Belenky, D.M., & Nokes, T.J. (2010). Optimizing learning environments: An individual difference approach to learning and transfer. In S. Ohlsson & R. Catrambone (Eds.), *Proceedings of the 32nd Annual Conference of the Cognitive Science Society* (pp. 459-464). Austin, TX: Cognitive Science Society.

Optimizing Learning Environments: An Individual Difference Approach to Learning and Transfer

Daniel M. Belenky (dmb83@pitt.edu)
Timothy J. Nokes (nokes@pitt.edu)
Learning Research and Development Center
University of Pittsburgh, 3939 O'Hara Street
Pittsburgh, PA 15260

Abstract

Prior work has found that the type of learning activity (direct instruction or invention) interacts with achievement goals (mastery or performance-oriented) such that invention tasks can help facilitate mastery goal adoption and knowledge transfer (Belenky & Nokes, 2009). In the current study, we investigated how robust the effect is, and whether explicit manipulations of the task goals can produce a similar effect. We conducted an experiment with 98 college students in which achievement goals were measured, while task goals and task structure were manipulated between subjects. Results indicated that task structure was generally a more effective way of influencing which achievement goals are adopted within a learning activity. However, task goals that promoted an evaluative context interfered with transfer for mastery-oriented learners from invention activities. The results are interpreted in relation to theories of regulatory fit and multiple goal hierarchies.

Keywords: learning; transfer; skill acquisition; motivation; achievement goals.

Student's achievement goals have a large influence on their behaviors and experiences in academic settings. The literature surrounding Achievement Goal theory shows that these goals lead to very different patterns of affect, interest and achievement (e.g., Harackiewicz et al., 2005). However, this literature has not focused on how the goals influence what is learned. That is, although "achievement" is frequently measured as an outcome, it is almost always done at a coarse-grain level, such as final grades in a course. It is not clear how different achievement goals (mastery versus performance) are related to different kinds of learning, such as learning procedural skills, simple facts, or conceptual knowledge.

To begin to address this gap, Belenky & Nokes (2009) examined how achievement goals impact the type of knowledge gained from different kinds of instruction. That study found that mastery-oriented learners do better on transfer measures, regardless of whether the mastery-orientation came from a stable predisposition or whether the open-ended structure of an "invention" task led to mastery-oriented feelings and goals in the specific context. Conversely, performance-oriented learners did better on skill acquisition when the instruction seemed to match their goals, by presenting a well-structured, simple task through direct instruction.

This initial work has provided evidence that task structure interacts with existing achievement goals to influence learning. In the current work we examine whether direct manipulations of task goals through instructions can change the ways students learn, similar to the effect of task structure. If directly manipulating task goals produces similar effects, it would offer a more direct way of encouraging students towards desired learning outcomes (whether towards transfer or skill). However, it is possible that achievement goals within a learning activity are not under conscious control, and task structure has more influence on how a student engages than instructions that attempt to prompt a particular achievement goal. It is also possible that task structure and task goals operate independently, leading to a three-way interaction in the adoption of achievement goals based on students' prior dispositions. This study explores these possibilities.

Background

Research on achievement goals has focused on two main aims; classifying what the goals are and then correlating those with predictors and outcomes. The prevailing classification is a 2 x 2 framework that has been well-validated (Elliot & McGregor, 2001). This framework separates the evaluative criterion (mastery or performance) from the valence (approach or avoidance), which results in four separable goals (mastery-approach, mastery-avoidance, performance-approach, performance-avoidance). Mastery goals refer to ones in which a person is basing his evaluation on the skill or competence he is trying to develop (that is, in comparison to an expectation or prior ability), while performance goals refer to evaluating oneself based on a normative standard (that is, in comparison to others). Approach goals refer to seeking out positive outcomes, while avoidance goals refer to averting negative ones. For example, a mastery-approach goal is one in which a person is seeking to improve his ability or knowledge, based on an internally-referenced criterion ("My aim is to completely master the material in this class"), while a performance-avoidance goal is one in which a person is seeking to not look bad compared to others ("My aim is to avoid doing worse than other students;" see Elliot & McGregor, 2001). Students can have different levels of each of these goals, and studies have validated that these four goals are separate factors (Elliot & McGregor, 2001). Because we are most interested in studying how different paths of *successful* learning affect what knowledge is gained, our work focuses on mastery-approach and performance-approach goals.

Mastery-approach (MAP) goals have been correlated with a host of positive outcomes, such as intrinsic motivation, interest, better self-regulation, and deeper strategy use.

However, MAP goals are generally unrelated to achievement scores (usually operationalized as exam scores or final grades). Performance-approach (PAP) goals have been correlated with some positive outcomes, such as perseverance, task engagement, enjoyment and topic interest, as well as some negative outcomes, such as test anxiety, poor help-seeking and, most importantly for the current study, shallow cognitive strategies. PAP goals are also positively correlated with achievement scores (see Elliot, 1999 for a review). One potential reason that MAP goals are unrelated to achievement scores but PAP goals correlate positively could be due to the type of knowledge being assessed on achievement measures. Scores on a final exam, for example, may reflect a person's ability to recall factual information (a performance-oriented outcome) more than a deep conceptual understanding (a mastery-oriented outcome).

Belenky & Nokes (2009) examined the possibility that different types of knowledge were being generated due to students' existing achievement goals. This work found that MAP goals led to more flexible use of knowledge on a transfer assessment, while PAP goals led to better procedural skill with a formula. That study also showed that an ill-structured, invention-based learning task promoted mastery-goal adoption. This goal adoption was particularly beneficial for those who entered the study low in MAP goals; this group performed one standard deviation better than low-MAP counterparts who completed a well-structured, direct instruction-based approach.

The paradigm and tasks used in the current work are based on previous work on "Preparation for Future Learning" (Belenky & Nokes, 2009; Schwartz & Martin, 2004). This paradigm allows one to measure a person's ability to use knowledge acquired in one situation to help learn in a new situation. The work that established this paradigm contrasted two types of learning activities, direct instruction and invention (Schwartz & Martin, 2004). The direct instruction condition was similar to a well-structured, "tell-and-practice" style of pedagogy, where a student is shown a method and asked to use it to solve similar problems. The invention activity was modeled on a form of "discovery learning" where students are asked to construct their own knowledge in an open-ended, ill-structured problem. All students were given a subsequent transfer problem on an extension of the concept. Half of the students from each of the group were given a learning resource (a worked example) embedded in the assessment, while half were not (see Figure 1). They found that only those students who had invented and were given the worked example showed large improvements in the ability to solve a transfer problem. Belenky & Nokes (2009) found evidence that this benefit was due, at least in part, to the adoption of different goals. Support for this view was based on a questionnaire given right after the learning activity showing that students given invention activities were more concerned about their understanding of the concepts than those given tell-and-practice activities, as well as a benefit in transfer performance among those students who entered the study low in MAP goals but performed invention activities as opposed to those who performed tell-and-practice activities.

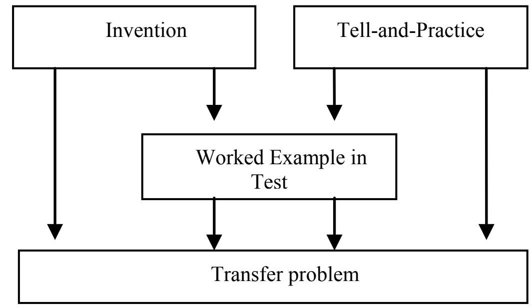

Figure 1. The general Preparation for Future Learning paradigm which measures the ability to transfer from an initial learning experience to a later one.

The current study further explores the issue of how achievement goals influence and are influenced by different task structures, by using the same tasks and adding task instructions to try and directly manipulate the achievement goals students adopt. We have evidence that goal adoption, whether due to individual dispositional differences or task structure, influences the form of the knowledge gained from instruction. The robustness of this effect is being examined, by adding the factor of task instructions to the framework.

Hypotheses

We are investigating how explicit, task-based goal instructions influence the effects of existing achievement goals and task structure on learning and transfer. We predict that the task structure will have a larger effect than the task-based instructions on which achievement goals students adopt and their impact on learning and transfer. However, there are several possibilities for interesting interactions between task structure, task goals, and existing achievement motivations. We are most interested in exploring whether there is evidence for a "multiple-goals" framework (Harackiewicz et al., 2002) or a "regulatory fit" viewpoint (Higgins, 2005).

The multiple goals hypothesis states that a mix of performance-approach and mastery-approach goals is optimal for learning (Harackiewicz et al., 2002). If this hypothesis is correct then we would expect that crossing mastery and performance goals would promote the best overall learning. Alternatively, a "regulatory fit" view states that the alignment of task structure, task-based goals and dispositional orientations should facilitate the best learning (Higgins, 2005). This suggests that matching students' goals with instructions that support pursuit of those goals would improve learning, (e.g., mastery-oriented achievement goals with invention tasks and mastery-oriented instructions). If this is the case, we would also expect to see that incongruous goals would harm learning. We will examine these possibilities by assessing three-way interactions between existing goals, task structure and instructions on different learning outcomes.

We also predict replications of our basic prior findings that invention activities promote more mastery-oriented learning behaviors and feelings than tell-and-practice. We also predict that existing mastery-approach goals will lead to better transfer, while performance-approach goals will lead to better performance on a skill measurement.

Methods

This study closely followed the methods and materials of prior studies (Belenky & Nokes, 2009; Schwartz & Martin, 2004). Participants completed a pre-test, went through a series of learning activities on basic statistical concepts, and then took a post-test in a 2-hour laboratory session. They were given questionnaires at the beginning, middle, and end of the experiment.

Participants

Ninety-eight undergraduates from the University of Pittsburgh participated in this study (*M* = 19.4 years old, *SD* = 2.5 years) in exchange for course credit.

Design and Materials

This study had a 2 (task structure: invention or tell-and-practice) x 2 (task goal: mastery-oriented or performance-oriented instructions) between-subjects design. Materials were presented as packets in binders. These packets consisted of an initial questionnaire; a pre-test; learning activities (with an activity questionnaire after the first one); a post-test; and final questionnaires, including a demographic sheet.

Learning Materials

The learning materials consisted of one activity on mean deviation, instruction on the correct calculation of mean deviation (a narrated PowerPoint video), and then a new activity on standardization. The first learning activity presented four different data sets representing the spread of a number of pitches thrown by different pitching machines. The students' task was to decide which of the machines is the most reliable. The invention and tell-and-practice students both attempted this problem, but the instructions each received was different. The tell-and-practice group was given a worked example demonstrating how to calculate mean deviation immediately prior to attempting to solve this problem. The invention group was not given this worked example, and was instructed "Your task is to invent a procedure for computing a quantity that expresses the variability for each of the pitching machines and decide which is most reliable. There is no single way to do this, but you have to use the same procedure for each machine, so it is a fair comparison." Both groups were given access to scrap paper and a calculator during this activity.

The video demonstrated a brief introduction on variability before walking through the calculation of mean deviation in a worked example. This was followed by two simple problems to work on, with solutions demonstrated to make sure students understood the basics of the formula.

The structure of the standardization activity was similar to the mean deviation activity. The invention group was asked to evaluate which of two world records was "more shattered," given two small data sets and the corresponding world record for each. The tell-and-practice group was given the exact same problem to solve, but was first shown how to graphically arrive at a solution, by divvying up a visual representation of the distributions through a worked example.

We manipulated task goals through instructions. Immediately below the tell-and-practice or invention instructions, participants saw instructions about the purpose of the task. These were constructed to spur participants towards either a mastery or performance orientation while working on the learning activities. The motivation goal instructions were modeled on previous work that had experimentally manipulated goals (i.e., Elliott & Dweck, 1988; Elliot & Harackiewicz, 1996). The mastery goal instructions were: "Many people see problems like this one as a challenge, and feel like they are developing their skill to solve these types of problems. While working on this problem, you may make mistakes and feel a little confused at times, but in the end you will have learned some things and developed your skill to solve problems like this one." The performance goal instructions were: "This problem assesses your mathematical ability. People who can solve this problem generally have the capability to solve similar problems. While working on this activity, you can gauge how good you are at these types of math problems." These were presented underlined, to make them more salient.

Test Materials

The pre-test consisted of three items: a skill measure, a data representation problem, and a transfer problem. The post-test contained isomorphic versions of these problems, as well as an adaptive use and a qualitative reasoning problem. We will only focus on the transfer and skill measures here. The skill measures presented small data sets and explicitly asked the participants to calculate mean, mode, median and mean deviation. The transfer problems were both word problems that presented descriptive statistics for two data distributions and one exceptional score from each, and asked which individual score was more impressive. While similar to the standardization activity, this problem assesses transfer because it requires reasoning from descriptive statistics, not raw data, and because simpler heuristic processing would lead to an incorrect answer (i.e., reasoning about range, or because one value is higher than the other).

The skill measure was scored dichotomously as a 0 if incorrect, while a 1 was awarded if the student flawlessly used the formula. All other problems were coded as a 0 if incorrect, 1 if conceptually correct but there was a computational error, and 2 if the answer was conceptually and computationally correct. For the transfer problem, this meant calculating standardized scores and using them to decide which value was more impressive.

The test also included a worked example on how to calculate a standardized score. This was presented just like the other test problems, and described a situation in which one would want to calculate a standardized score to compare values from different samples. The text then introduced the formula to do so and computed the values for the data presented. This was followed by a second, very simple data set, and asked the students to use the formula on these data. The worked example always came at least two problems before the transfer problem, so if a student used the formula on the transfer problem, it was because they noticed that it applied and could recall it, not due to temporal contiguity.

Motivation Measures

To assess achievement goals, we used the Achievement Goal Questionnaire (AGQ; Elliot & McGregor, 2001). This 12-item scale has 3 items for each of the 4 achievement goal constructs. We focus only on mastery-approach and performance-approach goals in this study. The questions were phrased to be specifically about math classes, and were assessed on a 7-point Likert scale. Cronbach's alphas for each of the scales was high (MAP = .839, PAP = .932). There was also an activity questionnaire, which was the same as that used in Belenky & Nokes (2009). This 8-item measure asked about student's focus and affect while working on the first learning activity. At the end of the study, the participants completed additional questionnaires – the AGQ again, and a questionnaire we developed to assess how students reflected on and our goal manipulations.

Procedure

The study was run in groups of up to six participants in a two-hour laboratory session, with all participants working individually. Inside the packet was: an initial questionnaire; a pre-test; a learning activity; the activity questionnaire; space to work on problems presented in the video; another learning activity; a post-test; a final set of questionnaires; a demographics sheet. Participants took as long as they needed to complete the questionnaires, with no one taking longer than three minutes for any one questionnaire. Both learning activity and the video took fifteen minutes each. Participants were given five minutes for each test item.

Results

Our hypotheses focused on the conceptual replication of earlier findings, as well as exploring how explicitly manipulating task goals and task structure would influence learning. First, we assess whether task structure and instructions influenced students' self-reported experiences during the learning activity. Then, we examine the competing "regulatory fit" and "multiple-goals" hypotheses by looking at the interactions of existing goals, learning tasks, and instructions.

Motivational effects on questionnaires

On the activity questionnaire, administered directly after the first learning problem, there were telling differences between responses from the invention and tell-and-practice groups (see Table 1). Namely, students who invented felt more concerned with understanding, challenged, and frustrated, Fs (1, 94) > 6.87, ps < .01. While frustration is generally considered a negative outcome, it may signal to a student a lack of understanding that drives further cognitive engagement. Notably, there were no differences due to our task goal manipulation, nor interactions. This is evidence that that the structure of the task led to very different experiences for the learner, both in terms of affective response and goals adopted, and that this structure has more influence on goal adoption than explicit instructions about the purpose of the task.

We also asked a similar set of reflective questions at the end of the experiment. The only observed significant differences were that students who completed invention tasks reported being able to be more creative, while those who received mastery instructions reported enjoying the activity more. There was also an interaction between the experimental manipulations on the item "it was a challenge to come up with the correct solution." This interaction was driven by the lower ratings on this item provided by those who completed the tell-and-practice activities with performance- oriented task goal. This manipulation seems to have led students to feel less challenged, even compared to the students who completed the same activity but were told that the task could help them improve. The results from the reflective questionnaire seem to suggest that manipulations of goals do not produce large differences in students' conscious, reflective experience of learning.

Interactions

We predicted that we would see three-way interactions between existing achievement goals, task structure, and task goals, specifically looking at performance-approach goals (PAP) when considering skill acquisition and mastery-approach goals (MAP) for transfer. To assess the effect of existing achievement goals, we performed a median split within each experimental group on the performance-approach and mastery-approach construct scores, based on the initial questionnaire, administered at the start of the experiment. We used the mean deviation skill measurement as the dependant variable in a 2 (task structure) x 2 (task

Table 1. Activity questionnaire items with differences between Tell-and-Practice (TP) and Invent. † $p < .1$, * $p < .05$.

Questionnaire Item	*Invent*	*TP*	*Effect Size d*
I was concerned with how well I understood the procedure I was using.*	3.7 (1.2)	3.0 (1.4)	0.54
I was concerned that the procedure I was using was not correct.*	4.0 (1.1)	2.5 (1.4)	1.17
I felt engaged. †	3.1 (1.3)	3.6 (1.0)	-0.37
I felt frustrated.*	3.5 (1.4)	2.6 (1.3)	0.65
I felt challenged. *	3.7 (1.1)	2.8 (1.2)	0.86

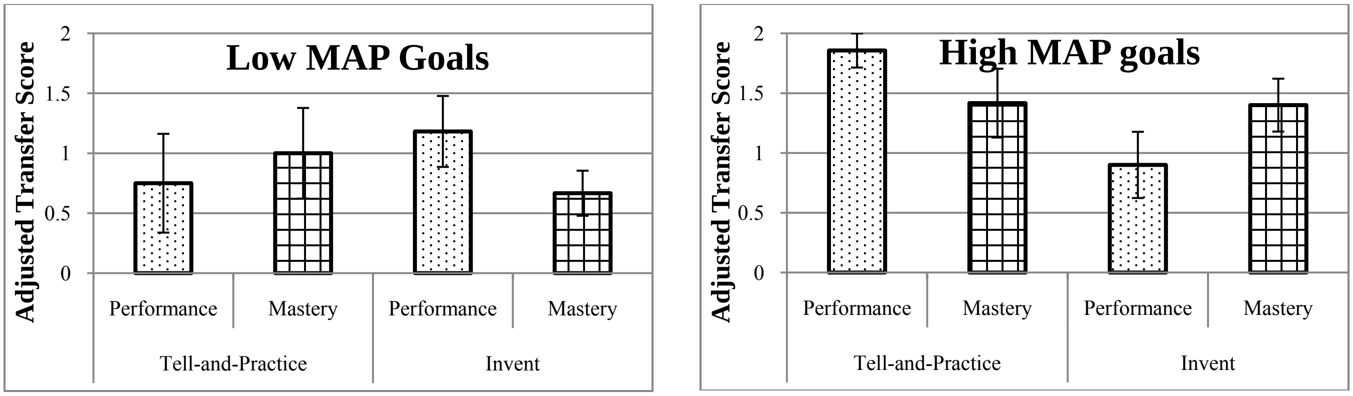

Figure 2. Adjusted transfer score (out of 2, post-test minus pre-test), for those below median in mastery-approach goals (MAP) in the left graph, and those above median in mastery-approach goals on the right, +/- 1 S.E.

goal instructions) x 2 (high or low performance-approach) ANOVA. There was no effect of instructional goal, F (1, 88) < 1, *ns*, but there was an effect of task structure, F (1,88) = 4.09, $p < .05$, with tell-and-practice leading to better performance (77% correct) than invention (57% correct). There was also a significant effect for PAP goals; those high in PAP goals (77% correct) did better than those who entered low in such goals (56% correct, F (1, 88) = 4.59, $p < .05$. However, the experimental manipulations did not interact with existing goals, nor was the three-way interaction significant, Fs (1, 88) < 2.48, *ns*. For learning a simple skill, the dispositional goal of being able to perform leads to better acquisition, as does a well-structured learning task.

To look at transfer, an adjusted score was calculated. This value is the post-test score minus pre-test score on two isomorphic transfer problems, which were counterbalanced. Participants who were at ceiling on pre-test were taken out of subsequent analyses of transfer, as they already had knowledge of the formula, and when it applies. No effect of either task structure or task goals on transfer performance was found. Invention activities produced the same level of transfer as tell-and-practice, as did mastery task goals and performance task goals, and there was no interaction between these factors, Fs (1,70) < 1.16, *ns*. However, students high in MAP goals did better (M = 1.36, SD = .84) than those low in MAP goals (M = .90, SD = .94) on the measure of transfer, F (1, 70) = 5.85, $p < .05$. When paired with the result that PAP goals lead to better skill, this is clear evidence that existing student goals affect the type of learning students exit instruction with. There was also a significant three-way interaction, F (1, 70) = 4.36, $p < .05$ (see Figure 2). This interaction is due to the invent structure and performance task goals condition behaving differently than all of the other conditions in terms of the effect of existing MAP goals. In the other conditions, high MAP goals produce better transfer than low MAP goals, as per the main effect. However, within the invention/performance-oriented cell, those low in MAP goals did slightly better than those high in MAP goals (Ms = 1.18, .9, SDs = .98, 88, respectively). It appears that an ill-structured environment like the invention activity, when presented in an evaluative context, changes the pattern of learning a mastery-oriented student might normally engage in, harming their ability to flexibly transfer their knowledge.

Discussion

This study examined whether explicit task goals would affect learning the same way task structures have been shown to. The evidence shows that they do not. Task structure was a bigger determinant in the learning outcome, as evidenced by the main effect of a well-structured domain helping with skill acquisition, and the difference in goal adoption on the activity questionnaire, which showed differences based on task structure, but not based on the task goals received. It seems that telling students which goal to adopt is not a particularly effective way to change learning behaviors or affect, and that task structure is a better way to change a student's focus during an instructional event.

The results provide strong evidence that existing student achievement goals change what a student learns during the course of instruction. Those high in performance-approach goals did better on skill measures, regardless of which task structure they learned with, while those high in mastery-approach goals did better on the transfer measure. These differences on their own are illuminating in light of the null effects for mastery-approach goals on achievement measures (see Harackiewicz et al., 2002). This may be due to the types of measures used when achievement is measured in schools. A grade on the test at the end of the semester may reflect factual knowledge that a performance-oriented student has been more focused on attaining than the deep understanding a mastery-oriented student focused on.

Finally, this work also examined whether having multiple goals would help one to learn, or if incongruous goals would produce a poor "fit" and interfere with successful learning. Within a given outcome (i.e., skill or transfer), we do not see evidence that multiple goals are best. On the skill measure, we see no interaction between existing goals, task structure and instructional goals. For transfer, we see evidence for a "regulatory fit" model, where a task goal that did not fit with the task structure and an existing achievement goal harmed learning. The focus on ability created by performance-oriented instructions may have changed the way in which these students would have normally processed the material, which, based on the performance of their equivalent group who received mastery-oriented instructions, and from our prior study, would have performed much better. The evaluative context brought on by the performance instructions appears to have fundamentally changed the way these students engaged in the learning activity, producing worse transfer. Though we have discussed this in terms of regulatory fit, this adverse effect could also be due to anxiety, or the use of simpler learning strategies related with performance goals (Elliot, 1999).

Conclusion

The study of student motivation has much to offer researchers focused on *how* students learn. It seems naïve to believe that the goal a student has in a learning environment does not influence the form and utility of the knowledge generated, and this line of work is bearing this out empirically. Specifically, research has found that the way information is represented and processed will have an effect on how that information is used to solve new problems and learn new concepts (e.g., Nokes & Ohlsson, 2005). A student's goals for a learning environment can be a catalyst for different types of processing and representations, as demonstrated by performance-approach goals predicting performance on measures of skill and mastery-approach goals predicting performance on transfer tests.

That we can reliably show these differences in a laboratory setting seems to be evidence that these results have high external validity, as student goals should be even more salient in authentic academic environments than in the lab. Also, while our attempts to manipulate these goals did not produce effects similar to existing achievement goals, stronger interventions conducted over longer periods of time may have very profound effects on student learning, especially when they succeed in changing the goal itself (i.e., Blackwell, Trzesniewski, & Dweck, 2007).

Within a task, it appears that making the learning goals, task structure, and instructions as coherent as possible promotes better learning. Those high in mastery-approach goals who invented were disrupted by instructions which placed the task into a more evaluative context. While the results within a given learning measure (e.g., skill or transfer) do not support a "multiple-goals" viewpoint, the study as a whole does. The separation of skill and transfer seems to illustrate one potential mechanism in support of a multiple-goals viewpoint, which claims that a combination of performance-approach and mastery-approach goals is optimal (Harackiewicz et al., 2002). Our work illustrates that having both goals would allow for the development of efficient skills at routine aspects (i.e., performance-approach goals lead to improved formula use) as well as innovative ability to use the underlying concept (i.e., mastery-approach goals lead to improved transfer). This combination of efficiency and innovation is hypothesized to be necessary for adaptive expertise, a desired outcome of education (Hatano & Inagaki, 1986). If the multiple-goals viewpoint is correct, a critical question left to examine is if students flexibly switch between these goals, and how they do so. It may be that the time course is a critical factor in explaining why, across a semester, multiple goals may be optimal. A more fine-grained, microgenetic study of when and how each goal contributes could shine light on this possibility.

Future work should also further explore the mechanisms by which achievement goals affect the type of learning done. It could be due to different representations used and formed and/or different types of learning strategies (i.e., self-explanation versus rehearsal). Achievement goals remain an important individual difference to consider, and one that could greatly impact how we can use cognitive science to improve education.

Acknowledgments

This work was supported by the National Science Foundation, Grant Number SBE-0354420 to the Pittsburgh Science of Learning Center (http://www.learnlab.org).

References

Belenky, D.M., & Nokes, T.J. (2009a). Motivation and transfer: The role of achievement goals in preparation for future learning. In N.A. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 1163-1168).

Blackwell, L.S., Trzesniewski, K.H., & Dweck, C.S. (2007). Implicit theories of intelligence predict achievement across and adolescent transition: A longitudinal study and an intervention. *Child Development, 78* (1), 246-263.

Elliot, A.J, & Harackiewicz, J.M. (1996). Approach and avoidance achievement goals and intrinsic motivation: a meditational analysis. *Journal of Personality and Social Psychology, 70* (3), 461-475.

Elliot, A.J., & McGregor, H.A. (2001). A 2 X 2 achievement goal framework. *Journal of Personality and Social Psychology, 80*(3), 501-519.

Elliot, A.J. (1999). Approach and avoidance motivation and achievement goals. *Educational Psychologist, 34* (3), 169-189.

Elliott, E.S., & Dweck, C.S. (1988). Goals: An approach to motivation and achievement. *Journal of Personality and Social Psychology, 54* (1), 5-12.

Harackiewicz, J.M., Barron, K.E., Pintrich, P.R., Elliot, A.J., & Thrash, T.M. (2002). Revision of achievement goal theory: Necessary and illuminating. *Journal of Educational Psychology, 94* (3), 638-645.

Hatano, G., & Inagaki, K. (1986). Two courses of expertise. In H. Stevenson, H. Azuma, & K. Hakuta (Eds.), *Child development and education in Japan* (pp. 262-272). New York: Freeman.

Higgins, E.T. (2005). Value from regulatory fit. *Current Directions in Psychological Science, 14* (4), 209-213.

Nokes, T. J., & Ohlsson, S. (2005). Comparing multiple paths to mastery: What is learned? *Cognitive Science, 29*, 769-796.

Schwartz, D.L., & Martin, T. (2004). Inventing to prepare for future learning: The hidden efficiency of encouraging original student production in statistics instruction. *Cognition and Instruction, 22*(2), 129-184.