 2: No participation

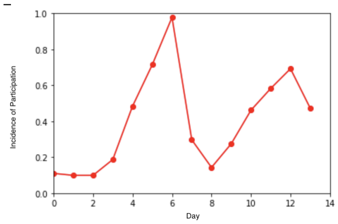

Figure 3. Profile 3: Late participation

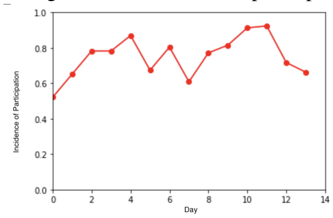

Figure 4. Profile 4: Constant participation

always available) per 2 homeworks, there are only 8 data points. Also, homeworks are due on day 4/8 and assigned on day 0 and 4.

B. Analysis

We identified the following profiles based on office hours attendance patterns using a K-means algorithm (applying the algorithm once to both pre and post-intervention):

1) **Early attendance**: These students attended office hours a couple of days before the due date. $n = 17$ pre-intervention, and $n = 17$ post-intervention.
2) **No attendance**: These students did not come (or came once or twice) to office hours. $n = 34$ pre-intervention, and $n = 48$ post-intervention.
3) **Late attendance**: These students came to office hours close to (2 days before) the due date. $n = 15$ pre-intervention, and $n = 6$ post-intervention.
4) **Regular attendance**: These students tended to go to office hours most of the days between when the homework is released and when it is due. $n = 11$ pre-intervention, and $n = 6$ post-intervention.

C. Results

We looked into the proportion of the sub-population that moved from one group to another group between the first half and second half of our time period of interest. Tables III and V summarize the results.

We found that more participants moved from group 2 to group 3 in the treatment group (7%) compared to the control group (no movement). Moving from group 2 to 3 represents a shift from people not interacting with the teaching staff to interacting with the staff quite a bit. Giving that the email notification encouraged participants to attend office hours, this can be interpreted as an indication of the intervention's impact, albeit weak.

Table III
MOVEMENTS BETWEEN OFFICE HOURS ATTENDANCE CLUSTERS: CONTROL GROUP

	Group 1	Group 2	Group 3	Group 4
Group 1	-	0.03	0.07	0.02
Group 2	0.02	-	0	0.01
Group 3	0.01	0.02	-	0.04
Group 4	0	0.02	0.02	-

Table IV
*

Each cell represents the proportion of participants who moved from a cluster (row) to another one (column) after the intervention.

Table V
MOVEMENTS BETWEEN OFFICE HOURS ATTENDANCE CLUSTERS: TREATMENT GROUP

	Group 1	Group 2	Group 3	Group 4
Group 1	-	0.15	0.04	0.02
Group 2	0.05	-	0.07	0.02
Group 3	0.01	0.04	-	0.05
Group 4	0	0.04	0.04	-

Table VI
*

Each cell represents the proportion of participants who moved from a cluster (row) to another one (column) after the intervention.

VII. LATE DAYS USAGE

We looked at the average number of late days used by the students. We looked at these numbers on a per-homework basis. We investigated if there were diverging patterns before and after the intervention between the two groups and between the low performers in either group.

We are not particularly concerned about running into the edge case where a student used up all their late days in the first two assignments- this is unwise, and the student would have likely dropped the class before it got to that point.

A. Data

The teaching staff records every late day used by the students. We computed the average counts of how many late days were used on a per-homework basis. The results are shown in table VII.

Table VII
NUMBER OF LATE DAYS USED FOR EACH HOMEWORK

	Hw1	Hw2	Hw3	Hw4
Treatment	0.06	0.2	1.5	0.4
Control	0.1	0.65	0.6	0.1

B. Analysis

We plotted the control and treatment groups' late day usage to see if they followed a similar pattern. We found that both groups followed a similar trend except for homework 3 (see figure 5).

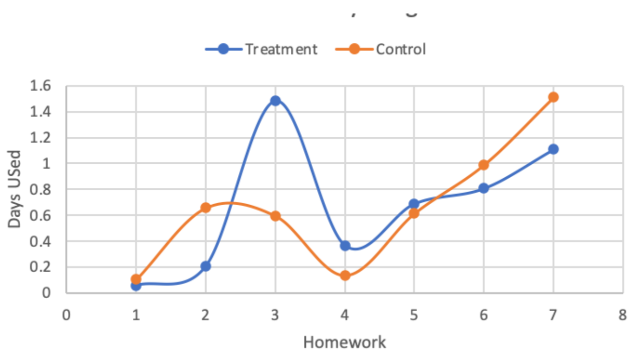

Figure 5. Late days usage trends treatment vs. control

C. Results

The particular point of interest is week 3, the homework after our treatment. During this week, we saw a statistically significant spike in usage of late days (see figure 5 and table VIII). A driver of this could be an inherently harder homework, but again, all students were subjected to the same academic conditions at this point, so differences in difficulty are unlikely to drive the differences between the treatment and control groups. Instead, we believe that our treatment increased late day usage. Within the treatment population, the average number of late days used for homework 3 was 0.45 for T_H (on par with the control group), and 0.82 for T_L; and that difference was statistically significant (z = 1.96, p-value = 0.017). We compared late days usage between T_L students and control students who performed below the median. Below median students in the control group used on average 0.41 days, and the difference was statistically significant (z = 1.96, p-value = 0.012). The above findings confirmed that the surge in late days usage was driven by students in the T_L group following the email nudge.

We saw that students in the treatment group stopped using more late days after homework 3, suggesting that the impact from the email nudge was short-lived. We included homework 5 to 7 in our graph to show that students in the treatment group consistently used the same or fewer late days except for the homework that followed the email – even after doing so had no penalty. As explained in the methods section, all students had a 48-hours grace period after homework 4. Which means that the differences in late days usage after homework 3 cannot be explained by the fact that students in the control group ran out of late days.

Table VIII
STATISTICAL ANALYSIS OF LATE DAYS USAGE: CONTROL VS. TREATMENT GROUPS

	Hw1	Hw2	Hw3	Hw4
z-statistic	0.508	-0.150	1.770	0.430
p-value	0.306	0.440	0.0383*	0.333

Table IX
*

* significant at level $p < 0.05$

VIII. PERFORMANCE

We sought to investigate the impact of our intervention on student's performance. To understand the differences between the treatment and control groups with respect to their overall performances, we first looked at the difference in average grade between the groups and the proportion of students in each treatment condition above and below the median for each homework.

A. Data

The average grade is shown in table X, and the proportion of treatment participants is presented in table XI.

Table X
AVERAGE GRADE ON HOMEWORK

	Hw1	Hw2	Hw3	Hw4
Treatment	42.38	44.32	40.51	44.90
Control	41.06	45.73	43.30	45.15

Table XI
PROPORTION (%) OF TREATMENT STUDENTS ABOVE THE MEDIAN

Hw1	Hw2	Hw3	Hw4
0.42	0.48	0.52	0.47

B. Analysis

We found that the difference between the average grades on each homework between our control group and treatment group was not significant (see table XII). This was somewhat expected because of the tight distribution of grades on each homework. Students have an unlimited number of submissions on each assignment, and they receive immediate feedback. We also observed that the number of treatment students who performed above the median increased for homework 3 (see table XI).

Table XII
STATISTICAL ANALYSIS OF AVERAGES GRADES, CONTROL VS TREATMENT GROUP

	Hw1	**Hw2**	**Hw3**	**Hw4**
z-statistic	0.169	-0.109	-0.089	-0.492
p-value	0.86	0.913	0.929	0.623

C. Results

Our analysis showed no difference in both groups' average grades throughout our experiment, even though we saw a slight increase in the number of the treatment population above the median for homework 3.

IX. DISCUSSION

A. RQ1: What is the impact of the email notification on students' learning behaviors?

Our study sought to investigate an email nudge's impact on the following learning behaviors: procrastination on Piazza and office hours. The email explicitly reminded students to use Piazza if they need help. However, when we looked at Piazza interaction data, we found few differences between the treatment and control groups. We also found that students in the treatment group were less likely to demonstrate consistent negative behavior.

When looking at office hours attendance, we found that 7% more students in the treatment group went from no office hours attendance to early office hours attendance. Early office hours attendance is a marker of starting the homework early.

We will note that the email nudge contained a graph showing a sharp increase in office hours wait time closer to the homeworks due date, but no visual information for the Piazza reminder. It is possible that using a visualization might result in a stronger effect as reported by Ilves et al. [19]. They divided students in three groups: one with no nudge (control), one with a textual visualization nudge, and one with a radar-like visualization nudge. The visualizations allowed students to compare their progress to the progress of others in the course. They found that the students with the radar visualization performed better than the textual visualization group. And, in some cases, the students with the textual visualization performed worse than the students in the control group with no visualization [19]. Further investigations are needed to tease out what should be included in the email nudge to get a stronger effect.

B. RQ2:What is the impact of the email notification on students' performance?

We found no significant difference in homework grades post-intervention. However, more students in the treatment group performed above the median for homework 3 (due after the nudge).

The late days usage and the early office hours attendance move seems to indicate that some students in the treatment group started the homework early, while others took advantage of the late days. We need to investigate these associations further. Some of the questions that need to be answered are: "Do the days usage and the early office hours attendance equally contribute to the improved performance?", or "is there a causal relationship between late-day usage and early office hours attendance and homework grade?".

C. RQ3: How long does the email notification impact on students behavior last?

We observed that the increase in late days usage did not last beyond the homework immediately following the intervention. This finding is in line with Damgaard and Nielsen's results. They found that the impact of reminder nudges is mostly short-lived [14].

X. THREATS TO VALIDITY

One major consideration in this study is the limitations of the data available to identify student behavior. Since we did not have access to IDE data, we relied on Piazza and office hours data to estimate when a student started the homework. However, the data collected might not give the most accurate indication of when a student began to work on the assignment. For example, a student might visit Piazza before starting a homework. We do not have a good way to estimate and identify the students who displayed that behavior. Also, the office hours queue might not capture all the students' interactions. For example, two students with the same question can attend office hours together with only one student entering the queue.

XI. CONCLUSION AND FUTURE WORK

Our experiment revealed that a single nudge to encourage students to participate more in office hours and start homework earlier resulted in mixed results. In terms of student outcomes, we saw no significant difference between groups. The main observed effect was the larger usage of late days by the treatment group. However, the impact did not persist beyond the week of treatment. The significant spike of (free) late-day usage could imply a longer time on task and engagement by students in the treatment group during the week of the intervention.

The results of this experiment may have been weakened by the relatively low dosage of the intervention (a single nudge). An obvious follow-up experiment would involve surveying students to better understand what makes this type of nudge effective or not. We also have plans to send an email to students after every single homework; we hypothesize that more frequent nudges will keep the class more at the forefront of student minds and thus result in higher engagement.

In addition, a different cost function for the k-means algorithm could improve performance, or some sort of boosted learning or time series analysis (ARIMA) might also elucidate certain characteristics of students that educators may find useful.

Lastly, we would like to look at lagging indicators to analyze the "stickiness" of these treatments; we hope to find methods that stay with students beyond a week or two. By establishing good habits, we hope we are able to set up students for success into the future.

References

[1] T. Beaubouef and J. Mason, "Why the high attrition rate for computer science students: some thoughts and observations," *ACM SIGCSE Bulletin*, vol. 37, no. 2, pp. 103–106, 2005.

[2] S. Bergin, R. Reilly, and D. Traynor, "Examining the role of self-regulated learning on introductory programming performance," in *Proceedings of the First International Workshop on Computing Education Research*, ser. ICER '05. New York, NY, USA: Association for Computing Machinery, 2005, p. 81–86. [Online]. Available: https://doi.org/10.1145/1089786.1089794

[3] K. Falkner, R. Vivian, and N. J. Falkner, "Identifying computer science self-regulated learning strategies," in *Proceedings of the 2014 Conference on Innovation Technology in Computer Science Education*, ser. ITiCSE '14. New York, NY, USA: Association for Computing Machinery, 2014, p. 291–296. [Online]. Available: https://doi.org/10.1145/2591708.2591715

[4] P. R. Pintrich, "A conceptual framework for assessing motivation and self-regulated learning in college students," *Educational psychology review*, vol. 16, no. 4, pp. 385–407, 2004.

[5] C. A. Wolters, "Understanding procrastination from a self-regulated learning perspective." *Journal of educational psychology*, vol. 95, no. 1, p. 179, 2003.

[6] V. Aleven, I. Roll, B. M. McLaren, and K. R. Koedinger, "Automated, unobtrusive, action-by-action assessment of self-regulation during learning with an intelligent tutoring system," *Educational Psychologist*, vol. 45, no. 4, pp. 224–233, 2010.

[7] G. Schraw, K. J. Crippen, and K. Hartley, "Promoting self-regulation in science education: Metacognition as part of a broader perspective on learning," *Research in science education*, vol. 36, no. 1-2, pp. 111–139, 2006.

[8] V. Aleven, B. Mclaren, I. Roll, and K. Koedinger, "Toward meta-cognitive tutoring: A model of help seeking with a cognitive tutor," *International Journal of Artificial Intelligence in Education*, vol. 16, no. 2, pp. 101–128, 2006.

[9] R. S. Baker, S. M. Gowda, and A. T. Corbett, "Towards predicting future transfer of learning," in *International Conference on Artificial Intelligence in Education*. Springer, 2011, pp. 23–30.

[10] E. National Academies of Sciences, Medicine *et al.*, *How people learn II: Learners, contexts, and cultures*. National Academies Press, 2018.

[11] C. Szabo, N. Falkner, A. Knutas, and M. Dorodchi, "Understanding the effects of lecturer intervention on computer science student behaviour," in *Proceedings of the 2017 ITiCSE Conference on Working Group Reports*, ser. ITiCSE-WGR '17. New York, NY, USA: Association for Computing Machinery, 2018, p. 105–124. [Online]. Available: https://doi.org/10.1145/3174781.3174787

[12] D. M. Hausman and B. Welch, "Debate: To nudge or not to nudge," *Journal of Political Philosophy*, vol. 18, no. 1, pp. 123–136, 2010.

[13] P. G. Hansen and A. M. Jespersen, "Nudge and the manipulation of choice: A framework for the responsible use of the nudge approach to behaviour change in public policy," *European Journal of Risk Regulation*, vol. 4, no. 1, pp. 3–28, 2013.

[14] M. T. Damgaard and H. S. Nielsen, "Nudging in education," *Economics of Education Review*, vol. 64, pp. 313–342, 2018.

[15] S. H. Edwards, J. Martin, and C. A. Shaffer, "Examining classroom interventions to reduce procrastination," in *Proceedings of the 2015 ACM Conference on Innovation and Technology in Computer Science Education*, ser. ITiCSE '15. New York, NY, USA: Association for Computing Machinery, 2015, p. 254–259. [Online]. Available: https://doi.org/10.1145/2729094.2742632

[16] H. C. Schouwenburg, "Academic procrastination," in *Procrastination and task avoidance*. Springer, 1995, pp. 71–96.

[17] J. W. You, "Examining the effect of academic procrastination on achievement using lms data in e-learning," *Journal of educational technology & society*, vol. 18, no. 3, pp. 64–74, 2015.

[18] D. Croft and M. England, "Computing with codio at coventry university: Online virtual linux boxes and automated formative feedback," in *Proceedings of the 3rd Conference on Computing Education Practice*, ser. CEP '19. New York, NY, USA: Association for Computing Machinery, 2019. [Online]. Available: https://doi.org/10.1145/3294016.3294018

[19] K. Ilves, J. Leinonen, and A. Hellas, "Supporting self-regulated learning with visualizations in online learning environments," in *Proceedings of the 49th ACM Technical Symposium on Computer Science Education*, 2018, pp. 257–262.