

Help Avoidance: When students should seek help,
and the consequences of failing to do so

Ma. Victoria Q. Almeda

Teachers College Columbia University

Ryan S. Baker

Teachers College Columbia University

Albert Corbett

Carnegie Mellon University

About the Authors

Ma. Victoria Q. Almeda
mqa2000@tc.columbia.edu
240-406-6275
517 W. 113th St. Apt 82, New York, NY, 10025

Victoria Almeda is a PhD student in Cognitive Studies in Education at Teachers College, Columbia University. Her research interests include student engagement, Intelligent Tutoring Systems, and math learning. Here is a sample of her recent work:

Godwin, K.E., Almeda, M.V., Petroccia, M., Baker, R.S., Fisher, A.V. (2013) Classroom activities and off-task behavior in elementary school children. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 2428-2433.

Almeda, M. V., Scupelli, P., Baker, R. S., Weber, M., & Fisher, A. (2014). Clustering of design decisions in classroom visual displays. In *Proceedings of the Fourth International Conference on Learning Analytics And Knowledge*, 44-48.

Ryan S. Baker
baker2@exchange.tc.columbia.edu
212-678-8329

Department of Human Development, Box 118, 525 W. 120th St., New York, NY
Ryan Baker is Associate Professor of Cognitive Studies and Program Coordinator for Learning Analytics at Teachers College, Columbia University. His research interests include student engagement, robust learning, meta-cognition, educational data mining, and learning analytics. His recent publications include:

Shute, V.J., D'Mello, S., Baker, R., Cho, K., Bosch, N., Ocumpaugh, J., Ventura, M., Almeda, V. (2015) Modeling how incoming knowledge, persistence, affective states, and in-game progress influence student learning from an educational game. *Computers & Education*, 86, 224-235.

Paquette, L., Baker, R.S., de Carvalho, A., Ocumpaugh, J. (2015) Cross-System Transfer of Machine Learned and Knowledge Engineered Models of Gaming the System. *Proceedings of the 22nd International Conference on User Modeling, Adaptation, and Personalization*, 183-194

Albert Corbett
corbett@cmu.edu
412-268-8808

Human-Computer Interaction Institute, 5000 Forbes Ave., Pittsburgh PA, 15232
Albert Corbett is an associate research professor, emeritus, in the Human-Computer Interaction Institute, Carnegie Mellon University. His principal research interest is in cognitive models of problem solving, and he has conducted extensive empirical evaluations of students learning

programming, mathematics, and genetics with intelligent computer tutors. He has recently published:

Corbett, A.T., Kauffman, L., MacLaren, B., Wagner, A., and Jones, E. (2010). A Cognitive Tutor for genetics problem solving: Learning gains and student modeling. *Journal of Educational Computing Research*, 42, 219-239.

The present article investigates whether there is an optimal point for determining whether a student needs help. Findings reveal that it is better for students within intelligent learning systems to seek help early within the learning process.

Abstract

Background: Across computer-based and traditional classroom settings, recent studies have identified motivational orientation, prior knowledge, self-regulation, and cognitive load as possible factors that impact help-seeking behaviors and their impact on learning. However, the question of whether there is an optimal point for determining when a student needs help has not been fully explored.

Purpose of Study: Using data from two modules of the Genetics Cognitive Tutor, the present study investigates this question by examining whether the relationship of help avoidance (failing to seek help when it is needed) and student learning is dependent on the student's level of prior knowledge. We also investigate how the relationship between help avoidance and student learning is mediated by the amount of prior practice, or the number of attempts at a problem step.

Research Design: We obtained existing data from the use of the Genetics Cognitive Tutor. We conducted a series of correlational analyses to better understand the relationship between help avoidance and student learning. We correlated students' proportions of help avoidance at different levels of knowledge with measures of robust learning. We also analyzed the relationship between students' proportions of help avoidance and measures of robust learning, taking the amount of practice or the number of attempts at a problem step into account.

Results: Our findings suggest that, except at very high or very low knowledge, help avoidance is generally stably (negatively) related to robust learning outcomes. Our results also indicate that help avoidance is more strongly associated with learning outcomes early in the practice sequence, suggesting that students should be encouraged to seek help on problem-solving skills on the first problem, rather than in waiting until later problems. Similarly, our results reveal that help avoidance is more negatively associated with learning outcomes on early attempts at a problem step than later attempts, indicating that students should be encouraged to seek help on the first attempt if help is needed.

Conclusions: These findings represent a step towards understanding when students should seek help, results with the potential to improve the design of meta-cognitive support within adaptive learning systems.

Keywords: help avoidance, help-seeking, Cognitive Tutor

Executive Summary

There have been concerns for several decades on why students fail to seek help when they are struggling, and what the consequences are of this choice. Students fail to seek help both in traditional learning settings, and when using adaptive learning systems. Previous research has determined that the failure to seek help when it is needed is associated with negative learning outcomes. However, previous studies have also identified specific factors (i.e., motivational orientation, prior knowledge, self-regulation, and cognitive load) that play a role in determining whether students to seek help more proactively. Within the context of intelligent tutoring systems, a number of efforts have also been made to build personalized support for effective help-seeking behaviors. However, these efforts have been unsuccessful at enhancing domain learning, possibly because previous research has not sufficiently delineated when students need help.

The focus of the present study is to fill in the gap within the existing literature and investigate when in the learning sequence is optimal for students to seek help. A student's level of knowledge has been used in the previous literature as the key factor for determining whether a student needs help and is avoiding it. As such, we investigate how the relationship between help avoidance and learning changes, as we change the threshold for student knowledge used to identify help avoidance. Additionally, we investigate how the relationship between help avoidance and learning changes, depending on the student's amount of prior practice on similar problem-solving steps or by the number of attempts the student has made on the current problem-solving step. The results of these research questions shed light on whether or not students should seek help immediately, or after a delay (either in terms of the number of problems solved, or after making a few attempts at the current problem step).

The data used for this paper were taken from students' use of the Genetics Cognitive Tutor, an intelligent tutoring system that supports students in learning key abductive reasoning skills in the domain of genetics. Specifically, we conducted a series of correlational analyses on datasets from two modules of the tutoring system covering 3-factor cross, a gene mapping technique that allows students to infer the order of three genes, and gene interaction, which engages students in reasoning about the various ways two genes interact to determine a single phenotypic trait. Participants were enrolled in genetics or introductory biology courses at Carnegie Mellon University, with 72 undergraduates using the three factor cross module and 52 undergraduates using the gene interaction module. All participants completed the following tests of robust learning: preparation for future learning (ability to acquire new knowledge based on existing knowledge), transfer (ability to use existing knowledge in new situations), and retention (ability to remember and apply knowledge one week later).

In analyzing the role of students' knowledge in the relationship between help-avoidance and learning, we first calculated the proportion of help avoidance for different knowledge thresholds to understand how common help avoidance was for various thresholds. Next, we

correlated students' proportions of help avoidance at different knowledge thresholds with the measures of robust learning. Patterns of results differed between the two modules. For gene-interaction, the correlation between help avoidance and learning declined sharply but then came back up somewhat, as students' knowledge increased. In contrast, for 3-factor cross, the correlation between help avoidance and robust learning went up briefly and then dropped sharply, as students' knowledge increased. Despite these differences, findings across datasets show a generally negative relationship between help-avoidance and robust learning, irrespective of students' knowledge.

Additionally, we analyzed how the correlation between the proportion of help avoidance and learning differed based on the student's amount of prior practice on the current problem-solving skill. In both datasets, the relationship between help avoidance and robust learning is relatively negative on the first practice opportunity for the current problem-solving skill. At the last practice opportunity for that skill, the relationship between help avoidance and robust learning is neutral or positive. In between these practice opportunities, patterns of results are more inconsistent across datasets, but generally negative. This suggests that the relationship between help avoidance and robust learning is unstable when differentiated in terms of the amount of prior practice. However, in general, it is more important to seek help on the first problem than on later problems.

Lastly, we compared the correlation between the proportion of help avoidance and robust learning between the first and subsequent attempts at a given step in a problem. Overall, the findings reveal that the negative relationship between help-avoidance and learning generally weakens by the third attempt. However, this shift was more evident in gene interaction than 3-factor cross. Further work is needed to explain the differences in the pattern of results between these two datasets. Overall, our findings reveal that help avoidance is more negatively correlated with robust learning in earlier attempts than later attempts, providing relevant implications for designing metacognitive support within adaptive learning systems. Specifically, for intelligent tutoring systems, it is important to encourage students to seek help during the student's first attempt at a problem-solving step. Given the gaps in the literature in this area, the present study represents a first step to a more conclusive understanding of the relationship between help avoidance and learning.

Help Avoidance: When students should seek help, and the consequences of failing to do so

Introduction

We all need help sometimes (Withers, 1972). We all have problems (in mathematics, genetics, or other domains) that someone else can help us to understand (Withers, 1972). And yet, many learners fail to seek help when it is needed to understand the learning material (Ryan et al., 1998; Butler, 1998; Aleven et al., 2006; Roll et al., 2014). A student's ability to effectively use help that is available is an essential skill to possess as a learner (Wood & Wood, 1999; Aleven et al., 2003); the reluctance or refusal to ask for or accept help when struggling is associated with poorer learning gains (Aleven et al., 2006; Baker, Gowda, & Corbett, 2011; Bartholomé et al., 2006).

The issue of help-avoidance has been present in classrooms for some time (Ryan et al., 1998; Butler, 1998; Newman, 1994; Newman & Goldin, 1990; Karabenick & Knapp, 1988) but may be even more pernicious in the case of students using computer-based learning technologies such as intelligent tutoring systems. Though a system may offer on-demand help facilities, these functions are not always utilized efficiently; many students neglect them completely (Aleven & Koedinger, 2000; Gräsel, et al, 2001; Renkl, 2002; Wood & Wood, 1999).

Most automated systems have limited capacity to deal with help avoidance. Whereas a classroom teacher can identify a struggling student and proactively offer help, the capacity of automated online systems for proactive help are limited. While some automated systems do offer proactive help to a struggling student (e.g. Murray & VanLehn, 2005), students do not always read or attend to these proactive help messages (Baker, Gowda, & Corbett, 2011). Thus, a student's individual choice to avoid help can have a larger effect in online learning than in traditional learning settings. To benefit more from online adaptive learning systems, students need to have the metacognitive abilities to understand and recognize when they need help on a question. The benefits of these systems can only be fully realized if students appropriately take advantage of these help functions. As such, intelligent tutoring systems need not only to promote domain learning, but also to promote the development of effective meta-cognition and skill for help-seeking (Roll et al., 2011). The development of these meta-cognitive skills depends on several individual factors of the learner. These factors include students' goals during learning, knowledge of how to use the learning environment, the ability to self-regulate one's own learning experience (Johnson, Archibald, & Tenebaum, 2010), and the potential cognitive load for processing information within tasks (Sweller, 1998, 1999).

Factors Impacting Help-Seeking

Motivational orientation. Previous research has shown a link between motivational orientation and help seeking (Ryan & Pintrich, 1997). For example, students with work-avoidance goals often seek help quickly (Butler, 2007). By contrast, students who adopt a performance-focused orientation, where the student attempts to demonstrate competence (Elliot

& Church, 1997), often avoid asking for help in fear that they will be viewed as less competent (Nelson-Le Gall, 1981; Ryan et al., 2001). Of particular relevance are performance-avoidance goals, where avoiding negative judgments of competence is likely to produce helpless patterns of responses (e.g., giving up in the face of failure) (Elliot & Church, 1997). Ryan and Pintrich (1997) found that students with learning goals reported being more likely to seek help, while those with performance goals reported that they were likely to avoid seeking help. Ryan and colleagues (1998) also found that students' self-reported help avoidance was higher when the classroom was perceived as performance oriented rather than learning oriented.

Another reason why students with performance goals may avoid help is out of a concern that they will receive less credit for a correct answer after asking for help. This concern is quite credible in modern online learning systems, which often penalize students for help requests. For example, the knowledge-tracing algorithm used in Cognitive Tutors treats help requests the same as errors (Corbett & Anderson, 1995). This is logical for knowledge inference, as help requests are (typically) evidence that the student does not know the skill. But at the same time, this means that the student's skill bar usually goes down when the student's first action at a problem-solving step is to seek help (subsequent actions at a problem-solving step do not affect the skill meter), an indicator that can be seen by the student, the teacher, and other students. Interview responses also indicate some students thought that their skill bars would continue to go down whenever they asked for further hint levels (Long & Aleven, 2011). As such, it is not surprising if some students may choose to avoid help to avoid this penalty.

Prior knowledge. Previous studies suggest that students' prior knowledge is a potential factor explaining the efficacy and frequency of help-seeking behaviors. For instance, Puustinen (1998) found that students with low prior knowledge were less effective help-seekers. These students did not seek help when they objectively needed it. Additionally, when they did seek help, their questions were aimed at confirming whether their answers were correct – rather than understanding the solution to the problem. Similarly, Wood and Wood (1999) found an interaction between students' prior knowledge and help-seeking behaviors. Findings reveal that students with high prior knowledge were more likely to ask for help after making an error, as compared to students with low prior knowledge. This suggests that high prior knowledge is associated with more accurate judgments of when to seek help.

Actively seeking help requires knowledge about how to use a system's help resources. Students who are less knowledgeable about a specific learning environment, or computer-based learning environments in general, tend to miss out on help opportunities because they do not know enough about the system to seek out these resources, while those who are more familiar with the system are able to use these resources in a more appropriate and effective manner (Hasebrook 1995). In some systems, help resources are clearly available; for example, Cognitive Tutors offer a large question-mark button for students to click (Roll et al., 2014). In other systems, however, help resources may be less obviously located. A student who does not know

how to seek help cannot do so. Therefore it is important that these learning environments provide instruction on the presence and use of help resources.

Self-regulation. A third core factor in whether students choose to seek help is self-regulation skill (Puustinen, Bernicot, & Bert-Erboul, 2011). Past theoretical accounts have argued that help seeking is deeply intertwined with self-regulated learning (Newman, 1998), since recognizing the need for help requires metacognitive and self-regulating abilities. One of the keys to successful self-regulated learning is attending to both internal and external feedback (Butler & Winne, 1995). Butler and Winne (1995) point out that self-regulated learners first provide themselves with internal feedback regarding the task. When a discrepancy exists between the learner's actual and desired performance, students must also seek external feedback from teachers and peers. The most effective self-regulated learners are those who attend to this external feedback (Bangert-Drowns, Kulik, Kulik, & Morgan, 1991; Kulhavy & Stock, 1989; Meyer, 1986).

Cognitive load. Many failures in self-regulating learning can also be tied to cognitive load. Learning gains can be compromised when attending to both the task itself and to the provided help leads to too much cognitive load (Sweller, 1988, 1999). When a student's cognitive resources are completely focused on the challenge of the activity, the student may not have the resources to recognize the need to ask for help. Wood (2001) suggests that cognitive load can have negative effects on self-regulatory skills of learners with lower prior knowledge especially. When problems are too difficult for a student, the ability to monitor one's comprehension can be compromised as a result. Wood's findings indicate that when problem difficulty was adjusted to match a student's own prior knowledge, there were no differences in effective help seeking between students with high or low prior knowledge. Other work has found that high achievers (based on class grade) were the best at self-regulation within help seeking activities (Puustinen, 1998).

Improving help-seeking behaviors in adaptive learning environments

In recent years, researchers have focused on using students' help-seeking behaviors to inform personalized instruction within adaptive tutoring environments.

For instance, there have been attempts to build online adaptive learning systems that encourage students to seek help more appropriately. Many intelligent tutoring systems provide various types of support functions intended to enhance learning, such as content-specific hints, hyperlinked textbooks, and online glossaries (see review by Alevan, Stahl, et al., 2003). Steps toward ensuring that students are seeking help sufficiently are also being made within intelligent tutoring systems. Roll and colleagues (2011) tested a system that offers metacognitive tutoring within the Geometry Cognitive Tutor learning environment. It includes self-assessment tutoring that helps students evaluate their own need for help, and help-use tutoring that offers feedback based on students' help seeking behavior. The results of one study show that these support

functions improved students' help seeking skill, inasmuch as the students behaved in a fashion closer to the prescriptive help model, however the support did not improve students' learning on the domain level (Roll et al., 2011). In a second study, the addition of an instructional video on help-seeking behaviors and support for self-assessment also showed similar results to the first study. Specifically, the additional support functions were found to facilitate the transfer of students' improved help seeking skills into new units of the Geometry Cognitive Tutor, but there were still no effects on domain learning. Roll and colleagues (2011) suggest that this may be due to excessive cognitive load imposed on students when they have to modify their help seeking strategies during problem solving tasks. It is also noted that while these metacognitive tutors can increase the probability that help is requested when it is needed, there is still no guarantee that the hint will be read by the student or that it will be understood. One limitation that may potentially explain the incomplete success of these help and metacognitive tutors is that we do not fully understand when students need help. The criterion for discriminating between needing help and not needing help within Alevan et al.'s (2006) model is the probability that the student knows the relevant skill, according to Bayesian Knowledge Tracing (BKT; Corbett & Anderson, 1995). BKT is successful at inferring student knowledge, predicting future student performance approximately as well as any other single algorithm (Pardos et al., 2011). The choice of knowledge to discriminate between students who need help and students who do not need help is reasonable. Less knowledgeable students clearly are more likely to need help than more knowledgeable students. In general, students with lower prior knowledge perform better when help is sought more often, while this does not hold true for those with high prior knowledge (Renkl, 2002; Wood & Wood, 1999). Given this pattern of results, it is possible that students with prior knowledge may overestimate their ability and process received help in a superficial fashion, as compared to students with low prior knowledge (Alevan et al., 2003).

However, the thresholds used in Roll et al.'s (2006) tutor are chosen by how well the models fit the data on student behavior; e.g. at the level of skill where students succeed on the first attempt if they do not receive help. This may not be the optimal skill level to suggest help at; for example, there may be a level of skill where students will not succeed immediately if they do not seek help, but will succeed eventually; alternatively, the conceptual support available in hints may be useful even if the student would get the problem correct without the hint.

In addition, there has been recent evidence that help-seeking may in some cases be more effective after first trying to solve a problem step without help. For instance, Zhu, Wang, and Heffernan (2014) found that students who make an attempt first before asking for a hint are more likely to try to figure things out for themselves. And in some cases, avoiding help can actually be associated with better learning, even when the student is attempting to answer problem steps where the student has low prior knowledge (Roll et al., 2014). As such, it appears that the relationship between help use and learning is more complicated than past theoretical models have suggested.

In this paper, we examine this issue more deeply, toward enhancing the support for help use in adaptive educational learning systems. Specifically, we investigate how the relationship between help avoidance and student learning is mediated by the definition of help avoidance. In particular, we start by looking at the role that estimates of student knowledge mastery play in defining whether a student's behavior involves help avoidance (Aleven et al., 2004, 2006). In Aleven et al.'s model (2004, 2006), whether or not a student needs help and should seek it is largely dependent on the student's degree of mastery of the skill involved in the current step of the tutor problem. Therefore, we analyze how changing the cut-off between situations where the student is viewed as needing help versus not needing help changes the relationship between help avoidance and learning. This will help us to understand whether there is an optimal point for deciding that a student needs help, and where this point lies. We will also look at whether the relationship between help avoidance and student learning is mediated by how many attempts the student makes on a problem-solving step, and how much prior practice the student has had on the relevant skill. Is it really best for students to immediately ask for help, as the design of Aleven et al.'s model suggests? Or should students sometimes try to solve a problem step first?

We investigate these issues within the context of data from the Genetics Cognitive Tutor (Corbett et al., 2010). Like all Cognitive Tutors, the Genetics Tutor provides fine-grained high quality data on student interactions. It is particularly advantageous to study these issues within the Genetics Tutor because it also has carefully crafted post-test measures of robust learning (Corbett et al., 2011). According to the KLI framework (Koedinger, Corbett, & Perfetti, 2012), robust learning is learning that is retained over time, transfers to new situations, and prepares students for future learning (cf. Bransford & Schwartz, 1999). By examining the relationship between help avoidance and robust learning (instead of immediate performance on the exact skills learned), we can work towards developing findings that can be used to develop adaptive learning systems that optimize long-term student success rather than short-term performance.

Method

Data Sets

As mentioned above, the data set used in the analyses presented here came from the Genetics Cognitive Tutor (Corbett et al., 2010). This tutor consists of 19 modules that support problem solving across a wide range of topics in genetics. Various subsets of the 19 modules have been piloted at 15 universities in North America. In this paper, we analyze data from two studies, conducted in successive years, involving lessons that were popular across many universities and generally successful at promoting robust learning. The first study focuses on a tutor module that employs a gene mapping technique called *three-factor cross*, in which students infer the order of three genes on a chromosome based on offspring phenotypes, described in Baker, Corbett, et al. (2010). The second study focuses on a tutor module on *gene interaction*, in which students reason about the various ways two genes can interact to determine a single phenotypic trait.

The participants in each study were enrolled in genetics or in introductory biology courses at Carnegie Mellon University; 72 undergraduates used the three factor cross module in the first study and 52 undergraduates used the gene interaction module in the second study. In both studies, the students engaged in Cognitive Tutor-supported activities for one hour in each of two sessions on successive days. In the first study all students completed standard three-factor cross problems in both sessions and in the second study all students completed gene interaction problems in both sessions. During the first session in each study, some students were assigned to complete other learning activities (e.g., worked examples and problem-solving) designed to support deeper understanding prior to the standard tutor problems; however, no differences were found between conditions for any robust learning measure, so in this analysis we collapse across the conditions and focus solely on student behavior and learning within the standard problem-solving activities.

As with many previous studies on student help avoidance in online learning (Alevin et al., 2006; Roll et al. 2011, 2014; Baker, Gowda, & Corbett, 2011), Bayesian Knowledge Tracing (Corbett & Anderson, 1995) is used to operationally identify student knowledge. Bayesian Knowledge Tracing computes the probability that a student knows a given skill at a given time, combining data on the student's performance up to that point with four model parameters. Bayesian Knowledge Tracing achieves comparable performance to other methods that infer whether students have learned skills in online systems (see review in Pardos et al., 2011), and has been shown to predict student post-test problem solving performance reasonably well within the Genetics Tutor (Baker et al., 2010).

Within these tutor lessons, student robust learning was measured using tests of three aspects of robust learning: preparation for future learning (PFL), transfer, and retention. PFL measures the ability to acquire new knowledge more quickly or effectively, based on existing knowledge. Transfer taps into students' understanding of underlying processes, requiring students to use their existing knowledge in new situations or fashions. Lastly, retention is measured through students' delayed performance, one week later in these studies, which demonstrates knowledge retained over time.

Results

Analyzing Knowledge's Role in the Relationship Between Help Avoidance and Learning

The key factor used to determine whether a student is avoiding help in Alevin et al.'s (2006) model of help-seeking (beyond whether the student asked for help) is student knowledge, which is measured using Bayesian Knowledge Tracing. Within that model, if student knowledge of a relevant skill for a problem-solving step is below a certain threshold, and the student fails to seek help and instead gets the wrong answer, the student is thought to be avoiding help. However, what is the appropriate threshold? Different papers have selected different thresholds,

ranging from 0.60 (Alevén et al., 2004) to 0.95 (Alevén et al., 2006). Presumably, a student with a 95% chance of knowing a skill should attempt it without asking for help, but what about a student with an 80% chance? A 60% chance? A 40% chance? In the graphs below, we compute the proportion of help avoidance using every potential knowledge threshold from 0.05% to 100% using a grain-size of 0.05% (e.g. 0.05%, 0.10%, 0.15%...100%).

Our first step is to look at how the proportion of help avoidance changes as the cut-off changes, where the proportion is calculated as the number of cases where knowledge is below threshold AND the student gets an incorrect answer without first asking for help, divided by the total number of cases where the student's knowledge is below the threshold. In doing this, we filtered out skills that were encountered multiple times per problem; in these tutor modules these skills generally represented basic execution of procedures rather than the more cognitively challenging parts of the domain. Figure 1 shows that as the knowledge threshold goes up within the gene interaction lesson, the proportion of help avoidance gradually drops, with a little more drop at 20% and at the highest levels of student knowledge. The graph for 3-factor cross slopes more downward, as shown in Figure 2, with the proportion of help avoidance slightly declining at 25%, then at 60%, and finally dropping again for the highest levels of student knowledge. However, there is no threshold where help avoidance is extremely rare, suggesting that any threshold is in principle valid for analysis.

<insert Figure 1 here>

<insert Figure 2 here>

As such, we can correlate each student's proportion of help avoidance – for each knowledge threshold – to their performance on the three aforementioned tests of robust learning. To do this, the proportion of help avoidance is computed for each knowledge threshold and student, and then the proportions of help avoidance are correlated to student performance on the three tests of robust learning. Pearson correlations are used as a simple indicator of the relationship between the proportion of help avoidance and each of the test measures. Since Pearson correlations are being used in an exploratory fashion rather than within statistical significance tests, normality assumptions are not checked. Violation of normality for this context can be expected to lead to correlation values closer to zero. Graphs showing these correlations are given in Figures 3 and 4.

<insert Figure 3 here>

<insert Figure 4 here>

There are noticeable differences for the graphs across datasets. First of all, for gene interaction, the correlations are near zero for the lowest threshold, with the exception of retention – but even in this case the correlation is still much closer to zero than for other knowledge thresholds. This makes sense for extremely low values, where there may be very little data.

Bayesian Knowledge Tracing assumes an initial knowledge level and best-fitting estimates of this initial probability are typically well above zero. The estimates of a specific student's knowledge can drop lower than the initial probability when the student makes incorrect answers, but it is uncommon for the estimate of a student's knowledge to drop down to 5% even after multiple incorrect answers. However, correlations for the lowest threshold across the measures of robust learning are more negative for 3-factor cross than for gene interaction, with values slightly less than -0.3.

In addition, the trend lines before the first inflection point also vary between data sets. In the case of gene interaction, as the knowledge threshold goes up, the correlation between help avoidance and robust learning becomes increasingly negative until it reaches the first inflection point. In this case, all graphs hit the first inflection point at a knowledge threshold of 0.15, with the following correlations between help avoidance and each measure of robust learning: PFL: -0.43, transfer, -0.43, and retention: -0.57.

The opposite pattern is seen for 3-factor cross. The correlation between help avoidance and robust learning becomes increasingly positive until it reaches the first inflection point. For the 3-factor cross lesson, first inflection points also occur at a knowledge threshold of 0.15. At this knowledge threshold, the correlation between each measure of robust learning and help avoidance is as follows: transfer: -0.18, PFL: -0.19, retention: -0.25. It is worth noting that at this inflection point, all three correlations are weaker within the 3-factor cross data set than in the gene interaction data set. It is also worth noting that help avoidance is much more strongly associated with retention than with the other measures of robust learning, a finding we have no a priori explanation for.

After the first inflection point is reached, the trend-lines also differ in their patterns as knowledge cut-offs go up. For gene interaction, the transfer and PFL graphs show a downward trend, then a slight peak in the graph, before becoming relatively flat. However, the values are fairly consistent, suggesting that that this relationship is robust to knowledge cutoff for all knowledge levels except very low knowledge, at least within the gene interaction lesson. In contrast, for 3-factor cross, these graphs show a slight peak at the cut-off of 45%, before generally trending downwards. The peak in the transfer graph occurs at the cut-off of 45% with a correlation of -0.23; then the correlation becomes more negative until it reaches a cut-off of 60% with a correlation of -0.39. In the case of the PFL graph, the peak occurs at the same point as the transfer graph at the cut-off of 45% with a correlation of -0.29, but then the correlation drops, reaching a more negative correlation of -0.42 at the cut-off of 50%.

After the first inflection point, the graphs for retention swing up to a greater degree than PFL or transfer, but then drop more sharply as well. In gene interaction, retention at one point (cut-off of 0.45) has a similar correlation to PFL and transfer, before becoming much more negative than the other robust learning indicators by the right side of the graph (e.g. cut-off of 1). In 3-factor cross, retention's correlation becomes less negative than PFL and transfer (cut-off

above 0.55) before again becoming more negative at the right side of the graph (e.g. cut-off approaching 1).

All of these findings suggest that the relationships between help avoidance and robust learning are somewhat unstable depending on the cut-off. On the other hand, regardless of data set, robust learning measure, or cut-off, the relationship between help avoidance and robust learning remains negative within these Genetics Cognitive Tutor modules.

Help Avoidance and the Amount of Prior Practice

As discussed above, Alevén et al.'s (2006) model parametrizes help avoidance in terms of the level of student knowledge. But it may also be relevant to consider the amount of prior practice a student has had. For example, the first time a student sees an entirely new skill, it may be more reasonable to try the problem-solving step before seeking help rather than immediately asking for help. By the time the student has had several opportunities to apply a skill across problems, however, there is an increasing chance that a student who does not know the skill is “wheel-spinning” (Beck & Gong, 2013), struggling with no chance of resolving the difficulty on their own, and in need of immediate assistance. As such, in this section we investigate the relationship between students' help avoidance and robust learning, in relation to the amount of practice.

For the remaining analyses, we used a knowledge cut-off of 0.60 in our definition of help avoidance. This value was selected for two reasons. First, it was the value used in the original version of Alevén's help-seeking model (2004), helping make the results in the remainder of the paper consistent with that work. Second, this cut-off captured a part of the space in both data sets where there was relatively little change in correlations between help avoidance and robust learning. Since any cut-off value would still capture the same negative overall relationship, choosing a relatively stable part of the space allows us to focus on other potential influences on the relationship. We studied this relationship for each of the first 7 (for gene interaction) or 8 (for 3-factor cross) times a skill was encountered by the student, focusing (as noted above) on skills encountered a single time by the student in each problem. The graphs for these correlations are provided in Figures 5 and 6.

<insert Figure 5 here>

<insert Figure 6 here>

The graphs for robust learning vary between datasets. Notably, the difference between the first practice opportunity and the second practice opportunity is different between data sets. Within gene interaction, the correlation for all three robust learning measures starts below zero, but trends towards zero on the second attempt. Within 3-factor cross, this relationship is inconsistent between the three learning measures.

After the second opportunity, in gene interaction, the relationship trends downwards for transfer and PFL, before generally trending back upwards, eventually pulling above 0 for both these measures on the 7th practice opportunity. By contrast, the graph for retention increases fairly steadily across the student's use of the tutor.

Within 3-factor cross, the correlation between help avoidance and retention generally trends downward after the second opportunity, going as low as -0.41, before trending back upward after the 4th opportunity. The relationship between help avoidance and PFL drops after the 3rd opportunity, and then approaches 0 for later opportunities. The relationship between help avoidance and transfer is generally unstable over time. All three robust learning metrics approach a correlation of 0 on the last practice opportunity, with the correlation between help avoidance and transfer pulling above 0 on the last practice opportunity (the same pattern as seen in gene interaction).

It is not entirely clear what these patterns mean. In both data sets, help avoidance is associated with negative outcomes on the first practice opportunity. By the last practice opportunity, the relationship between help avoidance and robust learning approaches zero. In between, the relationships are somewhat more complicated, and more unstable than when help avoidance was differentiated by the level of student knowledge.

Help Avoidance: First Attempts and Subsequent Attempts

Another factor we can consider in terms of help-avoidance is how many attempts the student has made on the current step within a problem. As we discuss above, Aleven et al.'s (2004, 2006) model does not make prescriptions about how many times the student should get something wrong before asking for help. Should students sometimes try a step before asking for help? Or is it better to ask for help right away? To investigate this, we analyze the relationship between help avoidance and robust learning for help avoidance occurring on a student's first, second, and third attempt at a problem-solving step. The second attempt at a problem step, as defined here, follows an error, and the third attempt follows two errors. (If the student had already sought help, it would not be correct to treat an incorrect attempt as representing help avoidance). As in the previous analysis, we use a knowledge cut-off of 60% to define help avoidance for this analysis, and filter skills as in that analysis. Refer to Figures 7 and 8 for the graphs showing these correlations.

<insert Figure 7 here>

<insert Figure 8 here>

In examining the correlation between help avoidance and learning for first and subsequent attempts, we can see that the relationship is generally negative on the first attempt for both data sets, in keeping with previous results (and the other analyses reported in this paper).

However, there are differences in these trend-lines across data sets when we compare later attempts (e.g., second and third attempts) to earlier attempts (e.g., first attempts). For gene interaction, the trend lines are closer to zero for later attempts than the first attempt (with less change for PFL than for the other measures of robust learning). For 3-factor cross, the second attempt is broadly similar to the first attempt, but all three learning graphs slope upwards substantially after the second attempt. Specifically, graphs for retention and PFL cross over into having a positive correlation, while the graph for transfer approaches zero. The key commonality across data sets is that the first attempt is the most informative about eventual student achievement.

Discussion

There has been considerable research on help-seeking and help avoidance over the last decade. However, many of the models of help seeking rely upon relatively ad-hoc parameters for inferring when a student needs help. As such, in this paper, we investigate whether there is an optimal point for deciding when a student needs help. We investigated this research question by examining the relationship between definitions of how much knowledge the student needs, to not need help, and whether the relationship between help avoidance and robust learning would differ according to these definitions. However, our findings suggest that, aside from extremely low levels of knowledge, this relationship is reasonably resilient to changes in the definition of how much knowledge the student needs; there are shifts across the graph, but values remain solidly negative with the exception of extremely low knowledge in gene interaction (which may be, as discussed above, an artifact of limited data).

We then investigated whether the relationship between help avoidance and learning is mediated by the amount of prior practice, or the number of attempts a student makes on a problem-solving step. In terms of the influence of the amount of prior practice, our findings suggest that the correlation between help avoidance and learning is stably negative for the first practice opportunity, and that the correlation is generally close to zero for the last practice opportunity. In between, patterns are more unstable across data sets and robust learning measures. This is surprising; one might expect that as the student begins to wheel-spin, repeatedly getting a problem-solving skill wrong and not seeking help (Beck & Gong, 2013), help avoidance would become more harmful rather than less harmful. However, it seems that help avoidance matters more, earlier in the learning sequence. This shift is stronger for 3-factor cross, where the correlations reach or pass 0, than for gene interaction. As such, help avoidance on later attempts is more of a problem in gene interaction than 3-factor cross. It is not yet clear whether this pattern of results is due to differences between learning environments, differences between populations, or some other factor. Additionally, our results seem to suggest that students who need help will be best served by immediately seeking help, rather than attempting to solve the problem on their own first.

The influence of the number of attempts a student makes on a problem-solving step on the relationship between help avoidance and robust learning is more stable than the amount of prior practice. For both datasets, the relationship between help avoidance and robust learning starts strong and negative, and then generally weakens for later attempts.

These findings, taken together, suggest that the general focus in past research on looking at help avoidance primarily on first attempts at a problem step is generally reasonable, and that help avoidance models are likely to be robust to differences in knowledge thresholds, except for extreme values. However, the results argue that not all practice opportunities are the same; help avoidance matters more at the beginning of learning than at the end.

While these findings help us to better understand when help is helpful, they have important implications for the future design of meta-cognitive support for help in adaptive learning systems. For instance, it appears that it is more important to encourage learners to seek help on first attempts than later, and that it is more important to encourage learners to seek help on the first problem than later, especially compared to the last problem.

More broadly, these findings are additional evidence that the general approach of attempting to encourage students to seek help when needed (e.g. Roll et al., 2011) is appropriate. In general, our findings underscore the importance of providing support for students in learning – recommending that they should seek help early within the learning sequence, and early in the process of tackling a difficult problem.

Most attempts to do so thus far have been reported to be unsuccessful at improving learning, but one possibility is that they have measured the wrong outcomes. Perhaps future attempts to encourage help should be considered not solely in terms of improving immediate domain learning, the primary construct measured in past research, but instead in terms of improving the robustness of learning. It is possible that a simple replication of past research on help-seeking support that measured robust learning (and particularly retention) would produce much stronger evidence for impact.

However, there remain several questions for future investigation. For example, given that help avoidance is more strongly associated with negative learning outcomes during earlier practice opportunities for skills, it may be worthwhile to investigate whether certain types of instructional support within adaptive learning systems are more beneficial during the beginning of the learning sequence. Specifically, it may be helpful for students to receive types of hints that elicit reflective thinking early on, to leverage the benefit of this type of learning experience at a point where students seem to generally be more receptive to help. Further work is needed to understand how instructional design within adaptive learning systems can possibly moderate the relationship between help avoidance and learning.

An additional question comes from our observation that the relationships we studied differ somewhat between the two lessons studied. In general, it would be valuable to extend the

research seen here by obtaining robust learning measures for a range of different content, or even learning environments, and seeing which aspects of tutor design and content impact the relationships between help avoidance and robust learning. Doing so would enable us to more conclusively understand this relationship in its full complexity. The work presented here is but one step in that direction.

Acknowledgements

We would like to thank the generous support of the National Science Foundation, under grant #DRL-0910188 and grant #SBE-0836012 to the Pittsburgh Science of Learning Center. We would also like to thank Angela Wagner for support with data collection.

References

- Aleven, V., & Koedinger, K. R. (2000). Limitations of Student Control: Do Student Know when they need help? In G. Gauthier, C. Frasson, & K. VanLehn (Eds.), *Proceedings of the 5th International Conference on Intelligent Tutoring Systems, ITS 2000* (pp. 292-303). Berlin: Springer Verlag.
- Aleven, V., McLaren, B., Roll, I., & Koedinger, K. (2004). Toward tutoring help seeking: Applying cognitive modeling to meta-cognitive skills. In J. C. Lester, R. M. Vicario, & F. Paraguaçu (Eds.), *Proceedings of Seventh International Conference on Intelligent Tutoring Systems, ITS 2004* (pp. 227-239). Berlin: Springer Verlag.
- Aleven, V., McLaren, B., Roll, I., & Koedinger, K. (2006). Toward meta-cognitive tutoring: A model of help seeking with a Cognitive Tutor. *International Journal of Artificial Intelligence in Education, 16*(2), 101-128.
- Aleven, V., Stahl, E., Schworm, S., Fischer, F., & Wallace, R. (2003). Help seeking and help design in interactive learning environments. *Review of Educational Research, 73*(3), 277-320.
- Baker, R. S., Corbett, A. T., Gowda, S. M., Wagner, A. Z., MacLaren, B. A., Kauffman, L. R., ... & Giguere, S. (2010). Contextual slip and prediction of student performance after use of an intelligent tutor. In *User Modeling, Adaptation, and Personalization* (pp. 52-63). Springer Berlin Heidelberg.
- Baker, R., Gowda, S. M., & Corbett, A. T. (2011). Automatically detecting a student's preparation for future learning: Help use is key. In *Proceedings of the 4th international conference on educational data mining*, 179-188.
- Bangert-Drowns, R. L., Kulik, C. L. C., Kulik, J. A., & Morgan, M. (1991). The instructional effect of feedback in test-like events. *Review of educational research, 61*(2), 213-238.

Bartholomé, T., Stahl, E., Pieschl, S., & Bromme, R. (2006). What matters in help-seeking? A study of help effectiveness and learner-related factors. *Computers in Human Behavior*, 22(1), 113-129.

Beck, J. E., & Gong, Y. (2013, January). Wheel-Spinning: Students Who Fail to Master a Skill. In *Artificial Intelligence in Education* (pp. 431-440). Springer Berlin Heidelberg.

Bransford, J. D., & Schwartz, D. L. (1999). Rethinking transfer: A simple proposal with multiple implications. *Review of research in education*, 61-100.

Butler, R. (1998). Determinants of help seeking: Relations between perceived reasons for classroom help-avoidance and help-seeking behaviors in an experimental context. *Journal of Educational Psychology*, 90, 630-643.

Butler, R. (2007). Teachers' achievement goal orientations and associations with teachers' help seeking: Examination of a novel approach to teacher motivation. *Journal of Educational Psychology*, 99(2), 241-252.

Butler, D. L., & Winne, P. H. (1995). Feedback and self-regulated learning: A theoretical synthesis. *Review of educational research*, 65(3), 245-281.

Corbett, A.T., & Anderson, J., R. (1995). Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling & User-Adapted Interaction*, 4(4), 253-278.

Corbett, A., Kauffman, L., Maclaren, B., Wagner, A., & Jones, E. (2010). A Cognitive Tutor for genetics problem solving: Learning gains and student modeling. *Journal of Educational Computing Research*, 42(2), 219-239.

Corbett, A., MacLaren, B., Wagner, A., Kauffman, L., Mitchell, A., Baker, R. & Gowda, S. (2011). Preparing students for effective explaining of worked examples in the Genetics Cognitive Tutor. In *Proceedings of the 33rd annual meeting of the Cognitive Science Society*, 1476-1481.

Elliot, A. J., & Church, M. A. (1997). A hierarchical model of approach and avoidance achievement motivation. *Journal of Personality and Social Psychology*, 72(1), 218-232.

Grasel, C., Fischer, F., & Mandl, H. (2001). The use of additional information in problem-oriented learning environments. *Learning Environments Research*, 3, 287-305.

Hasebrook, J. (1995). *Multimedia-Psychologie: Eine neue Perspektive menschlicher Kommunikation* [Multimedia psychology: A new perspective of human communication]. Heidelberg: Spektrum.

Johnson, T. E., Archibald, T. N., & Tenenbaum, G. (2010). Individual and team annotation effects on students' reading comprehension, critical thinking, and meta-cognitive skills. *Computers in Human Behavior*, 26(6), 1496-1507.

- Karabenick, S. A., & Knapp, J. R. (1988). Help-seeking and the need for academic assistance. *Journal of Educational Psychology, 80*, 406–408.
- Koedinger, K. R., Corbett, A. T., & Perfetti, C. (2012). The Knowledge-Learning-Instruction framework: Bridging the science-practice chasm to enhance robust student learning. *Cognitive Science, 36*(5), 757-798.
- Kulhavy, R. W., & Stock, W. A. (1989). Feedback in written instruction: The place of response certitude. *Educational Psychology Review, 1*(4), 279-308.
- Long, Y., & Alevan, V. (2011, January). Students' Understanding of Their Student Model. In *Artificial Intelligence in Education* (pp. 179-186). Springer: Berlin-Heidelberg.
- Meyer, L. A. (1986). Strategies for correcting students' wrong responses. *The Elementary School Journal, 227*-241.
- Murray, R.C., & VanLehn, K. (2005). Effects of dissuading unnecessary help requests while providing proactive help. In the *Proceedings of the 12th International Conference on Artificial Intelligence in Education*, 887–889.
- Nelson-Le Gall, S. (1981). Help-seeking: An understudied problem-solving skill in children. *Developmental Review, 1*, 224-246.
- Newman, R. (1994). Adaptive help seeking: A strategy of self-regulated learning. In Schunk, D., and Zimmerman, B. (eds.), *Self-Regulation of Learning and Performance: Issues and Educational Applications*, Lawrence Erlbaum Associates, Hillsdale, NJ, pp. 283–301.
- Newman, R. (1998). Adaptive help seeking: A role of social interaction in self-regulated learning. In Karabenick, S. (ed.), *Strategic Help-Seeking: Implications for Learning and Teaching*, Erlbaum, Hillsdale, NJ, pp. 13–37.
- Newman, R. S., & Goldin, L. (1990). Children's reluctance to seek help with homework. *Journal of Educational Psychology, 82*, 92–100.
- Pardos, Z.A., Baker, R.S.J.d., Gowda, S.M., Heffernan, N.T. (2011) The Sum is Greater than the Parts: Ensembling Models of Student Knowledge in Educational Software. *SIGKDD Explorations, 13* (2), 37-44.
- Puustinen, M. (1998). Help-seeking behavior in a problem-solving situation: Development of self-regulation. *European Journal of Psychology of Education, 13*(2), 271-282
- Puustinen, M., Bernicot, J., & Bert-Erboul, A. (2011). Written computer-mediated requests for help by French-speaking students: An analysis of their forms and functions. *Learning and Instruction, 21*(2), 281-289.

Renkl, A. (2002). Learning from worked-out examples: Instructional explanations supplement self-explanations. *Learning and Instruction, 12*, 529-556.

Roll, I., Alevan, V., McLaren, B. M., & Koedinger, K. R. (2011). Improving students' help-seeking skills using metacognitive feedback in an intelligent tutoring system. *Learning and Instruction, 21*(2), 267-280

Roll, I., Alevan, V., McLaren, B. M., Ryu, E., Baker, R. S., & Koedinger, K. R. (2006, January). The Help Tutor: Does Metacognitive Feedback Improve Students' Help-Seeking Actions, Skills and Learning?. In *Intelligent Tutoring Systems* (pp. 360-369). Springer Berlin-Heidelberg.

Roll, I., Baker, R.S.J. d., Alevan, V., Koedinger, K.R. (2014) On the Benefits of Seeking (and Avoiding) Help in Online Problem-Solving Environments. *Journal of the Learning Sciences, 23* (4), 537-560.

Ryan, A. M., Gheen, M. H., & Midgley, C. (1998). Why do some students avoid asking for help? An examination of the interplay among students' academic efficacy, teachers' social-emotional role, and the classroom goal structure. *Journal of Educational Psychology, 90*(3), 528-535.

Ryan, A. M., & Pintrich, P. R. (1997). "Should I ask for help?" The role of motivation and attitudes in adolescents' help seeking in math class. *Journal of Educational Psychology, 89*(2), 329-341.

Ryan, A. M., Pintrich, P. R., & Midgley, C. (2001). Avoiding seeking help in the classroom: Who and why? *Educational Psychology Review, 13*(2), 93-114.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science, 12*, 257-285.

Sweller, J. (1999). *Instructional design in technical areas*. Melbourne: ACER Press.

Withers, B. (1972) *Lean on Me. Still Bill*. Los Angeles, CA: Sussex Records.

Wood, D. (2001). Scaffolding, contingent tutoring, and computer-supported learning. *International Journal of Artificial Intelligence in Education, 12*, 280-292.

Wood, H., & Wood, D. (1999). Help seeking, learning and contingent tutoring. *Computers and Education, 33*(2), 153-169.

Zhu, L., Wang, Y., & Heffernan, N. T. (2014). The Sequence of Action Model: Leveraging the Sequence of Attempts and Hints. In the *Educational Data Mining Workshop held at the 7th international conference on Educational Data Mining*.

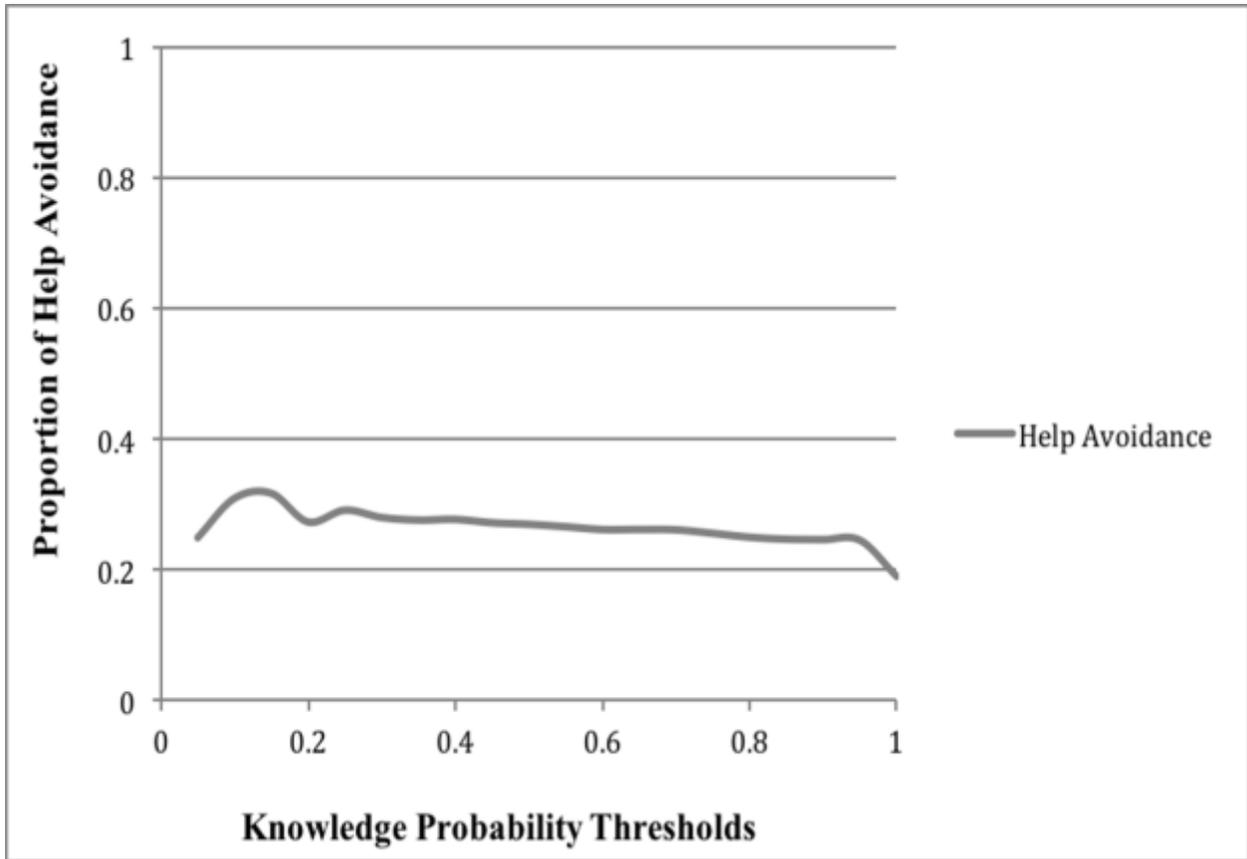


FIGURE 1. The proportion of help avoidance for different knowledge thresholds, for gene interaction

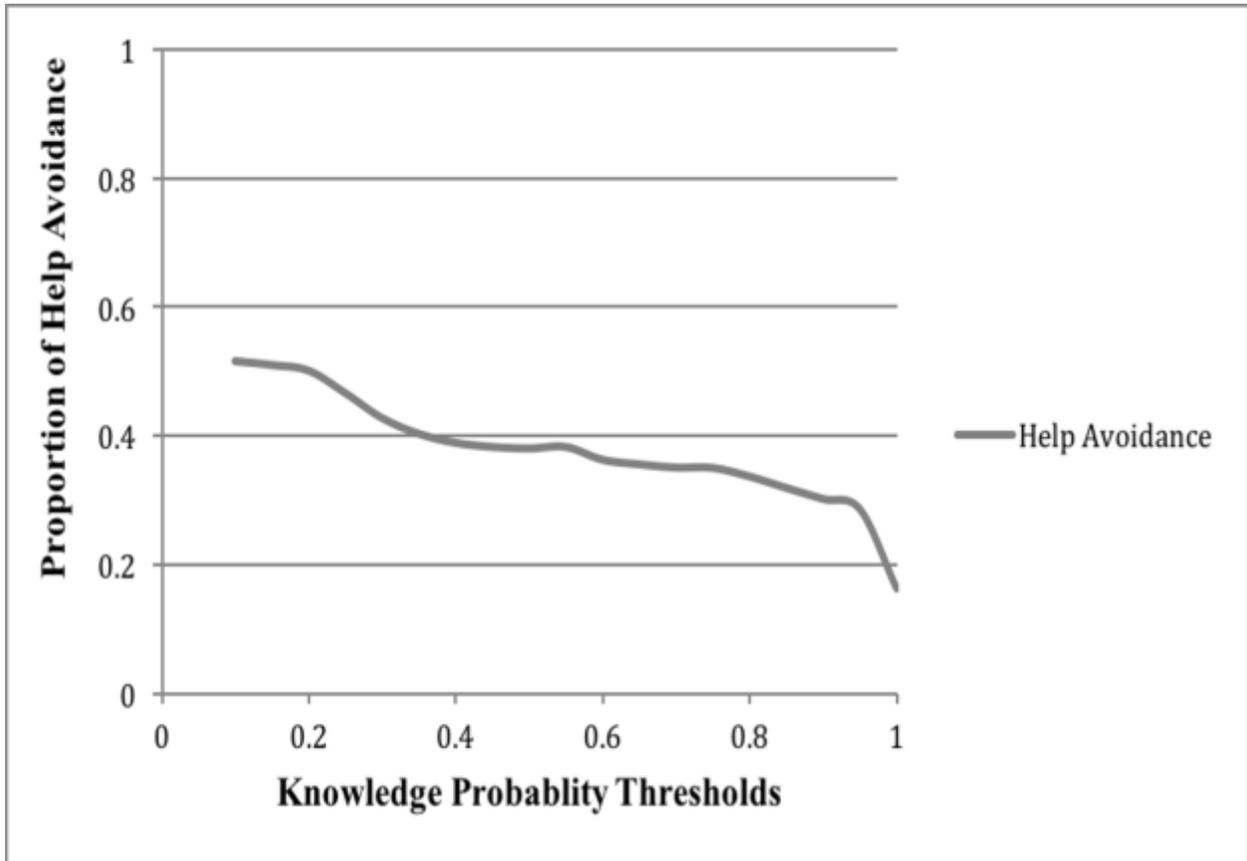


FIGURE 2. The proportion of help avoidance for different knowledge thresholds, for 3-factor cross

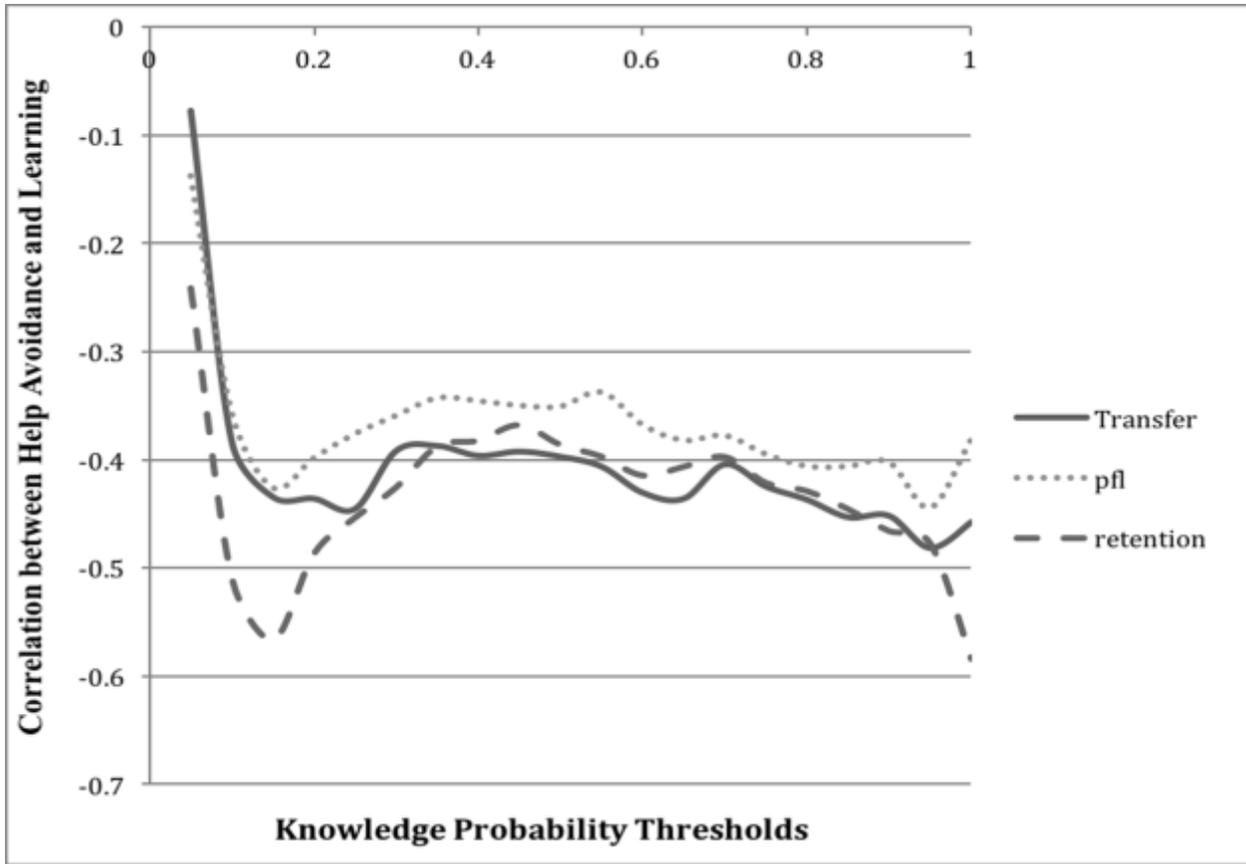


FIGURE 3. The correlation between help avoidance and robust learning for different knowledge thresholds, for gene interaction

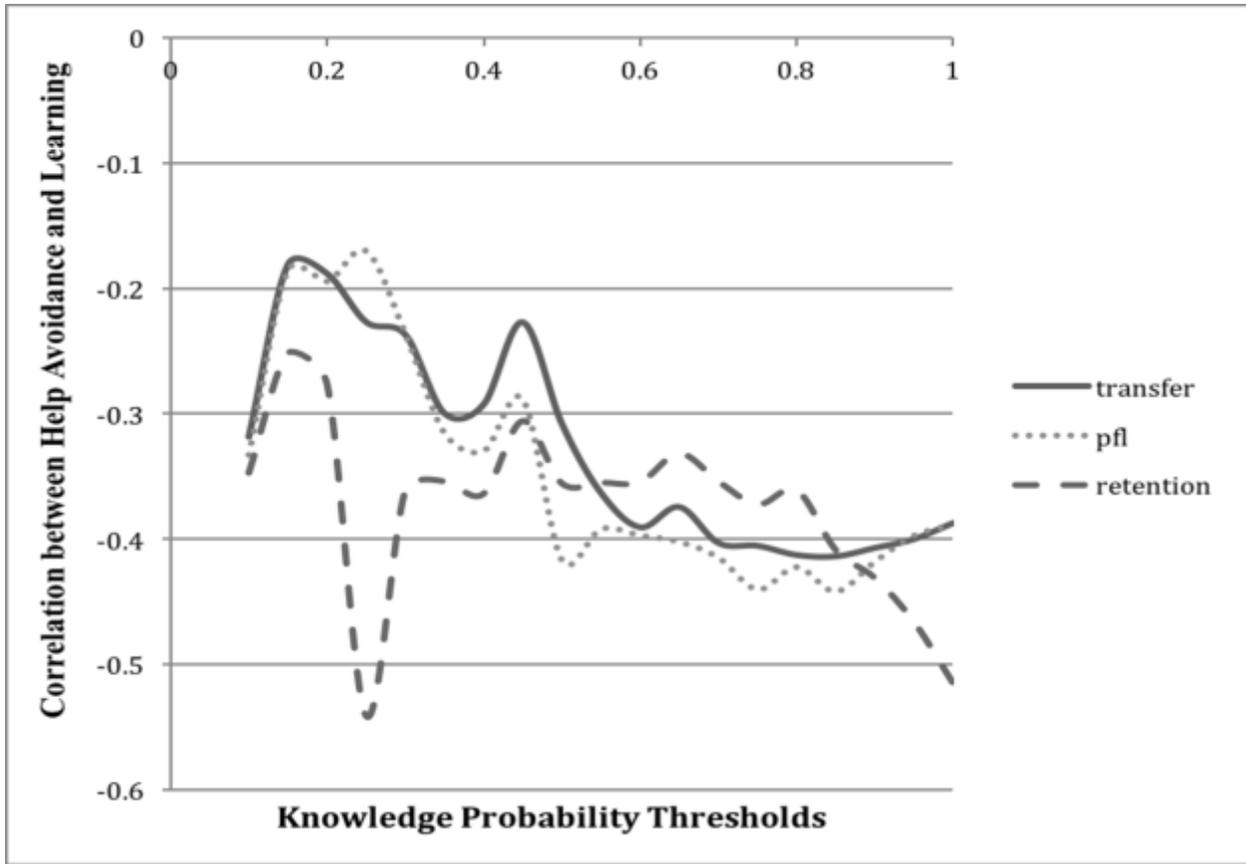


FIGURE 4. The correlation between help avoidance and robust learning for different knowledge thresholds, for 3-factor cross

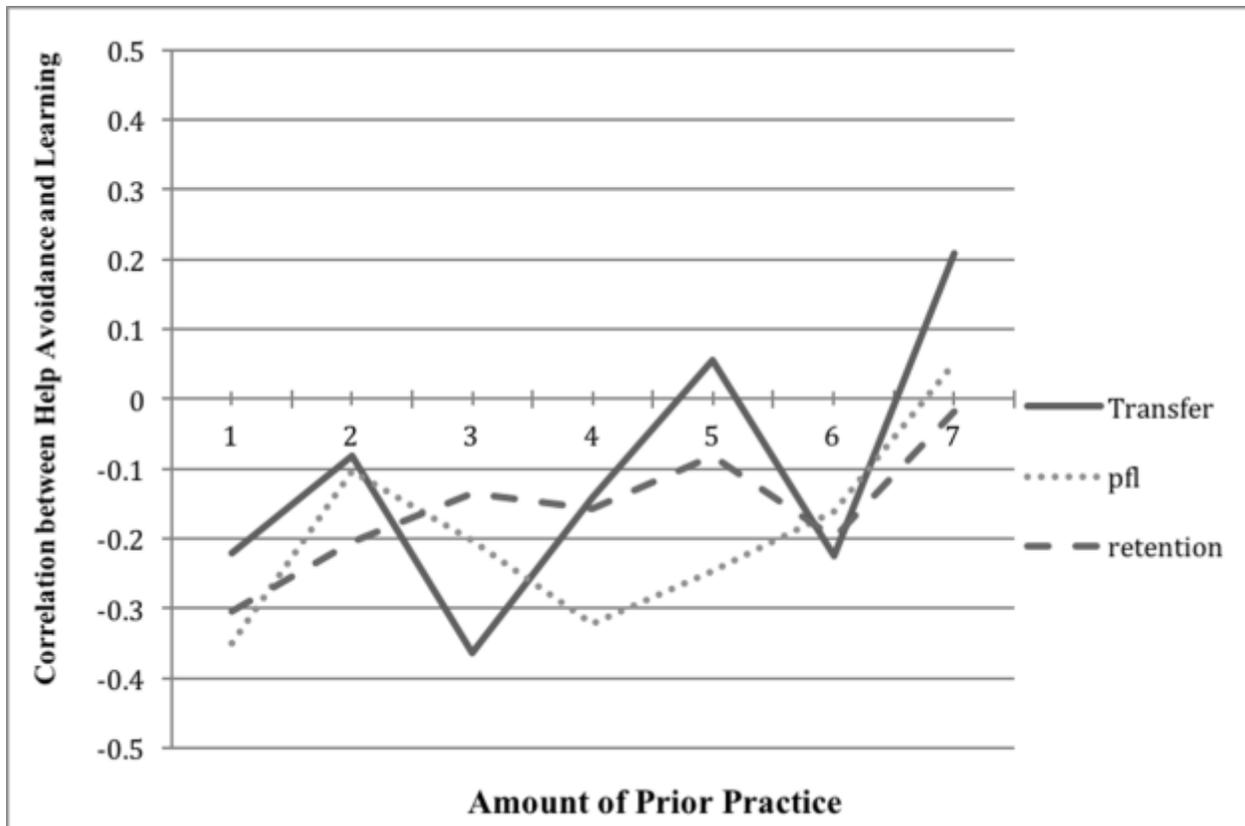


FIGURE 5. The correlation between help avoidance and robust learning for the amount of prior practice, for gene interaction

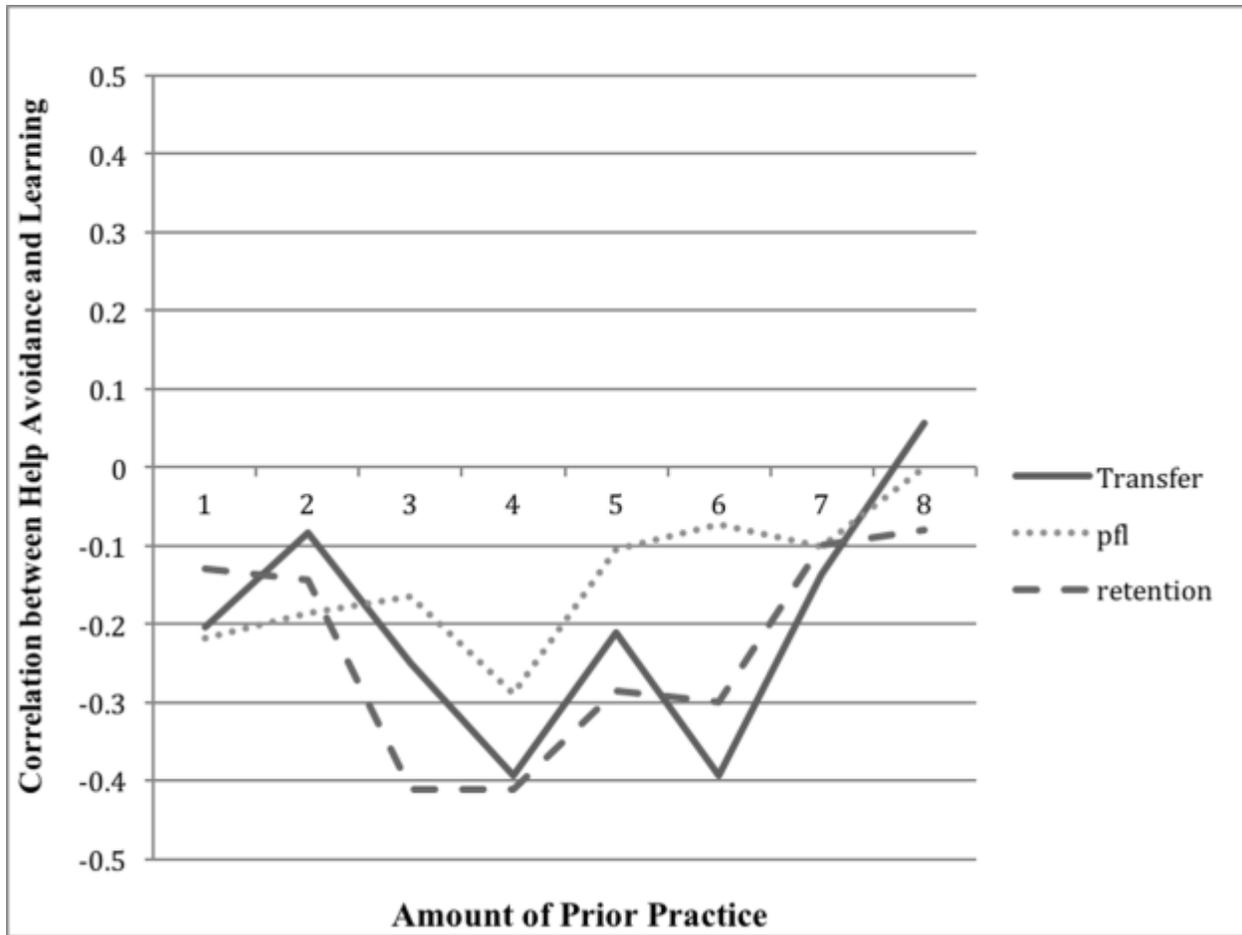


FIGURE 6. The correlation between help avoidance and robust learning for the amount of prior practice, for 3-factor cross

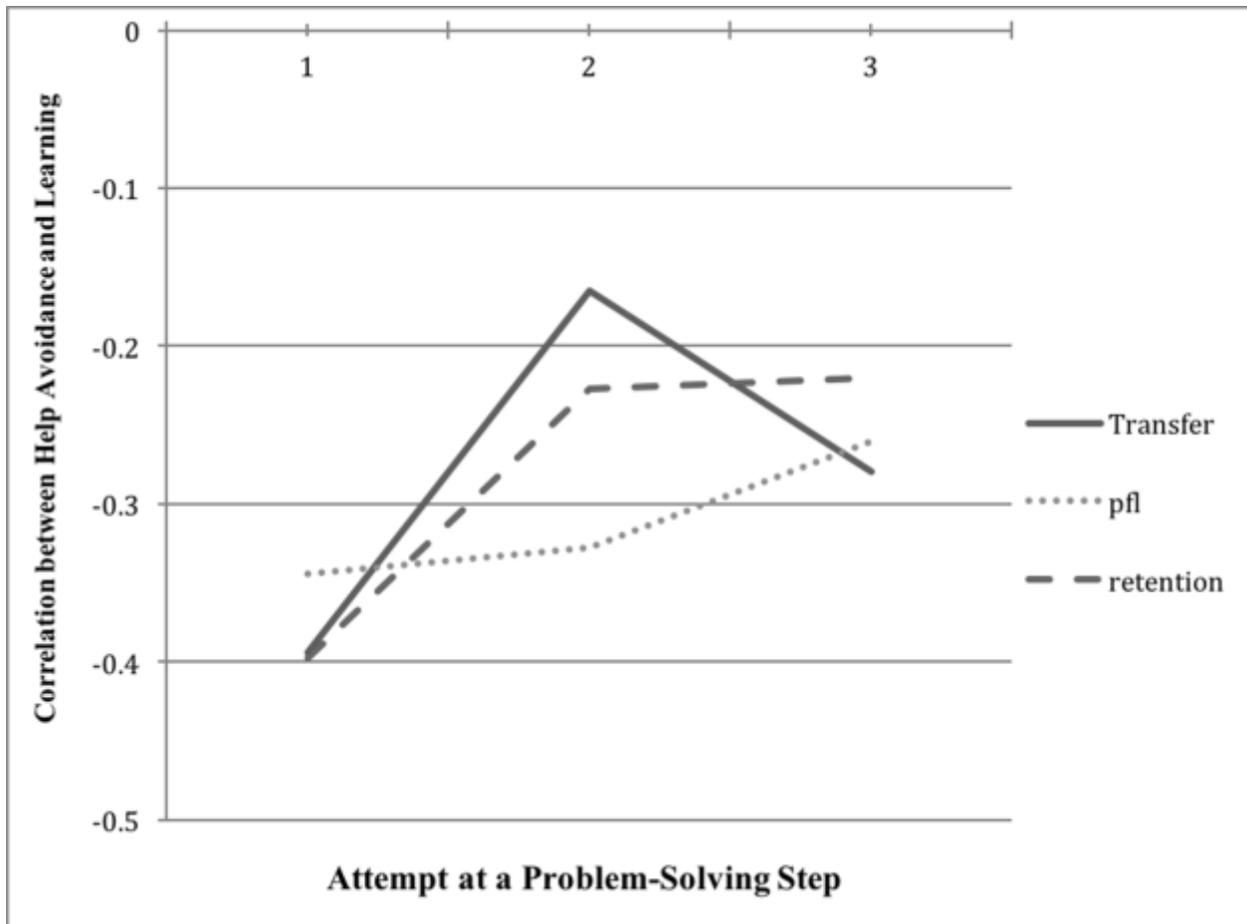


FIGURE 7. The correlation between help avoidance and robust learning, by how many attempts the student has made on the problem step (without having already sought help), for gene interaction

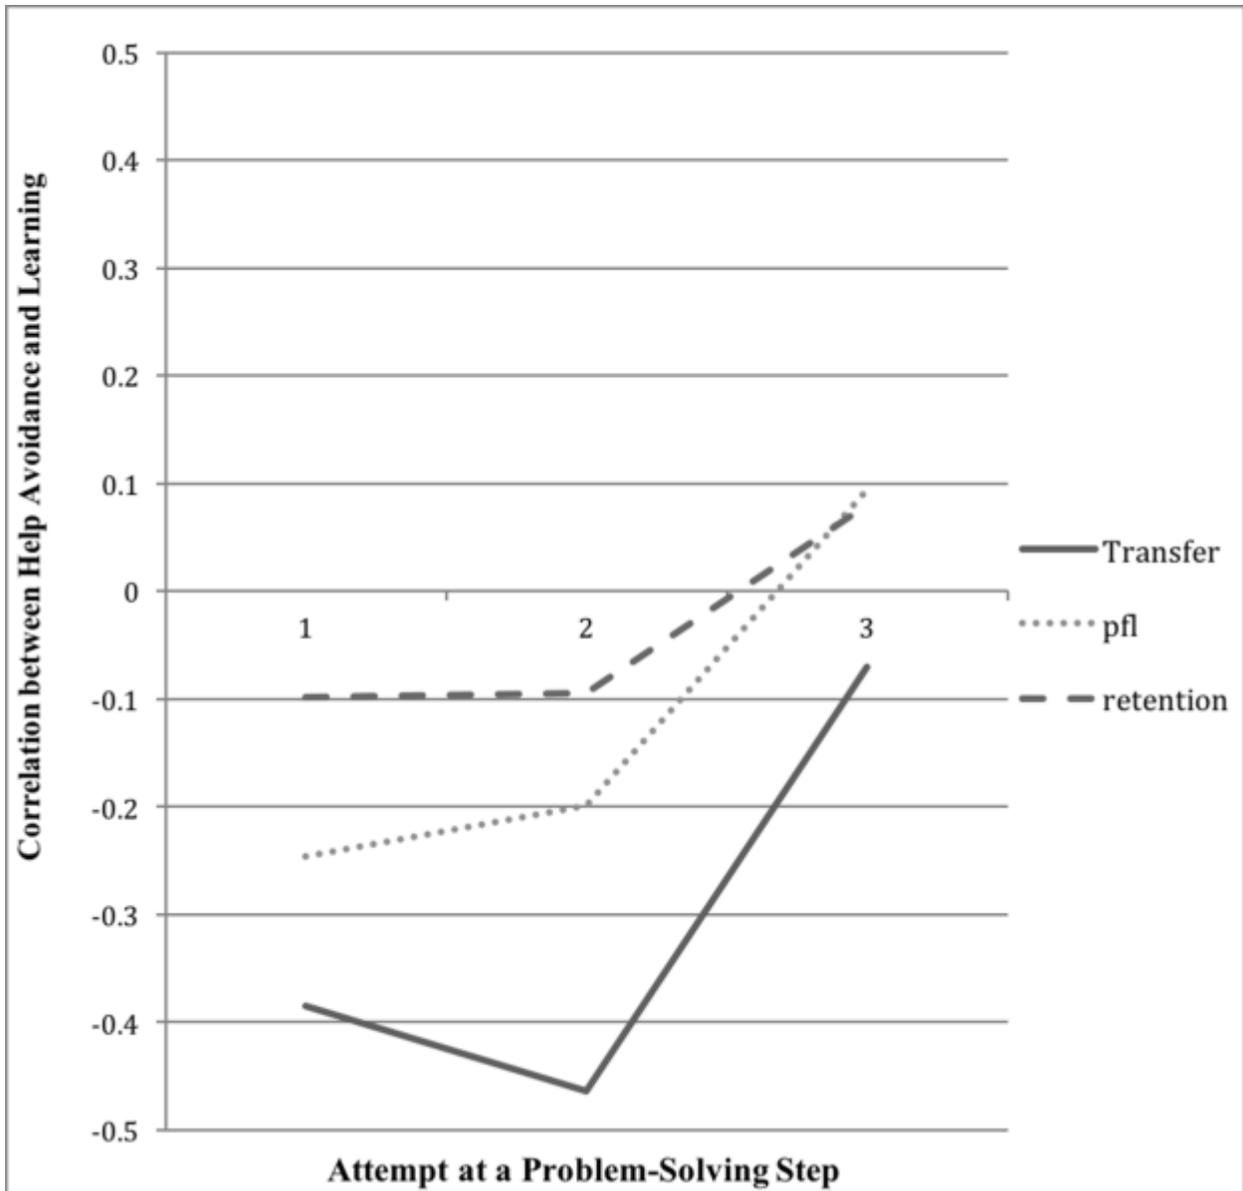


FIGURE 8. The correlation between help avoidance and robust learning, by how many attempts the student has made on the problem step (without having already sought help), for 3-factor cross