



141  
749  
THS

ESIS

009



This is to certify that the  
thesis entitled

EXPLORATION AND EXPLOITATION:  
A NEW EXPLANATION OF DIFFERENTIAL  
GOAL-SETTING EFFECTS

presented by

PAUL G. CURRAN

has been accepted towards fulfillment  
of the requirements for the

          M.A.           degree in           PSYCHOLOGY          

*Paul Schmitt*

Major Professor's Signature

*12/3/08*

Date

*MSU is an Affirmative Action/Equal Opportunity Employer*

**PLACE IN RETURN BOX** to remove this checkout from your record.  
**TO AVOID FINES** return on or before date due.  
**MAY BE RECALLED** with earlier due date if requested.

<b>DATE DUE</b>	<b>DATE DUE</b>	<b>DATE DUE</b>

**EXPLORATION AND EXPLOITATION:  
A NEW EXPLANATION OF DIFFERENTIAL GOAL-SETTING EFFECTS**

**By**

**Paul G. Curran**

**A THESIS**

**Submitted to  
Michigan State University  
in partial fulfillment of the requirements  
for the degree of**

**MASTER OF ARTS**

**Psychology**

**2008**



## ABSTRACT

### EXPLORATION AND EXPLOITATION: A NEW EXPLANATION OF DIFFERENTIAL GOAL-SETTING EFFECTS

By

Paul G. Curran

Goal-setting theory has been shown over the years to have a number of experimental outcomes inconsistent with previous theory. An experiment was designed in order to attempt to intentionally produce these outcomes in a controlled design, such that explanation could be provided a priori as opposed to post hoc. It is suggested that these results, as well as more standard results, could be explained using the constructs of exploration and exploitation. Results basically supported main predictions concerning these ideas.

## TABLE OF CONTENTS

LIST OF TABLES.....	v
LIST OF FIGURES.....	vi
INTRODUCTION.....	1
Complexity.....	5
Exploration and Exploitation.....	10
Exploration and Problem Representation.....	12
Exploration/Exploitation and Goal-Setting Effects.....	14
Exploration and Reinforcement Learning.....	18
Summary and Hypothesis.....	20
Individual Difference Effects.....	21
Curiosity.....	22
Goal Orientation.....	23
METHOD AND RESULTS.....	25
Study 1: Pilot.....	25
Participants.....	25
Procedure.....	25
Results and Discussion.....	27
Study 2: Pilot 2.....	29
Participants.....	29
Design.....	30
Procedure.....	31
Measures.....	32
Perceptual complexity.....	32
Perceptual difficulty.....	32
Results.....	32
Study 3: Test of Goal-Setting Effects.....	34
Participants.....	34
Design.....	35
Procedure.....	35
Measures.....	38
Trait level goal orientation.....	38
Trait level curiosity.....	38
Perceptual complexity.....	39
Perceptual difficulty.....	39
Exploration (operationalization).....	39
Results.....	40
Manipulation check.....	40
Main Analyses.....	40
Discussion.....	46
Implications.....	50

Complexity.....	50
Exploration and problem representation.....	51
Exploration and goal setting theory.....	51
Exploration and reinforcement learning.....	52
Other implications.....	53
Practical implications.....	54
Limitations.....	54
Next steps.....	56
Conclusions.....	57
APPENDICES.....	59
REFERENCES.....	71

## LIST OF TABLES

Table 1: Regression of cumulative exploration on performance over time by world.....	59
Table 2: Linear regressions of proposed moderators on exploration: results for hypothesis 3, 4a, and 4b (EX1) .....	60
Table 3: Linear regressions of proposed moderators on exploration: results for hypothesis 3, 4a, and 4b (EX2) .....	61
Table 4: Results of repeated measures analysis on performance and exploration over time by goal condition and world.....	62

## LIST OF FIGURES

Figure 1: 'Do your best' goals .....	63
Figure 2: Specific, difficult goals.....	64
Figure 3: Performance over time.....	65
Figure 4: World 1.....	66
Figure 5: World 2.....	67
Figure 6: World 3.....	68
Figure 7: Start of trial (zoomed out).....	69
Figure 8: After some exploring in trial (zoomed out).....	70

Images in this thesis/dissertation are presented in color.

## INTRODUCTION

The finding that specific, difficult goals improve performance relative to 'do your best' goals has been demonstrated in thousands of studies throughout the course of nearly 40 years (Locke, 1968; Locke & Latham, 1990; Locke & Latham, 2002; Locke, Shaw, Saari, & Latham, 1981). Beyond this finding, research has also addressed how goals produce such a shift in performance. Goals are said to accomplish this shift through four distinct mechanisms: goals "direct attention and effort toward goal-relevant activities and away from goal-irrelevant activities", "have an energizing function", "prolong effort" via persistence, and lead "to the arousal, discovery, and/or use of task-relevant knowledge" (Locke & Latham, 2002, p706-707). In short, goals increase and direct personal resources to create a boost in performance.

Despite such findings, goals do not produce a strong goal setting effect in every situation. There exist situations in which specific difficult goals produce a much weaker goal setting effect over do your best goals, or even a negative goal setting effect where those given 'do your best' goals outperform those given specific, difficult goals (Locke, Shaw, Saari, & Latham, 1981). For instance, in a study by Kanfer and Ackerman (1989), participants in a specific, difficult goal condition were outperformed early in the task by those in a 'do your best' condition. Huber (1985) demonstrated the same effect on a simple maze task with both easy and difficult versions of the maze. Sweller and Levine (1982) found that the presence of a goal (relative to a no-goal condition) led to such a detriment in task performance that a simple maze task was rendered virtually unsolvable in the goal condition. Earley, Connolly, and Lee (1989), as well as DeShon and Alexander (1996) showed a negative goal effect on similar, moderately complex,

versions of a multiple cue probability learning task. Burns and Vollmeyer (2002) also found a negative goal setting effect on transfer performance on a multiple cue probability learning task. While these studies do not reach the number of studies showing a positive goal setting effect they are still too numerous to dismiss as experimental flukes. What, then, is occurring?

Early speculation around this idea led to the suggestion that task complexity might account for weaker cases of the goal setting effect (Locke et al, 1981). Using Wood's (1986) model of complexity, Wood, Mento, and Locke (1987) performed a meta-analysis of the goal setting literature and showed that, in general, as task complexity increased the positive goal setting effect decreased. Thus, simple tasks showed the strongest positive goal setting effect, and more complex tasks showed the weakest. This finding has been explained in numerous ways, usually focusing on aspects of cognitive resources such as attention and effort (Kanfer & Ackerman, 1989; Locke et al, 1981).

While Wood, Mento, and Locke's (1987) explanation of task complexity moderating the goal setting effect does seem to work for a large number of studies there are three experimental results that are contradictory. First, the task complexity explanation cannot account for studies with simple tasks where specific, difficult goals show diminished performance in relation to those with 'do your best' goals. A number of these studies have been highlighted above (Burns & Vollmeyer, 2002; Huber, 1985; Sweller & Levine, 1982). Second, the task complexity explanation cannot account for studies with complex tasks where specific difficult goals show a strong goal setting effect - that is, notable increased performance over subjects given 'do your best' goals. While not as numerous, these effects do exist (Ivancevich, 1974; Pritchard, Jones, Roth,

Stuebing, & Ekeberg, 1988). Finally, the task complexity explanation does not account for situations in which a single task can produce vastly different goal setting effects through non-complexity based manipulation. While no studies have directly addressed this issue, results such as those of Kanfer and Ackerman (1989) can be interpreted from this framework and will be discussed later in this paper. Nevertheless, numerous studies have demonstrated the first two experimental situations. Given these findings, goals must be acting through a means other than task complexity.

Wood (1985) suggested that specific, difficult goals on complex tasks might lead to suboptimal search processes. Others, such as Earley, Connolly, and Lee (1989), suggested that specific, difficult goals might lead to excessive strategy search. Seijts and Latham (2001) suggested that performance (score) goals and learning (exploration) goals might have differential effects. These different explorations, at their roots, lead back to a basic psychological distinction between exploration and exploitation.

Exploration has been treated in many different ways in the various fields of psychology, but it may in fact be a simple dictionary definition that best encapsulates it – “the investigation of unknown areas” (Random House, 2006). When studied, exploration is often defined and studied in relation to its counterpart of exploitation. Exploitation is the behavior of acting on a system in the best known way. Exploitation comes at the cost of exploration the same that exploration comes at the cost of exploitation. While exploration is often said to be driven by motivations (such as an animal searching for food to survive), exploration is potentially motivated by different mechanisms (Voss & Keller, 1976); situations “for which there appears to be no goal object or condition to and for which the organism responds” (Fowler, 1965, p23). In terms of goal setting,



exploration can be defined as actively sacrificing known optimality toward a goal in order to search the unknown.

The purpose of this research is to suggest and show that the way in which specific difficult goals may be causing both positive and negative goal setting effects is through their effect on exploratory behavior. Specific goals, according to goal setting theory, “direct attention and effort toward goal-relevant activities and away from goal-irrelevant activities” (Locke & Latham, 2002, p706). Exploitation – action on the task toward the goal – is a goal-relevant activity and produces positive goal effects in thousands of research studies. It involves striving toward the goal in the best way one knows how. In this case, exploration for the sake of exploring is goal-irrelevant. Put simply, it is suggested that specific goals limit exploration relative to ‘do your best’ goals.

It is thus further suggested that this effect on exploratory behavior produces different goal effects due to the fact that performance on some tasks requires more exploration than others. If a task requires exploration to perform well (Burns & Vollmeyer, 2002; Sweller & Levine, 1982), specific goals that inhibit this exploration will decrease performance and produce a negative goal effect (Burns & Vollmeyer, 2002; Sweller & Levine, 1982). If exploration is not needed to perform well on a task (Latham & Kinne, 1974; Locke, 1968) then specific goals will help focus effort on the appropriate method of completing the task and produce a positive goal effect (Latham & Kinne, 1974; Locke, 1968).

This explanation also holds potential to more deeply explain the results of Wood et al (1987). Tasks that are generally more complex are likely to require more exploration to complete (producing a negative goal effect by limiting exploration) while

simpler tasks are likely to require less exploration to complete (producing a positive goal effect by limiting exploration). Thus, while complexity can appear to act as a moderator, it is suggested that what is occurring at a more basic level is the underlying effect of specific goals on task exploration.

To this end, complexity will first be examined before moving on to exploration and exploitation. With this foundation an experiment will be proposed. Proposed individual difference moderators of effects will also be discussed.

### Complexity

Before one can reasonably assess the validity of the task complexity moderation on goal setting effects, one must address a number of questions, such as ‘what is complexity?’ Much of the work on complexity is based on the framework laid out by Wood (1986). This framework is that on which the study by Wood, Mento, and Locke (1987) was carried out. It is centered on three aspects of complexity: component complexity, coordinate complexity, and dynamic complexity. In short, Wood suggests that these three aspects of complexity interact to determine how complex a task is.

In order to grasp Wood’s definition of overall complexity one must thus review Wood’s definitions of the aspects of complexity:

“Component complexity of a task is a direct function of the number of distinct acts that need to be executed in the performance of the task and the number of distinct information cues that must be processed in the performance of those acts” (Wood, 1986, p66). “Coordinative complexity refers to the nature of the relationships between task inputs and task products...[including] timing, frequency, intensity, and location requirements for performances of required acts” (Wood, 1986, p68). “Dynamic complexity is due to changes in the states of the world which have an effect on the relationships between task inputs and products” (Wood, 1986, p71).

Following from these definitions, Wood defines total task complexity as “a function of the component, coordinative, and dynamic complexities of the task” (Wood, 1986, p74). Thus, a clear and concrete definition of complexity at the general construct level remains unsupplied by Wood (1986). Moreover, a competing and no less valid typology of complexity was produced by Campbell (1988) just two years later and supplies a different view on complexity in which complexity is comprised of four dimensions different than Wood’s. Campbell (1988) claims that the four important dimensions of complexity are multiple paths, multiple outcomes, a conflicting interdependence among paths, and the presence of uncertain or probabilistic linkages. While there is slight overlap with Wood’s (1986) dimensions, this produces a much different picture of what complexity is and highlights the fact that complexity as a singular construct is incredibly difficult to define. Even a search in a dictionary produces a fruitless and circular result: “Complexity: 1. The quality or condition of being complex. 2. Something complex: *a maze of bureaucratic and legalistic complexities*” (American Heritage Dictionary, 2006).

Further, while component complexity is fairly concrete and coordinative complexity somewhere between concrete and abstract, dynamic complexity is placed as the most important dimension of task complexity while simultaneously possessing the highest level of abstraction. Transfer tasks are a useful example of this, and Wood’s take on dynamic complexity begs the question of how they might be rated on this scale. Is a task complex because of its characteristics, or because of how it is administered? If a transfer task is presented in which task rules change slightly, can it possess its own rating of complexity or must now the entire task and transfer task be rated with the same

values? Does the task itself possess complexity outside of the situational boundaries in which it is presented? The answers to these questions are well beyond the scope of this paper, but highlight a further difficulty in the conceptualization of complexity as a whole.

With such difficulty in conceptualization of complexity it is no surprise that operationalization of complexity proves to be equally difficult. Wood (1986) puts forth a series of concrete and quantifiable equations that can be used to calculate the values for the different dimensions of complexity, and in fact even presents two example cases in which he presents explanations and values for these figures. Despite this, however, Wood is forced to admit that “at this point, we cannot specify the exact form of the relationship between the different types of complexity and total task complexity” (Wood, 1986, p74). Without this relationship, Wood’s final equation of total task complexity (total complexity =  $\alpha$ \*component complexity +  $\beta$ \*coordinative complexity +  $\gamma$ \*dynamic complexity) possesses three unknowns – the different dimensions’ weighting variables  $\alpha$ ,  $\beta$ , and  $\gamma$ . Thus, while two tasks that possess values in only one of the three dimensions can be compared in terms of complexity, comparing values across the different dimensions is functionally equivalent to comparing apples and oranges.

Wood, Mento, and Locke (1987) further state that “task complexity is ordinal in nature” (p417), further distancing the field from Wood’s implication that complexity may be represented on a mathematically defined ratio scale. This is carried out in their operationalization of complexity, as even they (Wood, Mento, Locke, 1987) do not use the mathematical formulas set forth a year earlier by Wood (1986) and instead use aggregate ratings by coders on a ten-point scale. This coding on 125 tasks produces a range of complexity ratings from only one to seven, with one representing simple

reaction time tasks and seven representing “science and engineering” (Wood, Mento, Locke, 1987, p 418). Over 60 percent of the tasks are represented by ratings of two or three on the ten point scale. Beyond these problems, but possibly underlying them, is the fact that - regardless of the simplifying assumption that task complexity can be represented on an ordinal scale - Wood, Mento, and Locke (1987) have still not overcome the pitfall put forth by Wood (1986): “we cannot specify the exact form of the relationship between the different types of complexity and total task complexity” (Wood, 1986, p74). Thus, total task complexity as laid out by Wood cannot be represented even on an ordinal scale, by current measurement methods. Only effects can be measured, tasks that differ on different types of complexity have the potential to look the same on ratings. The fundamental assumptions underlying the operationalization in Wood, Mento, and Locke (1987) were shown to be unreachable a year prior by Wood (1986) himself.

As already discussed, complexity is suggested to be linked to the goal setting effect in such a way that larger positive goal setting effects are found on simpler (less complex) tasks while smaller positive goal setting effects will be found on more complex tasks. Wood, Mento, and Locke (1987) showed that this relationship holds for the most part using the framework of task complexity put forth by Wood (1986) and discussed above. The implication follows that goal-setting effects should not be prone to manipulation if such manipulation is not done on the complexity of the task.

Even accepting that in largely distinct tasks (very complex and minimally complex) complexity can be measured ordinally, data presented earlier in this paper still falls counter to the hypothesis of Wood, Mento, and Locke (1987) that less complex tasks

should possess more powerful positive goal setting effects than more complex tasks. There exist tasks that possess relative low levels of complexity (maze tasks, simplified MCPL tasks) in which participants still experience a negative goal setting effect (Burns & Vollmeyer, 2002; Huber, 1985; Sweller & Levine, 1982). There exist tasks that possess relative high levels of complexity (complex jobs in workforce, management) and show large positive goal setting effects (Ivancevich, 1974; Pritchard, Jones, Roth, Stuebing, & Ekeberg, 1988). Finally, Kanfer and Ackerman (1989) provide an accidental demonstration of the third situation postulated to show a failing of the theory of complexity as a moderator of the goal effect: a situation in which a single task produced different goal setting effects through non-complexity based manipulation. This effect will be discussed later in the paper.

A task, according to Wood (1986), and by extension Wood, Mento, and Locke (1987), is complex because of the three dimensions of complexity and these three ideas alone. Goals, unless they alter one of these three components, cannot make a task more or less complex. In Kanfer and Ackerman (1989), goals possess no complexity altering qualities. In this experiment, it is first shown that when participants are given goals at the start of the experiment it produces a negative goal setting effect – specific, difficult goal participants are outperformed by those in a ‘do your best’ goal condition. When this same goal manipulation is given later in the experiment (all participants receive ‘do your best’ goals in trials one to four, specific, difficult goals begin in goal condition in trial five), the goal setting effect is reversed and becomes a positive goal setting effect in which do your best goal participants are outperformed by those in the difficult specific goal condition. In Wood’s (1986) framework there is no possible way this manipulation

could have changed the complexity of the task, yet the goal-setting effect was changed nonetheless. From an objective task complexity standpoint this is impossible. Cumulative with the previous evidence, Kanfer and Ackerman (1989) show that examples exist to counter all three of the main implications of Wood, Mento, and Locke's (1987) finding. What, then, is actually the cause?

### Exploration and Exploitation

The psychological questions underlying the topic of exploration have been around nearly as long as the science of psychology. In attempts to explain behavior, early psychologists reasoned that just as instincts drive other animals, so too must human behavior be driven by instinct (James, 1890). This idea eventually evolved into the notion that motivational drive states, such as hunger, thirst, and pain, caused behaviors that would return the being to a state of homeostasis (Cannon, 1932; Carr, 1925). In the case of goal setting, this would reasonable account for actions taken to move oneself toward a goal state (i.e. exploitation). However, this idea failed to account for situations in which an organism would investigate a puzzle without the presence of a primary drive (Harlow, 1950). The following investigation of exploratory behavior became very much intertwined with the idea of curiosity as a possible drive of this behavior (Fowler, 1965). In summary of this following research it was recommended that “exploration should be used to describe observable behavior, whereas curiosity should be used to describe the corresponding hypothetical construct” (Voss & Keller, 1983, p150). Furthermore, exploratory behaviors “are visual, linguistic, haptic, motoric, or intellectual activities that make *new* information available to the individual” and “are used to reduce uncertainty” (Voss & Keller, 1983, p150). This new information has only the potential of being

useful, however, and the time spent collecting it is time *not* spent exploiting the information one already has. This balance of exploration and exploitation as a resource allocation issue is an issue of importance.

Over time, many different areas have come to study this dilemma of exploration versus exploitation, knowingly or not. The topics of game theory (Jervis, 1988), evolutionary computing and genetic algorithms (De Jong & Spears, 1993), and even dating and mate selection (Das & Kamenica, 2005) have all come to the conclusion that there exists a balance between exploration and exploitation that must be met to obtain optimality. One must spend time exploring and collecting information, but must also realize when the potential for future information must be sacrificed for immediate action with an incomplete knowledge set. This subtle balancing act is pushed to even farther limits in fields such as foraging theory (Pyke, 1984), where neglecting exploration for exploitation or exploitation for exploration can literally result in death at both the individual and species level. If a herd of antelope stay at the same watering hole and eventually eat all the vegetation around it without allowing for re-growth, they will eventually starve. A similar herd that moves outside of their known world to find other watering holes while using only some of the vegetation of the first will now not only have food, but the knowledge of where to find more.

The study of exploration and exploitation has also found root in organizational research. March (1991) noted the balance that must be kept between these two ideas, and that deviating from optimal in both directions could be disastrous:

“Adaptive systems that engage in exploration to the exclusion of exploitation are likely to find that they suffer the costs of experimentation without gaining many of its benefits....Conversely, systems that engage in exploitation to the exclusion of exploration are likely to find themselves trapped in suboptimal stable



equilibria. As a result, maintaining an appropriate balance between exploration and exploitation is a primary factor in system survival and prosperity. (March, 1991, p71)”

March (1991) also takes note of the vulnerability of exploration in that it produces less certain outcomes that only have a chance of paying off. Just as in the research mentioned above, exploration is not a sure bet. In do or die situations exploration may hold much more risk than benefit. Other organizational research has taken place in the context of exploration and exploitation, from research on innovation (Nooteboom, 1999), resource allocation (Garcia, Calantone, & Levine, 2003), and even strategic ventures and alliances (Rothaermel & Deeds, 2004).

#### Exploration and Problem Representation

Work in the field of problem space and problem representation provides a link between goals and the balance of exploration and exploitation. “A problem space has an initial state, a goal state, and a set of operators that can be applied that will move the solver from one state to another” (Dunbar, 1998, p 290). In this framework, the representation of the goal state can fundamentally alter the behaviors that individuals enact to move from an initial state to the goal state. As a concrete example, Miller, Lehman, and Koedinger (1999) used a task called electric field hockey in which participants tried to push a puck between obstacles and into a net. Participants were given either a specific goal, a nonspecific goal, or an appropriate specific goal. The appropriate goal was designed to push participants toward the optimal manner of problem completion through the accurate representation of the task. The appropriate specific goal and the nonspecific goal groups both produced better performance than the specific goal group, suggesting that participants could find the optimal behaviors by search and

exploration (no goal), or by specific direction toward it (appropriate goal). The linkage of exploration and search to performance in the no goal group was supported by significantly increased scores on a knowledge test relative to other conditions. In all, it was suggested that the lowered performance in the specific goal condition was due to the primes toward task-irrelevant actions built into the standard goals. In fact, it was suggested that one of the benefits of the no goals condition was that “the no goal condition was not ‘distracted’ by having to focus on irrelevant...relationships of standard [electronic field hockey] play” (Miller, Lehman, & Koedinger, 1999, p320). The appropriateness and inappropriateness of these goals can be viewed as the balance of exploration and exploitation they foster relative to the optimal balance for the particular task characteristics.

Strongly linked to the idea of appropriate and inappropriate goals is the idea of insight problem solving. Insight problems are defined as those problems that require restructuring of the initial problem representation to reach a solution (Chronicle, MacGregor, & Ormerod, 2004). This is to say that participants, when placed in the task, exhibit a general and (often unanimous) default representation of the task that is sub-optimal. Successful performance on an insight problem requires the individual to break away from the individual representation and search for a more functional representation. This required shift in representation is what allows Miller, Lehman, and Koedinger (1999) to make claims about the appropriateness or inappropriateness of a goal. In their study those in the no goal and appropriate goal conditions were able to represent the task differently to their benefit than those in the standard goal condition. Further, those in the

no goal condition were able to fully represent the task instead of just the areas of it relative to specific outcomes.

Also in the field of problem representation, Vollmeyer, Burns, and Holyoak (1996) showed that having a goal encouraged participants to deduce ways to simply meet that goal without creating hypotheses about the underlying structure of the problem. Put another way, those in the goal condition were encouraged by the goals to use a goal-oriented or means-end method of solving the task (i.e. exploitation). Those in the nonspecific goal condition had spent their time producing and testing hypotheses and learning the structure of the task (i.e. exploration) such that when given a different goal they were able to apply the same rules and reach the goal almost trivially.

#### Exploration/Exploitation and Goal-Setting Effects

This balance of exploration and exploitation is also present in some goal setting research, although not specifically treated as such. In a study by Seijts and Latham (2001), it was found that “a specific, difficult learning goal leads to higher performance than urging people to ‘do their best’ on a task where people initially lack the requisite knowledge to perform it” (p22). Conversely, it was found that “specific, difficult outcome goals...appear to have a detrimental effect on performance during the learning process” (p22). The finding by Seijts & Latham (2001, 2005) is consistent with the idea that specific, difficult (performance) goals will decrease exploration on a task, as learning is a form of exploration involving the collection of some new knowledge. This second result – that specific, difficult goals are detrimental during the learning process – is consistent with Kanfer and Ackerman’s (1989) findings on the air traffic controller task. Like the air traffic controller task, as well as a number of studies along this topic

(Kozlowski & Bell, 2006; Seijts & Latham, 2005), the task in Seijts & Latham (2001) is a complex task where learning is required. Thus, while this result of specific, difficult performance goals affecting performance dependant on stages of learning has been shown on complex tasks where learning is required, little, if anything has been said of predicted effects on less complex tasks or tasks where no learning is required. A more basic explanation involving the promotion or inhibition of exploration by goals works not only in these cases of complex tasks demonstrated by Seijts and Latham (2001), but also in situations of simple tasks where exploration on task – not just learning – can be required or detrimental.

Moreover, the findings of Seijts and Latham (2001) require a shift to a qualitatively different paradigm than what has been used in the traditional goal setting studies that they are using as foundation. Mastery goals are qualitatively different in structure and purpose than traditional performance goals. It was shown that mastery goals can promote learning and exploration. In situations where learning and exploration is useful mastery goals should be useful, and in situations where learning is not required and/or exploration is detrimental mastery goals should not be useful. A discussion much further into the realm of mastery goals is beyond the scope of this paper and the issue at hand. The inconsistent goal setting effects found in traditional goal setting literature (Burns & Vollmeyer, 2002; DeShon & Alexander, 1996; Early et al, 1989; Huber, 1985; Kanfer & Ackerman, 1989; Sweller & Levine, 1982) all have to do with performance goals, and thus performance goals and their relation to ‘do your best’ goals is what must be first explained. While Seijts and Latham’s (2001) work provides an explanation of effects in complex tasks using mastery goals, the promotion and inhibition of exploration

by specific, difficult performance goals leads to the same predictions in traditional goal setting theory from a more basic yet more encompassing standpoint, explaining not only complex tasks but also simple tasks and all those in between. Moreover, it is done without needing – but also allowing for – mastery goals.

Earlier in the paper, the results of Kanfer and Ackerman (1989) were shown to be in direct contradiction to the findings of Wood, Mento, and Locke (1987) in that the goal setting effect was moderated with no change in complexity. How, then, does exploration and exploitation fare in explaining this result? Early in the task exploration is a useful strategy as participants learn the rules of the task. Thus, if goals reduce exploration, they will deleteriously affect performance. This is consistent with results from Kanfer and Ackerman's (1989) study one. Later in the task, after all participants have explored enough to learn the rules of the task, further exploration is not useful and exploitation is the favored strategy. Thus, if goals reduce exploration, they will positively affect performance. This is consistent with the results of Kanfer and Ackerman's (1989) study two.

Other tasks in the goal setting literature also fit into this framework. Sweller and Levine (1982) used an experimental design in which the manipulation consisted of giving the participant the location of the exit of a maze (southwest of the start of the maze). This drove participants to exert all their effort at moving southwest while the *path* to the exit actually existed to the northeast. Participants without knowledge of where the goal was located had no choice to explore the maze, and outperformed the goal condition by an overwhelming margin. So powerful was this effect, Sweller and Levine (1982) also

provided a replication, of which they noted: “The presence of a goal has again rendered the problem essentially insoluble” (p 469).

In a study by Earley, Connolly, and Ekegren (1989), it was found that specific difficult goals were harmful to initial performance on a complex task, similar to the findings of Kanfer and Ackerman (1989). Earley, Connolly, and Lee (1989) completed a study in order to both replicate these results and try to go one step further. The task in this study consisted of a stock market prediction task in which participants were to predict stock prices for 75 fictional companies based on performance numbers from different departments. This is an example of a multiple cue probability learning task, or MCPL task. Participants were placed in a specific, difficult goal condition or a ‘do your best’ goal condition. Participants were also placed into one of three groups in terms of strategy interventions. The first group received no strategy training, the second received a list of ten possible prediction strategies (search restriction condition), and the third received a step by step guide to developing their own accurate prediction strategy (search training condition). It was found that, relative to others, subjects with specific difficult goals did significantly worse in predicting stock prices on the first block of 15 trials (exploitation before adequate knowledge has been acquired), but significantly better on all trials following. It was also found that the strategy interventions had a significant impact on performance and search for those in the goal condition, but not in the do your best condition. This is perhaps due to the fact that those in the do your best condition were already engaged in exploration of different strategy, while those in the specific, difficult goal condition were not.

In another MCPL task, DeShon and Alexander (1996) implemented a series of manipulations, the first between goal and do your best conditions, the second between implicit and explicit strategy development. In this situation implicit pattern recognition and learning can be conceptualized as an exploratory behavior, while a push toward explicit learning can also be seen as a push toward exploitation. The outcomes of this study suggested that when a task is processed explicitly, goals can be used to positive effect by focusing attention on the task (exploitation). When the task is one where an implicit strategy is best (exploration), goals will hurt performance.

#### Exploration and Reinforcement Learning

The distinction between exploration and exploitation is also illustrated in the field of reinforcement learning. Reinforcement learning is “the problem faced by an agent that learns behavior through trial-and-error interactions with a dynamic environment” (Kaelbling, Littman, & Moore, 1996, p237). Task-participant interaction in reinforcement learning can generally be described a situation in which a participant receives a perceptual indication of the current state of some system, takes an action on that system, and is presented with feedback as to what result the action produced. Over time, the participant will choose actions which maximize the value of the result of interest. For example, to use a task from the goal setting literature discussed earlier, Kanfer and Ackerman’s (1989) air traffic controller task can be described in this way, incorporating the balance of exploration and exploitation. At any given time the participant can visually determine the state of the system of planes and runways, take actions - such as landing a plane, and then see the result of that action on their score. The

participant may explore the rule set and learn about the task or apply whatever knowledge is held at the moment to immediate performance.

While even tasks as complex as these can be described in this way, much of reinforcement learning focuses on simpler tasks so that distinctions such as those between exploration and exploitation can be made clear, and so that computer algorithms can be written to solve them. One such simple task is the  $k$ -armed bandit problem (Berry & Fristedt, 1985). This task is summarized by Kaelbling et al (1996):

“The agent is in a room with a collection of  $k$  gambling machines (each called a ‘one-armed bandit’ in colloquial English). The agent is permitted a fixed number of pulls,  $h$ . Any arm may be pulled on each turn. The machines do not require a deposit to play; the only cost is wasting a pull playing a suboptimal machine. When arm  $i$  is pulled, machine  $i$  pays off 1 or 0, according to some underlying probability parameter,  $p_i$ , where payoffs are independent events and the  $p_i$ s are unknown.” (p234)

In a situation where  $k$  is greater than one, the agent is faced with a dilemma. If the agent believes that one of the machines has the best underlying pay probability, optimal strategy should dictate that all pulls be on this machine (Kaelbling et al, 1996). However, there is a chance that a different machine has a better pay probability and that it is simply the case that not enough information has been collected to know it. Thus, the agent can choose to explore the other machines in order to collect new information at the cost of acting on what it believes to be optimal. This is the struggle between exploration and exploitation.

Another fairly simple task arises out of foraging theory (Pyke, 1984), mentioned earlier. Foraging theory examines the behavior of agents (often animals) who must interact with their world in order to find and utilize resources for survival. However, foraging theory is also applicable to human behavior, both in our distant past as hunter



gatherers (Bird, Smith, & Bird, 2001) and even in recent times in terms of decision making (Rode, Cosmides, Hell, & Tooby, 1999).

In foraging theory, an agent must exploit resources in order to survive, but also at the cost of depleting those resources. As such, the agent must also explore the world in order to gain knowledge about previously unknown resources that may be more or less rich or abundant than those already known. Thus, the agent is faced with a dilemma: the choice of exploiting known resources versus the choice of exploring for new resources.

What determines how this balance will be decided? In a situation without constraint it is easy to imagine that a person will try to reach a balance that allows for investment in both exploration of alternatives and exploitation of presumed optimality. However, constraints on the agent may cause the balance to swing one way or the other. As stated earlier, it is suggested that specific, difficult goals are one such constraint.

### Summary and Hypotheses

How then, do goals produce these constraining effects? Recall that goals “direct attention and effort toward goal-relevant activities and away from goal-irrelevant activities” (Locke & Latham, 2002, p706). Exploration is cases “for which there appears to be no goal object or condition to and for which the organism responds” (Fowler, 1965, p23). Exploratory behavior is goal-irrelevant behavior, such as wandering away from resources in a foraging theory world. Goals should inhibit this behavior and instead swing the balance of exploration and exploitation to the side of exploitation, as shown in numerous studies in the previous section (DeShon & Alexander, 1996; Earley, Connolly, & Ekegren, 1989; Earley, Connolly, & Lee, 1989; Kanfer & Ackerman, 1989; Miller,

Lehman, & Koedinger, 1999; Seijts, Latham, Tasa, & Latham, 2004; Sweller & Levine, 1982; Vollmeyer, Burns, & Holyoak, 1996). In short, goals inhibit exploration on a task.

*H1: The presence of a difficult, specific goal, relative to a condition of a 'do your best' goal, will inhibit exploratory search behaviors, negatively impacting the pattern and/or quantity of such behavior.*

Further, this effect on exploratory behaviors will manifest itself as a goal setting effect. As it is found in the literature, this goal setting effect will sometimes be strongly negative, sometimes be strongly positive, and sometimes be differing levels in between. The direction of the effect (positive, negative, or neutral) will depend on characteristics of the task, and can be well conceptualized in the framework of insight problems as discussed above (e.g. Chronicle, MacGregor, & Ormerod, 2004):

*H2a: If the default method for completing the task is the most effective, goals will limit exploration and direct effort toward this method, creating a positive goal setting effect.*

Conversely:

*H2b: If the default method for completing the task is not the most effective and exploration is required to find the most effective method, goals will create a negative goal setting effect by limiting exploration and thus lessening the chances of finding the most effective method.*

### Individual Difference Effects

There are individual difference constructs in psychology that are linked to the ideas that have been presented so far as the foundation of this basic process. The constructs that are seemingly most related to these ideas are those constructs of curiosity

and goal orientation. There are reasons to believe that these individual differences in participants may cause effects in concert with the above hypotheses. These proposed individual difference constructs will be discussed, and predictions involving their effects will be hypothesized.

### *Curiosity*

Curiosity has been defined as “a desire to acquire new knowledge and new sensory experience that motivates behavior” (Litman & Spielberger, 2003, p118). It has also been recognized as a critical motive that influences human behavior (Loewenstein, 1994), and has been further conceptualized as “information seeking, or cognitive curiosity, which stimulates information-seeking, exploratory behavior” (Reio, Petrosko, Wiswell, & Thongsukmag, 2006). This linkage to exploratory behavior is of use to this study through the proposed linkage between exploration and goal setting. As stated, the study of exploratory behaviors has long been intertwined with the study of curiosity as a possible drive of these behaviors (Fowler, 1965). Curiosity has also been shown to be a stable individual difference trait that can be measured (Kashdan, Rose, & Fincham, 2004). In all, some people may be generally more curious in their behaviors than others, and thus may naturally engage in more exploratory behaviors (Litman, Hutchins, & Russon, 2005; Reio, Petrosko, Wiswell, & Thongsukmag, 2006). Thus:

*H3: Those who score higher on ratings of trait level curiosity will engage on average in more exploratory behaviors than those who score lower. If this is found, trait level curiosity will be controlled for in above effects involving exploratory behavior.*

### *Goal Orientation*

Goal orientation is generally defined as the way that individuals interpret and approach situations in which they are expected to reach an achievement outcome (Dweck, 1986). Early conceptualizations of the topic outlined two distinct styles – learning and performance goals. Dweck (1986) defined learning orientation as that of individuals that approach a task with the intent of learning and increasing competence for their own sake, and performance orientation as that of individuals who approach a task with the intent of gaining positive (or avoiding negative) judgments of their result. Performance goal orientation has further been distinguished between the ideas of gaining positive or avoiding negative judgments of performance such that performance approach goal orientation has been defined as “the desire to prove one’s competence and to gain favorable judgments about it” (VandeWalle, 1997, p. 1000), and performance avoid goal orientation has been defined as “the desire to avoid the disproving of one’s competence and to avoid negative judgments about it” (VandeWalle, 1997, p. 1000).

For the purposes of this study, learning goal orientation and performance avoid goal orientation are of note. Learning goal orientation is generally behavior that is not motivated toward some end state – simply behavior for the sake of behavior. It is important to note the distinction between behaviors driven by an individual’s learning goal orientation and those driven by learning goals (Seijts & Latham, 2001). This links well to the ideas of exploratory behaviors in that they are behaviors without motivation stemming from a concrete end state. Performance avoid goal orientation is also potentially related through its effects on fear of failure (VandeWalle, 1997). Individuals with goals have concrete methods by which they can evaluate themselves (against the

goal) while those in do your best conditions do not (accept against their own prior performance).

Due to the non-outcome focused nature of exploratory behavior – the idea of exploration simply for the sake of gathering new knowledge, it is possible that there may be a linkage between exploratory behavior and goal orientation. Hints of this linkage can be found in the goal setting literature (Seijts et al, 2004). Specifically:

*H4a: Those high on learning orientation will engage in more exploratory behaviors than those low on learning orientation.*

*H4b: Those high on performance avoid orientation will show a stronger inhibiting effect of goals on exploratory behavior due to desire to reach those goals and avoid failure.*

## METHOD AND RESULTS

### *Study 1: Pilot*

To explore the feasibility of the proposed manipulations and hypotheses a small pilot was completed using a reduced version of the final method on a different but conceptually similar task.

#### *Participants*

Participants consisted of six graduate student volunteers from the psychology department of a large Midwestern university. None of the participants had knowledge of the task or principles underlying the task prior to the experiment. Participants were exclusively Caucasian and predominately male (66%).

#### *Procedure*

The task in this study was a k-armed bandit with four arms presented to participants through a computerized interface. Upon arrival at the lab, participants received a brief training on the task. This training involved explanation of how to operate the simulation as well as background information informing them that they should imagine themselves as a member of the human resources division of a large company. As a member of human resources, it was their job to hire workers from four different temp agencies – each of these agencies was represented in the task as an arm of a k-armed bandit as summarized by Kaebbling et al (1996) and discussed above. In this situation the four arms of the bandit are represented by four buttons. In order to remain consistent in terminology these four buttons will be referred to as the ‘arms’ of the task – as that is what they are representing in a reinforcement learning framework. A view of

the screen presented to participants in the goal condition and clearly showing these four buttons can be seen in figure 4.

Put simply, each participant was faced with four arms of which at any given point one could be activated. This activation led to success (score increments by one) or failure (score stays the same) based on predetermined underlying probabilities of the task. Over the course of the experiment participants completed 20 trials of the task, each task consisting of 50 'pulls' on the k-armed bandit. As stated, the k-armed bandit in this experiment had four arms (hereafter referred to as arm A, arm B, etc.). The probabilities on arms A, B, and D were fixed throughout the experiment at 70%, 20%, and 40% probability of return, respectively. The probability of arm C was differential in such a way that it was a poor arm in the first half of the experiment (30% return for trials 1-10), but the best arm in the second half of the experiment (90% return for trials 11-20).

Participants were also placed in one of two conditions for the experimental manipulation. In both conditions, text was displayed above the four bandits throughout the entire task. This text is the method by which the manipulation was enacted. In the first condition (do your best condition), the text read 'Do your best to maximize score.' In the second condition (specific, difficult goal), the text read 'Achieve a score of 33 out of 50.' A score of 33 out of 50 (i.e. a 66% return rate) represents a specific, difficult goal in the fact that in order to achieve it (in the first 10 trials), one must invest nearly all their pulls on the highest probability arm (arm A). No other arm in the first ten trials would deliver on this goal. As such, the first half of the task rewards exploitation on arm A. In the second half of the experiment, exploitation on arm A is not an optimal strategy, as

exploration can reveal that arm C is in fact a better choice and where effort should be invested.

After the task participants were also verbally debriefed by the experimenter as well as posed a series of task knowledge questions such as ‘what arm had the best probability?’

### *Results and Discussion*

Two variables of interest were examined as part of this pilot study: probability of pulling any given arm on any given trial, and performance. Probability was computed by counting the number of times a participant pulled any given arm during a trial, and then dividing by 50. This yielded four percentages for each trial. Performance was the number of returns a participant accumulated during a trial.

Results for probability were consistent with hypotheses, as a repeated measures analysis showed that the three way interaction of time, condition (goals vs do your best), and probability of arm choice was significant ( $F(57,304) = 2.359, p < 0.001$ ). This finding is also apparent in figures 1 and 2. Figure 1 displays the probabilities for the ‘do your best’ condition and reveals that participants in this condition were more likely to find the better arm in the second half (compared to the goal condition represented in figure 2).

As well, results for performance were also consistent with hypotheses. While a repeated measures analysis showed that the interaction of performance and condition was not significant ( $F(19,76) = 1.25, p = 0.239$ ; represented in figure 3), significant differences were found on the last two trials after the participants in the ‘do your best’ condition explored enough to determine the better arm to pull ( $F(1,4) = 11.63, p < 0.05, F$



= 10.321,  $p < 0.05$ ). Additionally, while the first few trials were marked by exploration on the task in both conditions due to the need for all participants to learn the probabilities, the last few trials of the first half represented a situation where exploitation was the favored strategy. Performance differences in these trials reached marginal significance on one of the trials (trial 8,  $F(1,4) = 4.985$ ,  $p = 0.089$ ), suggesting that a difference might be present and require more statistical power to reveal (or dismiss).

The results of verbal debriefings also revealed that of the participants in the ‘do your best condition’, all were able to respond to the question of ‘which was the best arm’ with the split answer of A in the first half, C in the second. Only one participant in the specific, difficult goals condition was able to make close to this distinction, stating that A seemed better at first, but that at some point it seemed like C might have gotten to about the same probability as A. The other participants in the specific, difficult goals condition believed that A was the best arm throughout the entire task.

Overall, this pilot study supports the idea that specific difficult goals limited exploration on this k-armed bandit task; these initial results are strongly consistent with the proposed hypotheses. However, the bandit task as presented is incapable of addressing the concern of the need for a complexity manipulation that produces results counter to that of Wood et al (1987). In order to address the concerns, it is necessary for the task to be manipulated in both complexity and need for exploration in a way that does not confound the two. As well, goals must have a way to work on the task through their effects on effort. The bandit task, since there are a set number of pulls in each trial, does not allow differential effort to manifest itself over different conditions. Thus, we have

determined that a similar task from the area of exploration/exploitation literature would be better suited to find the results we are proposing.

### *Study 2: Pilot 2*

The task chosen to better examine the relationship between exploration and task complexity is a computer simulation modeled after the ideas of foraging theory. In the simulation, the participant interacts with a two-dimensional world which they can navigate in search for resources (in foraging theory these would be food, water, etc). The world, except for a small area within a radius around the participant (in this case 3 squares in any direction), is hidden from view until explored. The participant is able to move in any direction with the arrow keys to explore the world and find resources, which are scattered throughout the map in patterns dependent upon experimental manipulation. Once a participant encounters a resource they then must devote time to collecting it (using the spacebar), accumulating a score. There is both a move delay and a collect delay such that participants cannot move again after moving for .5 seconds, or collect again after collecting for .2 seconds. The particulars of the experiment are described below.

### *Participants*

Data was collected from 47 undergraduate students at a large Midwestern university. Participants were roughly evenly split on gender and predominately Caucasian, all between the ages of 18 and 24. Participants received credit in the psychology research pool for their participation.

## *Design*

The purpose of this pilot was to determine at what level the goals should be set for conditions involving difficult specific goals, as well as show that the different world maps being used would differ correctly not only on objective complexity, but also on perceptual complexity, while not differing on difficulty.

For these purposes, all participants engaged in a 'do your best' version of the task. Additionally, participants were placed into one of three between subject conditions, each consisting of a different version of the world with which they interacted. Each world consisted of a 20x30 square space; the differences arose in the placement of resources throughout each individual world.

In condition 1 (World 1, figure 4), there were small pools of resources within sight of each other near the place where the participant starts. The participant starts near the 'bottom' of the world, at the center. In this way they are surrounded by the dispersed pools of resources. This section of small pools of resources accounts for slightly more than half of the task's world. In the rest of the world there is one large (relative to the smaller pools) pool of resources that the participant can work on without moving. This pool can be said to thus be a richer source of resources.

In condition 2 (World 2, figure 5), the world was mirrored such that participants still start at the 'bottom' of the screen, but this part of the world is now empty except for the pool of rich resources located to their immediate right. They are one square away from being in sight of the larger pool of resources, and the small pools of resources are on the top half of the screen.

In condition 3 (World 3, figure 6), a complexity manipulation was used relative to the first condition. Participants started in the same place as in the first world, at the center of the 'bottom' of the screen, surrounded by dispersed resources. In order to make the task more complex in the most clear and powerful way possible, dynamic complexity was manipulated. It was manipulated by making the pools of resources non-stationary across trials. The smaller pools of resources change position within their half of the screen, as does the larger pool of resources on the top half of the screen. This occurs by random placement with each of three areas of the screen being 33.3% likely of having the pool of resources on any given trial. As the resources on the bottom half of the screen were generated pseudo-randomly, a new set was simply generated for each trial.

As the participant moves around each of these three worlds more and more of the world becomes visible. Any square that has been seen during that trial remains visible for the remainder of that trial, but participants return to the start and to their initial scope of vision at the start of each trial. Examples of the amount of the world that can be seen early and mid trial by participants can be found in figures 7 and 8.

### *Procedure*

Upon arrival at the lab, participants were randomly assigned to one of the three conditions of the experimental design. Participants then received a brief training on the task (approximately 5 minutes). This training informed them of what they would be doing during the session (how many trials they would complete, etc), and how to move around the world and collect resources. This training was uniform across conditions. Participants then engaged in a one-minute practice trial on a much smaller world in order to be sure that they knew how to move and collect resources properly. Following this,

participants completed five two-minute trials on the task as outlined above. At the end of each trial, participants received score feedback on the previous trial.

Upon completion of the experiment, participants completed a series of exit measures including a number of measures assessing both the complexity and difficulty of the different versions of the task.

### *Measures*

Perceptual Complexity – Perceptual complexity was measured using a three item scale containing one question about each of the different facets of complexity as outlined by Wood (1986). These three items, measuring component, coordinative, and dynamic complexity, respectively, are: “Rate the number of unique physical actions that had to be completed in order to accomplish the computer task you completed during this experiment”, “Rate the degree to which actions which you completed during the computer task were linked to each other in terms of frequency, timing, and coordination”, and “Rate the degree to which characteristics of the task and the actions best suited for the task varied from trial to trial.” Each item was measured on a 5-point scale with tailored anchors.

Perceptual Difficulty – Perceptual difficulty was measured using a one item scale meant to address how difficult and challenging participants believed the task was. This item was: “How difficult was the computer task in this study?”, with anchors on a 5-point scale ranging from very low difficulty to very high difficulty.

### *Results*

The purpose of this pilot was twofold. First, participants were placed in this version of the task in order to set a goal for those in the specific, difficult goal condition.

Keeping with the goal setting literature, a goal was calculated for the 85<sup>th</sup> percentile of those in the do your best condition. However, in order to limit the potential effects of score boost by exploration as would occur later in the task, the goal was calculated only on the first trial of the do your best condition. For world 1 this produced a goal of 107.9, which for practicality was rounded to 108. As there are 28 patches of resources, each containing 4 resources on the bottom of the world – that which is easily exploitable without exploration – this goal of 108 out of a maximum of these 112 is not unreasonable or impossible, but still quite difficult. For world 2 this produced a goal of 142.6 (rounded to 143), which is reasonable due to the fact that participants start on the part of the world with a much richer area of resources. For world 3 this produced a goal of 108.0, which is in practicality identical to the goal set for world 1. This is likely due to the fact that without exploration world 1 and world 3 are both randomly generated to identical specifications in the first trial.

The second purpose of this pilot was to show that results of perceptual complexity would be consistent with our assessments of objective complexity. It was proposed that neither world 1 and world 2 should be more or less complex than the other – the transformation between them is a non-complexity based transformation. However, world 3 was transformed from world 1 with a complexity based transformation, specifically a transformation involving dynamic complexity. Thus, world 1 and 2 should be perceptually similar in terms of complexity, while world 3 should differ from both in dynamic complexity. This is precisely what was found.

A one-way ANOVA on all three groups showed no differences on component ( $F(2,56) = .464, ns$ ) or coordinate ( $F(2,56) = .095, ns$ ) complexity, but did show a

difference on dynamic complexity ( $F(2,56) = .3828, p < .05$ ). Post hoc tests revealed that this effect was manifested in mean differences between world 1 and 3 and world 2 and 3, with participants in world 1 ( $M = 2.38, SD = .94$ ) and world 2 ( $M = 2.31, SD = .79$ ) reporting significantly lower dynamic complexity than those in world 3 ( $M = 3.13, SD = 1.06$ ).

Additionally, the result that perceptions of difficulty did not differ between the three groups ( $F(2,56) = 1.474, ns$ ) was found. Thus, while the sample size is relatively small, it hints that our complexity manipulation is not confounded with a difficulty manipulation.

### *Study 3: Test of Goal-Setting Effects*

#### *Participants*

Data was collected from 332 undergraduate students at a large midwestern university. Of these participants, 47 were those of the previous pilot study included because of the identical nature of their cells of the experimental design. Participants received credit in the psychology research pool for their participation or extra credit in a psychology class. Of these 332 participants, 20 participants were removed from the dataset for having large amounts of incomplete data on the task. This was due to a number of possible random factors such as computer error. Of the remaining 312 participants, 2 were removed for having a score of less than 50 on any trial. As only two participants received scores this low on any trial it is to be taken as a showing of low motivation or a misunderstanding on the task, not of any other important process. Thus, the final data set consisted of 310 participants.

w

w

D

su

re

ex

Th

wa

be

stu

wo

Pr

onl

onl

and

wo

to t

Age

Eac



Of these 310 participants, the majority (65%) were female. As well, the majority were Caucasian (81%). The mean age for participants was 20.05 years, and the majority was between the ages of 18 and 22.

### *Design*

The design of the study was a three by two between subjects. The first between subjects experimental manipulation was identical to that of the pilot studies – one group received difficult specific goals (difficult specific goal condition) throughout the experiment while another received do your best goals (do your best goals condition). These goals will be displayed on the screen in the same location as the text ‘do your best’ was displayed to those in the ‘do your best’ goal condition of the initial pilot. The second between subjects manipulation was identical to the manipulation in the second pilot study. That is, participants performed the task in one of three worlds, world 1 (figure 4), world 2 (figure 5), or world 3 (figure 6).

### *Procedure*

Trait level measures (e.g. trait level curiosity, goal orientation) were measured online prior to the study in order to obtain results unbiased by any lab experience. This online measurement took place immediately after participants signed up for the study, and was thus separated from the lab session by anywhere from one day to one week’s worth of time.

Upon arrival at the lab, participants received a brief training on the task identical to that of the second pilot study, as well as an identical one-minute practice session. Again identical to the pilot participants next completed five two-minute trials on the task. Each of these trials was followed by score feedback and a series of process measures. A

final set of measures (perceptual complexity, perceptual difficulty) was completed after the last set of process questions.

In conditions 1 & 2 (world 1, figure 4), there were small pools of resources within sight of each other near the place where the participant starts. By devoting time to moving from each of these to the next, the participant stood a chance of reaching the difficult goal presented to them. However, as found in the pilot, this was still a difficult prospect. This section of small pools of resources accounted for roughly half of the task's world. On the other half of the world, there was one large (relative to the smaller pools) pool of resources that the participant could work on without moving. Thus, in this first version of the task, those that explored and found the larger pool should outperform those that spend their time exploiting. Thus, if difficult specific goals do in fact limit exploration as hypothesized, this should produce a negative goal setting effect, as those with goals will exploit and those without (do your best conditions) will explore (hypothesis 2b).

In the second version of the task (conditions 3 & 4, world 2, figure 5), the world was mirrored such that participants started near the larger pool of resources. Thus, in this version of the task - with equal complexity to the first - any exploration (time spent not collecting resources) should lead to a decreased score. Those that exploited on the task (difficult goal condition) should thus outperform those that spent time exploring. Again, if difficult specific goals limit exploration then those that are given goals will exploit and those without will explore. This should produce a positive goal setting effect (hypothesis 2a).

For the third version of the task (conditions 5 & 6, world 3, figure 6) a complexity manipulation was used relative to the first condition. In order to make the task more complex in the most clear and powerful way possible, dynamic complexity was manipulated. It was manipulated by making the pools of resources non-stationary across trials. The smaller pools of resources changed position within their half of the screen (by trial), as did the larger pool of resources on the top half of the screen. This occurred by random placement, with each of three areas of the screen being 33.3% likely of having the pool of resources on any given trial (figure 6). Thus, while participants gain some benefit of exploring and finding the pool of resources, they can also never gain concrete knowledge about its position through this exploration. Through this effect the benefits of finding the pool (conditions 1 & 2) are offset by the price of finding it (i.e. time spent exploring). While this has little effect on those that exploit (some small pool should by chance always be in the participant's sight when they start), it will have an effect on those that explore in that more time is required (on average) while exploring to find the larger pool. This should lower the score of those that explore, thus shifting the task to one that again favors exploiting. Since it is hypothesized that those that are given specific, difficult goals will be those more likely to exploit, this should shift the benefits of the task back toward their favor – on a task that is unquestionably more complex by Wood 1986 criteria than the simpler version that produced a negative goal setting effect. It is unclear then if this will be powerful enough to produce a positive goal setting effect or if it will simply shift the negative goal setting effect to lesser negative or neutral.

## *Measures*

Trait level goal orientation – Goal orientation is defined as the method and motivation behind individuals' work toward different types of goals. Two aspects of goal orientation are proposed to have effects in this study; learning goal orientation and performance avoid goal orientation. Both were measured using a nine item measure developed by VandeWalle (1997). Five items are designed to measure learning goal orientation and four items are designed to measure performance avoid goal orientation. Example items are “I often look for opportunities to develop new skills and knowledge” (learning) and “Avoiding a show of low ability is more important to me than learning a new skill” (performance avoid). This measure has been shown to have high internal consistency ( $\alpha=.88$  for learning,  $\alpha=.83$  for performance avoid), and reasonable test-retest reliability ( $\alpha=.66$  for learning,  $\alpha=.57$  for performance avoid). Validity of the measure has also been shown through correlations with similar measures (VanderWalle, 1997).

Trait level curiosity – Curiosity was measured using a four item scale from Kashdan, Rose, and Fincham (2004) designed to measure a trait level model of curiosity involving exploration. Example items include “Everywhere I go, I am out looking for new things or experiences” and “I frequently find myself looking for new opportunities to grow as a person (e.g. information, people, resources)” (Kashdan, Rose, and Fincham, 2004). This measure has been shown to have a reliability of  $\alpha=.74$ , and a test-retest reliability over a month period of  $r=.78, p<.001$ . It has also been shown to possess convergent validity with other curiosity measures, openness to experience, and positive affect and discriminant validity with conscientiousness and agreeableness (Kashdan, Rose, & Fincham, 2004).

scale

by W

had

exp

con

and

the

tai

sc

Th

ch

in

d

c

c

n

v

v

f

s

**Perceptual Complexity** – Perceptual complexity was measured using a three item scale containing one question about each of the different facets of complexity as outlined by Wood (1986). These three items are: “Rate the number of unique physical actions that had to be completed in order to accomplish the computer task you completed during this experiment”, “Rate the degree to which actions which you completed during the computer task were linked to each other in terms of frequency, timing, and coordination”, and “Rate the degree to which characteristics of the task and the actions best suited for the task varied from trial to trial.” Each item was measured on a 5-point scale with tailored anchors.

**Perceptual Difficulty** – Perceptual difficulty was measured using a three item scale meant to address how difficult and challenging participants believed the task was. These three items are: “How difficult was the computer task in this study?”, “How challenging was the computer task in this study?”, and “How easy was the computer task in this study?” (reverse-scored). Each item was measured on a 5-point scale.

**Exploration (operationalization)** – Exploration can be operationalized in two different ways: 1) how much of the map (by percentage) was visible to the participant, cumulatively across trials, during any given trial, and 2) whether or not the participant crossed the (conceptual) line between the side of the map containing the dispersed resources and the side of the map containing the rich resources. From this point on these variables will also be referred to as EX1 (percentage of map visible) and EX2 (crossing). Where appropriate and in situations when these two variables do not differ in their effects they will be addressed simply as exploration. When they do differ, or if the effects are similar but differ in interpretation, results will be discussed separately.

## *Results*

Results will be ordered as follow. To begin, manipulation checks involving complexity and benefit of exploration will be examined. Following this, individual difference effects will be tested in order to determine whether or not these individual differences should have a role in further analyses. Following this, full tests of the proposed model will be tested.

Manipulation Check – In order to test whether or not participants found any world to be more complex or difficult than the others, a series of analysis of variance was completed on self report complexity and difficulty measures by world. While an effect was found for dynamic complexity in pilot data such that participants self reported world 3 to be more dynamically complex than worlds 1 and 2, this effect was not found in the full sample. Further, no effects of any level of perceptual complexity or perceptual difficulty were found.

Additionally, in order to determine if exploration was more or less useful in any world as predicted (exploration beneficial in worlds 1 and 3, detrimental in world 2), a number of linear regressions spilt by world were performed predicting score at each trial by the amount of exploration by that trial. The results of these regressions can be found in table 1 and show that exploration is positively related to performance in worlds 1 and 3 and negatively related to performance in world 2, supporting the need for exploration manipulation.

Main Analyses - In order to test hypotheses 3, 4a, and 4b, a series of linear regressions were performed. These analyses tested the predictive power of trait level curiosity, trait level learning goal orientation, and trait level performance avoid goal

orientation on exploration. Both operationalizations of exploration were tested, with the 'gone across' variable in table 2 and the cumulative exploration variable in table 3, both tested at each individual trial. No significant effects were found. Thus, hypothesis 3, 4a, and 4b are not supported, and these individual differences will not be included in further analyses.

In order to begin testing remaining hypotheses, repeated measures analysis was first performed on performance (score) and exploration (as outlined above) by trial, goal condition (do your best vs. specific difficult), and world (world 1, 2, and 3). For performance there are a number of significant effects, including the three way interaction of time by goals by world ( $F(8,1216)=3.01, p<.01$ ). This three way interaction is not significant for exploration (EX1,  $F(8,1216)=1.48, ns$ ; EX2,  $F(8,1216)=1.06, ns$ ). The remaining results of these two analyses can be found in table 4 and will be examined in turn as a means of explaining these more complex interactions and testing the predicted hypotheses.

First, in terms of time effects, participants' performance was generally increasing, with a mean score of 106.19 on the first trial and a mean score of 135.05 on the last trial. Conversely, participants' exploration was generally increasing over time, as more participants were more likely to explore the map as time went on. While performance is explained by participants getting better on the task, exploration likely rests on the fact that at least some participants will see something new each trial through exploration. As participants see new parts of the map their cumulative exploration increases, as does the chance that they will cross the dividing line between sections of the worlds.



There is also an effect of world on both performance and exploration. Further examination of this effect shows that participants in world 2 scored significantly higher on average than those in world 1 ( $F(1,211)=37.23, p<.01$ ) and world 3 ( $F(1,209)=39.99, p<.01$ ). This effect is likely due to the fact that participants in world 2 start near the richest area of resources. Thus, even if they do explore, they are more likely to return to this area to maximize their score. Further, those collecting the rich resources simply have less distance to travel to get to them. This score differential is also consistent with pilot data on do your best only conditions (pilot study 2). Additionally, differential exploration effects appear. Participants in world 2 are more likely to travel across to the other side of the map than in conditions 1 or 3. This is due in part to the nature of the worlds – the division between the distributed resources and the rich resources is at the two-thirds point of the map. That is, the area with distributed resources takes up two-thirds of the map, and the area with rich resources takes up one-third. Thus, participants in world 2 start closer to the line due to the fact that they start on the side with rich resources. In terms of cumulative exploration, participants in world 3 tend to explore more than participants in the other worlds. This is likely due to the changing nature of the world – participants need to explore more areas of the world in order to continue doing the same thing, whereas participants in the first two worlds can maintain the same pattern across trials.

The non-significant effect of goals on performance across all worlds is to be expected – there are different predictions for each world in terms of goal-setting effects. While in world 1 participants with do your best goals should perform better than those with specific, difficult goals, participants with ‘do your best’ goals in world 2 should

perform worse than those with specific, difficult goals, effectively canceling out effects across conditions.

In a direct test of hypothesis 1, the effect of goals on exploration was tested. This effect is in the predicted direction; those given difficult specific goals explore less (as measured by both of the exploration variables) than those given do your best goals. While it is not measured directly, it can be assumed that those given specific, difficult goals exploit more than those given 'do your best' goals. These results are supportive of hypothesis 1.

The next result is the interaction of time by world on performance and exploration. In terms of performance this result is somewhat unclear in interpretation. It appears that world 2 has a slightly larger slope over time than world 1 or world 3, and world 3 has less consistent slope over time. This first effect is due to the fact that people who eventually find the rich resources can score higher in world 2 than in the others. The second effect is likely due to the shifting of the rich resources, as there is a chance that a participant might find them right away during a trial or spend opportunity cost again searching for them. The effect on exploration seems driven at the 'gone across' variable by earlier effects of world 2 (participants more likely to cross early in the experiment), and at the cumulative exploration variable by world 3. Specifically, those participants in world 1 and world 2 seem to slow or stop exploring new parts of the map by the third trial, whereas participants in world 3 continue to keep exploring new areas. This is due to the fact that those that have found the rich resources need to search new areas in order to find it. They pay for this continuing need for exploration in continuing opportunity costs of it.

The time by goals interaction on both performance and exploration is also somewhat difficult to interpret. Those given 'do your best' goals seem to have a slightly higher slope in terms of cumulative exploration and going across the map over time, though in terms of performance these same participants seem to have a more varied slope over time. It is suspected that without taking into account what is happening in each different world this result may simply be nonsensical.

In a direct test of hypotheses 2a and 2b, analysis was performed to test the effect of a world by goals interaction on performance. The effect on performance is significant and shows that in world 1, participants perform better when given difficult specific goals relative to do your best goals, whereas in world 2 participants perform worse when given specific, difficult goals relative to 'do your best' goals. These two results are supportive of hypotheses 2b and 2a, respectively. While no direct hypotheses have been made about world 3, it would stand to reason from Wood, Mento, & Locke (1987) that due to the fact that world 3 is simply a more complex version of world 1, the goal setting effect in world 3 would be predicted by their results to be tending toward lower positive or greater negative than whatever the goal setting result is in world 1. In world 3 there is no difference in score between those given 'do your best' and specific, difficult goals. This is likely due to the fact that the effects of performance play out over time with the gap between goal conditions widening with each trial. Since those participants in world 3 must effectively re-learn the world every trial they never have a chance for this effect to play out – thus a neutral goal setting effect is the result. This is counter to the prediction of Wood, Mento, & Locke (1987) – that an increase of complexity would shift a goal

setting effect in the negative direction – and consistent with exploration as an explanation of goal setting effects.

The effect of world by goals on cumulative exploration is not significant, but this is to be expected. There is no reason that the effects of goals should differ by world. The significant effect on exploration (in terms of the 'gone across' variable) seems to again be driven by those participants in world 2. Those in the goal condition who spend their time exploiting are somewhat likely to exhaust the rich resources before the very end of the trial, at which point they will cross to the side with distributed resources. If this is in fact the case, it would represent continuing exploitation of the world, and not a shift to exploration, as they are only exploring as a means to continue what they know how to do.

One final exploratory analysis was completed on a subset of the participants of world 3. It was observed during the coding of data that some participants who had found the rich resources on previous trials crossed on following trial to where they once were in order to collect them. Finding the rich resources no longer in the spot they last left them, participants had a choice to explore more of the top. This increased exploration was at an opportunity cost of not collecting for an additional and unknown time period, if anything was even ever to be found. Participants may have only found the rich resources the trial prior, or thought that removal of them was a valid possibility of the task design. Put simply, participants were placed again in a situation where they were unsure if exploration would pay off. Participants in this world knew that they could return to the bottom and exploit the dispersed resources there, as they would have had to have passed through them from the start to the top of the screen. As a further test of our hypotheses it should be the case that this new subset of participants should be subject to the same

effects – those who are given difficult specific goals should be more inclined to return to exploiting and those given do your best goals should be more inclined to continue exploring. Statistical power is diminished by the fact that this analysis is on groups within subsets of the population on only one of the worlds (sample size for this analysis in trial 1 is only  $N=8$ ), but as the trials progress more and more participants are likely to be crossing sides of the map.

The only marginal effect found is on trial 5 ( $F(1,68) = 2.96, p=.09$ ), but it is in the right direction – 47% of those in the difficult specific goal condition returned to the dispersed resources compared to only 27% in the do your best goal condition. In effect, the difference between exploring and exploiting in the difficult specific goal condition was akin to a coin flip, as each participant had a roughly half and half shot of continuing to explore. Chances of returning in the do your best goal condition were cut almost in half, as roughly 3 out of every 4 participants chose to stay and explore. Further, this impacted these participants scores in a way which produced a negative goal effect (as should be the case where exploration is beneficial); those who chose to continue exploring outscored those who returned to the dispersed resources by a mean difference of over 50 points ( $F(1,68) = 48.10, p<.01$ ). While this is not a definitive replication of the above results, it suggests that with these results can be generalized to other situations where individuals face choices between exploration and exploitation.

### Discussion

In summary, a number of hypotheses were supported and a number were unsupported. The hypothesis relating difficult specific goals to a decrease in exploration were supported – those participants given difficult specific goals explored less than those

given do your best goals. Further, by changing the nature of the world in order to alter the task's need for exploration differential goal setting effects were created. In the first world, where exploration was beneficial, a negative goal setting effect was found. This effect was reversed in the second world where exploration was detrimental, producing a positive goal setting effect, and negated in the third world where exploration was started anew every trial, producing no goal setting effect. However, those hypotheses relating to individual differences as predictors of exploration were unsupported. Trait level curiosity, trait level learning goal orientation, and trait level performance avoid goal orientation failed to predict exploration at any trial of the experiment.

Globally, these results speak directly to the purpose of this paper, namely to “show that the way in which specific difficult goals may be causing both positive and negative goal setting effects is through their effects on exploratory behavior”. Complexity was not a manipulation between world 1 and world 2, but rather by manipulating the need for exploration in those worlds was the goal setting effect reversed from a negative to a positive effect. When a complexity manipulation was put in place – making the third world more complex relative to the first world – the negative goal setting effect that had been found was neutralized, contradicting the current stance on goal effects. From Wood, Mento, and Locke (1987) the goal setting effect should move from high positive to low positive when complexity is increased, not from negative to neutral.

In short, and in summary, difficult specific goals were shown to decrease exploration by driving exploitation. If exploration is needed and beneficial on a task this suppression effect will produce a negative goal setting effect. If exploration is not

needed, and rather exploitation is beneficial, this suppression effect will produce a positive goal setting effect. If exploration is beneficial but interrupted frequently enough by dynamics of the task neither effect will gain enough temporal traction to establish itself, and a neutral goal effect will result.

These findings are consistent, as shown earlier, with previous goal setting research in that complexity is likely – in many studies – to be driven by these need for exploration effects. If a task is simple exploration is likely to be detrimental to performance – thus difficult specific goals that suppress exploration will drive the numerous strong positive goal effects found on simple tasks. The more complex a task is, the more likely it is that exploration will be needed for high performance. When difficult specific goals are given in this situation, participants will still explore less than participants given do your best goals, driving a reduction of the positive goal setting effect. This reduction is again consistent with the findings of Wood, Mento, and Locke (1987).

What cannot be explained by the findings of Wood, Mento, and Locke (1987) are situations where negative goal setting effects arise, especially on simple tasks, or where the relationship between complexity and the level of the goal setting effect operates opposite to their predictions. The simple tasks where this negative goal setting effect occurs tend to be insight based in some way or another and thus possess a need for exploration. Huber's (1985) maze task and Sweller and Levine's (1982) maze task are both good examples of this where such an effect is found. Both of these cases can be explained using exploration as an explanation. Using this new theory as a template for experimental design also allowed both of these cases, including the reversal of the

relationship of complexity to goal setting effects, to be illustrated empirically in this experiment.

One question that should fall out of these results rests on the non-significant findings in terms of relevant individual difference measures predicting exploration in the task. Neither of the operationalizations of exploration was affected by any of the trait measures: curiosity, learning goal orientation, and performance avoid goal orientation. While Seijts and Latham (2001) showed that learning goals promote exploration, it was an extension of this result to suggest that individual goal orientation may also drive this effect. While it still may be the case that it does in weaker situations, it may be that the situational effects in this case were simply too strong. This would also explain why curiosity failed to predict exploration. It is possible that the trait measures taken were too distal, and that a collection of state measures during the experiment may have had a more proximal linkage to the behavioral measures and thus shown an effect. It is also true that while the operationalizations of exploration are the best that can be achieved in this task they are by no means perfect representations. More than anything it may be that the situation presented is fairly strong and decreased the effects of distal trait measure linkages to behavior.

Additionally, the worlds in this task were small enough that some participants may have crossed from one side of the world to the other haphazardly and driven only by chance. This was clearly not the case for all participants, as effects were found, but the size of the world may have led to more noise in the data than if participants had completed the task with substantially larger areas to explore and exploit. World 2 is likely to have fared the worst in this situation, as participants who moved directly to the



rich resources and followed a orderly pattern of collection had a reasonable chance of being able to collect all of the resources contained within. Faced with a now empty half of the world they would likely cross to the other side not out of exploration but out of lack of anything more to exploit. While this is problematic, it is not counter to the results, as if this did happen it would only reduce the power of the effects that were found, as people who were exploiting might inadvertently have higher exploration scores while still (technically) exploiting. The fact that this would be more likely toward the end of the task (as people become more skillful in collection) might actually explain the diminishing linkage between exploration and performance late in the experiment in world 2. Thus, while it is not a detrimental sticking point of the current research, future work on this task is likely to be benefited by the use of larger worlds.

### *Implications*

Complexity – Unlike previous studies and meta-analyses, this study was able to manipulate objective complexity in a way befitting Wood's (1986) original mathematical ideals. Due to the fact that only dynamic complexity was altered through the task design, it is possible to claim that world 3 was in fact ordinally more complex than world 1.

While this only mapped to individual's perceptions of the task in the pilot and was not found on the larger population, it still has implications for task complexity theory in providing a framework by which to make manipulations to task complexity. As complexity theory based on Wood (1986) currently stands, a scientific statement can only be made about the relative standing of one task to another when only one variable of the task (component, coordinative, or dynamic complexity) is changed, holding the other two constant. Instead of making claims across widely disparate levels of tasks, more work

should be applied in understanding the effects of each of these different types of complexity. Work on this level may eventually lead to the discovery of Wood's (1986) elusive  $\alpha$ ,  $\beta$ , and  $\gamma$ .

Exploration and problem representation – The main implications to the area of problem representation deal with the distinction of insight problems from other types of tasks. Insight problems are unique in the way in which they need to be solved – a solution must come from a distinctly different area than where one might start. For example, a golfer may spend years and thousands of hours swinging golf clubs in order to see small but steady improvements in performance. Conversely, a golfer may only need to make one change of mindset to realize they have been holding the club upside-down, and likely see immediate improvement in performance after their change. The first is an example of a problem solved by exploitation, the second a problem solved by exploration. Providing a specific, difficult goal to the first individual may help drive their performance, providing a specific, difficult goal to the second may simply frustrate them.

Insight problems in the real world are also likely to be without known solution. These types of problems – in whatever form they may take – may then be benefited by using a 'do your best' goal approach as opposed to specific, difficult goals.

Exploration and goal setting theory – In effect this entire paper is about the implications of exploration as it applies to goal setting theory. If a task can be said to require exploration, then specific, difficult goals should not be used in conjunction with it – performance will be decreased. If a task can be said to require exploitation, then specific difficult goals should be used in conjunction with it – performance will be improved. This has been shown on a simple task, and with future research on tasks of

medium and complex tasks it stands to reason that this advice can be applied without worry of complexity; only with concern for need for exploration.

Earlier in the paper the idea of mastery and performance goals was discussed briefly, somewhat dismissing mastery goals at the time in order to focus on performance goals. Mastery goals, through their effects on learning, may be a proper substitute for 'do your best' goal when specific, difficult (performance) goals are found to be detrimental due to their effects on exploration. More research is needed that specifically links these ideas to the processes of exploration while taking into account the amount of exploration needed on the task.

The use of mastery goals comes with a caveat, as well; the use of specific, difficult mastery goals may in the end have similar effects on unwarranted exploration the same as specific, difficult performance goals. That is to say that while an experimenter may be able to use a mastery goal to provoke a participant to explore parts of their world that they otherwise might not find, this provocation may deter them from other, more self-motivated and task-irrelevant exploration. This idea may not be worrisome in a lab setting where the experimenter likely knows the best way for a task to be completed, but in the real world, or with more complex tasks, any specific, difficult goals – mastery or performance – may stop an individual from 'thinking outside the box', so to speak.

Exploration and reinforcement learning – The area of reinforcement learning has already grappled with the ideas of exploration and exploitation to a large degree. One of the largest implications for this area is that goals can seemingly be used as a trigger to shift this balance. Almost impossible to predict (or perhaps imagine!) would be the

application of these ideas to reinforcement learning as told by automated agents, or computers. Is it possible to model this behavior in a way that a computer can follow? Are these drives of goals and curiosity innate to the human consciousness, or can they in fact be adequately replicated? These questions are obviously far beyond the scope of this paper, but answers to them may help to better understand just exactly how individuals are being driven in terms of goals and exploration.

Other implications – The fact that goals were set at the 85<sup>th</sup> percentile is in proper keeping with the goal setting literature. The fact that they were normed from the first trial exclusively is less in keeping. While the reasons for this use are sound, it is easy to imagine that using a specific, difficult goal from a later trial (and thus a higher score) would have had a different effect on individuals. This is an important distinction, and relies very heavily again on the distinction between insight and non-insight problems. Goals were normed from the first trial in this study because it was desired that participants strive to do the best they could do *while exploiting*. The insight to this problem was that one could stop exploiting and explore. Participants fail to make this jump because they are focused on their goals. In this experiment it was almost too easy to find the rich resources; the problems of the size of the world have and will continue to be discussed.

If the goal was set higher than a participant could achieve through exploitation, they would be triggered that something was wrong. In this way, specific, difficult goals set with full knowledge and control of the world may actually drive a behavior that mimics exploration. It is argued that this would no longer be exploration – this would simply be a continuation of exploitation.

Practical implications – When these ideas are removed from a lab setting a number of questions and possible implications arise. First and foremost, as stated earlier in this paper, using the incorrect balance of exploration and exploitation in the real world can, at its absolute worst, result in death of a species. At the corporate level it can result in bankruptcy or dissolution. At the very least it can lead to outcomes more unfavorable than were met in the lab by undergraduate psychology students who were simply having fun playing a game. It is likely that exploration in the real world holds a larger risk than could be replicated in a lab. Despite this, it is suggested that the above findings should still hold, possibly with a number of qualifications to be uncovered by future research. One idea is that while the specific, difficult goals were used in this study were strictly outcome based, similar results may be found when specific, difficult time goals are placed on individuals or organizations. It may also be the case that many of the problems faced by organizations as a whole can be conceptualized as insight problems – problems which no one may have a solid answer to. Simply making slide rulers 100 times faster at 100 times the efficiency might stop the slow onslaught of the electronic calculator, but it will never stop it. If exploration is needed, limit the use of specific, difficult goals. If exploitation is needed, use specific, difficult goals.

### *Limitations*

This study was completed on college students in a psychology subject pool, on a specific task, and as such it may be the case that future replications on diverse populations are needed to show full generalizability. That said, many of the results of the goal setting literature have first been found in college populations participating in lab settings, and some replication has also been shown on a different task in this study (pilot

1) as well as in a different situation (exploratory analysis on those returning to the dispersed resources in world 3). It would be useful to replicate the results of this study on a task of both medium and high complexity in order to show results across the full spectrum of task complexity. While the task used in this study was a relatively simple one, results fitting with the hypotheses proposed have already been found on a wide spectrum of tasks (Burns & Vollmeyer, 2002; Huber, 1985; Ivancevich, 1974; Pritchard, Jones, Roth, Stuebing, & Ekeberg, 1988; Sweller & Levine, 1982). What is missing in these studies are measures of exploration both needed and completed on the task and the effects of goals on these measures – the foundation for which has been put forth in this study. The results of this study should be generalizable to any number of tasks as long as the amount of exploration needed on the task can be qualified. If exploration is needed, specific difficult goals will hurt, and if exploration is not needed, specific difficult goals will help.

In terms of limitations of the task itself, and as already mentioned, it might also have been possible that the size of the world in which participants acted imposed constraints on the total amount of exploration that participants could engage in. This might also have produced a constraint from the opposite end, as those exploiting may have eventually gotten good enough at the task to run out of things to exploit, prompting exploration.

The orientation of the world as well, with participants being able to find three ‘walls’ of the world fairly quickly may have prompted them to find the final wall through some other motivational mechanism. Individual differences such as need for cognitive closure or tolerance of ambiguity may well have predicted some of this behavior, and

future work on this task should either address this issue by ‘removing’ the walls or examining these processes.

### *Next Steps*

As the task used is in its early stages, much was learned though this study which may be useful both in development of the task as well as redesign to look at additional effects and relationships.

By increasing the size of the world, and thereby pushing the walls of the lower half of the world away from the participant, participant behavior may be drastically altered. Predominantly, this would open up whole new areas for them to explore, instead of – in effect – ‘funneling’ them toward the top of the screen. This sort of world could be useful in manipulating the benefits of exploration by giving a larger area of space that could be made beneficial or detrimental via exploration. Though this, participants could be given a wider range of possible behaviors. While this may complicate scenario design it allows the task to address a wider range of questions, such as how much participants will continue to devote to exploration that may not be working (e.g. sunk cost paradigms).

Exploration was also examined to a large extent in this study, but exploitation was not concretely operationalized (other than as the absence of some exploration). There are a number of ways that exploitation can potentially be operationalized in future work. The number of dispersed pools that they visit or even the redundancy between pools they visit from one trial to the next could both be viable ways at developing a variable to incorporate exploitation into future work and analysis.

There is also some potential need to show that goals created in different ways (e.g. proximal goals normed at each trial, distal terminal goals from final trials) have similar effects. While the world structures of world 1 and world 3 may need to be changed a bit in order to accommodate this change, world 2 is the place where goals seem to have the most leeway in their location on the score scale. Goals in world 2 have the potential to be raised by almost 40 points before hitting the exploitation ceiling of 180 points. This task, and this design of the task possesses an interesting point on the score scale – this exploitation ceiling that is the maximum that can be scored without exploring. This point on the scale may be a cut point where effects can be further changed. Future research on larger versions of world 1 or 3, or world 2 with different levels of goals may give more insight into this unique property.

### *Conclusions*

Negative goal setting effects, as well as situations in which the relationship between complexity and goal effects operate contrary to the predictions of Wood, Mento, and Locke (1987), cannot be explained under current theories with the exception of dismissing them as experimental flukes. This paper has shown that these effects are not flukes, but situations in which the need for exploration on the task does not follow the simple ‘more complex tasks need more exploration’ relationship. Complexity works sometimes – when the above statement holds true – but exploration is a deeper and more encompassing explanation of effects. While this paper was aimed at explaining these ‘flukes’ of goal setting, even positive goal effects are explainable by exploration, allowing exploration as an explanation of goal setting effects to account for all the types



of empirical results in the goal setting literature, not just those that fit preconceived notions.

## Appendixes

**Table 1: Regression of cumulative exploration on performance over time by world.**

World	Trial Number	R <sup>2</sup>	B	SE B	Standardized $\beta$
1 (N=100)	1	.16	108.38	25.04	.40**
	2	.27	117.88	19.61	.52**
	3	.32	150.55	21.85	.57**
	4	.33	162.69	23.43	.58**
	5	.34	179.73	25.19	.59**
2 (N=112)	1	.70	-116.56	8.04	-.84**
	2	.18	-70.97	15.92	-.42**
	3	.11	-62.89	19.27	-.32**
	4	.03	-33.37	20.72	-.17
	5	.04	-47.25	23.06	-.21*
3 (N=98)	1	.33	119.86	17.63	.57**
	2	.26	99.51	16.95	.51**
	3	.37	143.73	19.19	.61**
	4	.29	151.67	24.18	.54**
	5	.13	126.67	33.82	.36**

Note: \* - Significant at the  $p < .05$  level; \*\* - significant at the  $p < .01$  level.

**Table 2: Linear regressions of proposed moderators on exploration: results for hypothesis 3, 4a, and 4b (EX1); (N=310).**

	Predictor								
	Trait Curiosity			Trait Learning Goal Orientation			Trait Performance Avoid Orientation <sup>a</sup>		
	B	SE B	$\beta$	B	SE B	$\beta$	B	SE B	$\beta$
EX1 Trial 1	.002	.017	.008	.011	.018	.037	-.023	.017	-.112
EX1 Trial 2	-.002	.019	-.007	.017	.020	.052	-.017	.020	-.072
EX1 Trial 3	.005	.019	.018	.022	.020	.070	-.021	.020	-.089
EX1 Trial 4	-.004	.019	-.012	.020	.020	.064	-.025	.021	-.102
EX1 Trial 5	-.007	.018	-.024	.016	.019	.051	-.018	.020	-.076

<sup>a</sup> Note: Performance-avoid orientation was only tested on conditions in which specific goals were present (N=158). Tests on full sample (N=310) were consistent with these results.

\* Significant at the  $p < .05$  level.

Table 3: Linear regressions of proposed moderators on exploration: results for hypothesis 3, 4a, and 4b (EX2).

	Predictor								
	Trait Curiosity			Trait Learning Goal Orientation			Trait Performance Avoid Orientation <sup>a</sup>		
	B	SE B	$\beta$	B	SE B	$\beta$	B	SE B	$\beta$
EX2 Trial 1	.077	.055	.086	.054	.058	.058	-.083	.057	-.124
EX2 Trial 2	-.016	.058	-.017	.046	.061	.046	.017	.062	.024
EX2 Trial 3	-.031	.058	-.033	.012	.061	.012	-.051	.062	-.070
EX2 Trial 4	.069	.057	.075	.035	.060	.036	-.032	.064	-.043
EX2 Trial 5	-.015	.056	-.017	-.019	.059	-.019	.009	.063	.013

<sup>a</sup> Note: Performance-avoid orientation was only tested on conditions in which specific goals were present (N= 158). Tests on full sample (N= 310) were consistent with these results.

\* Significant at the  $p < .05$  level.

Table 4: Results of repeated measures analysis on performance and exploration over time by goal condition and world.

			Dependant Variables		
Result	Numerator df	Denominator df	Performance <sub>a</sub>	Exploration (EX1) <sub>a</sub>	Exploration (EX2) <sub>a</sub>
Time	4	1216	123.05**	426.5**	37.65**
World	2	304	30.57**	5.46**	25.41**
Goals	1	304	.07	4.17*	20.68**
Time by World	8	1216	3.45**	10.62**	17.04**
Time by Goals	4	1216	3.84**	4.21**	2.89*
World by Goals	2	304	9.02**	2.21	8.98**
Time by World by Goals	8	1216	3.01**	1.48	1.06

\*Significant at the  $p < .05$  level.

\*\*Significant at the  $p < .01$  level.

a – F values.

Figure 1 - 'Do your best' goals




Figure 2 - Specific, difficult goals




Figure 3 - Performance over time






Figure 4 – World 1




Figure 5 – World 2




Figure 6 – World 3




Figure 7 – Start of Trial (zoomed out)

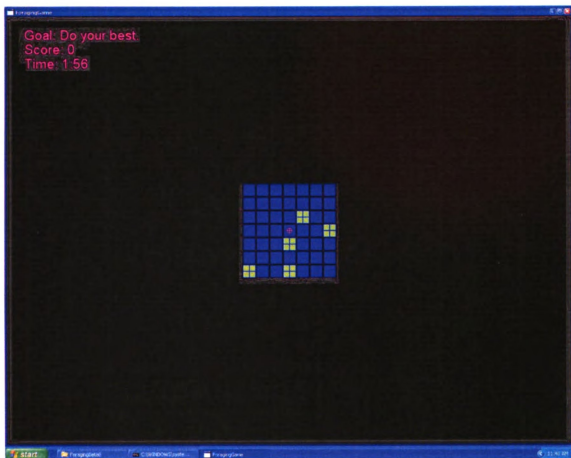
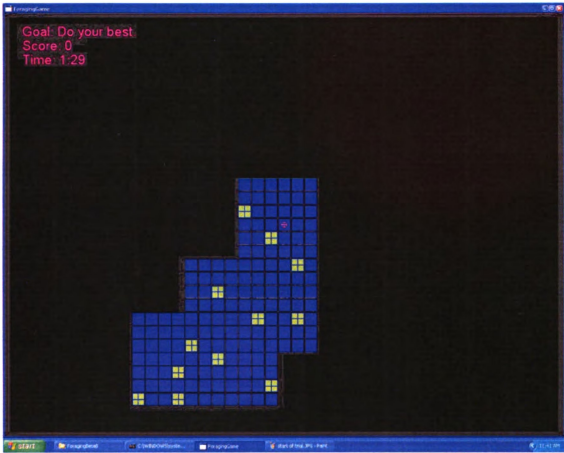


Figure 8 – After some exploring in trial (zoomed out).



## REFERENCES

- Berry, D.A., & Fristedt, B. (1985). *Bandit problems: Sequential allocation of experiments*. New York, NY: Chapman and Hall.
- Bird, R., Smith, E., & Bird, D. (1999). The hunting handicap: costly signaling in human foraging strategies. *Behavioral Ecology and Sociobiology*, 50, 9-19.
- Burns, B.D. & Vollmeyer, R. (2002). Goal specificity effects on hypothesis testing in problem solving. *The Quarterly Journal of Experimental Psychology*, 55, 241-261.
- Campbell, D. (1988). Task complexity: a review and analysis. *Academy of Management Review*, 13, 40-52.
- Carr, H. A. *Psychology: A study of mental activity*. New York: Longmans, Green, 1925.
- Cronicle, E.P., MacGregor, J.N., & Ormerod, T.C. (2004). What makes an insight problem? The roles of heuristics, goal conception, and solution recoding in knowledge-lean problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30, 14-71.
- Das, S. & Kamenica, E., (2005). Two-sided bandits and the dating market. *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, 19, 947-952.
- DeJong, K.A. & Spears, W.M., (1993). Using genetic algorithms to solve NP-complete problems. *Proceedings of the 3<sup>rd</sup> International Conference on Genetic Algorithms*, 124.
- DeShon, R.P. & Alexander, R.A. (1996). Goal setting effects on implicit and explicit learning of complex tasks. *Organizational behavior and human decision processes*, 65, 18-36.
- Earley, P.C., Connolly, T. & Ekegren, G. (1989). Goals, strategy development, and task performance: some limits on the efficacy of goal setting. *Journal of Applied Psychology*, 74, 24-33.
- Earley, P.C., Connolly, T., & Lee, C. (1989). Task strategy interventions in goal setting: the importance of search in strategy development. *Journal of Management*, 15, 589-602.
- Fowler, H., (1965). *Curiosity and exploratory behavior*. New York: Macmillan.
- Garcia, Rosanna, Calantone, Roger and Levine, Ralph (2003). The Role of Knowledge in

- Resource Allocation to Exploration versus Exploitation in Technologically Oriented Organization. *Decision Science*, 34, 323–349.
- Harlow, H.F. (1950). Learning and satiation of response in intrinsically motivated complex puzzle performance by monkeys. *Journal of Comparative Psychological Physiology*, 43, 289-294.
- Houghton Mifflin Company. (2006) *The American Heritage® Dictionary of the English Language, Fourth Edition*. Houghton Mifflin.
- Huber, V.L. (1985). Effects on task difficulty, goal setting, and strategy on performance of a heuristic task. *Journal of Applied Psychology*, 70, 492-504.
- Ivancevich, J.M., (1974). Changes in performance in a management by objectives program. *Administrative Science Quarterly*, 19, 563-574.
- James, W. (1890). *Principles of psychology*. New York: Holt.
- Jervis, R. (1988). *Perception and misperception in international politics*. Princeton, N.J.: Princeton University Press.
- Kaelbling, L.P., Littman, M.L., & Moore, A.W. (1996). Reinforcement learning: a survey. *Journal of Artificial Intelligence Research*, 4, 237-285.
- Kanfer, R. & Ackerman, P. (1989). Motivation and cognitive abilities: an integrative/aptitude-treatment interaction approach to skill acquisition. *Journal of Applied Psychology*, 74, 657-690.
- Kashdan, T.B., Rose, P., & Fincham, F.D. (2004). Curiosity and exploration: facilitating positive subjective experiences and personal growth opportunities. *Journal of Personality Assessment*, 82, 291-305.
- Klein, H.J, Wesson, M.J., Hollenbeck, J.R., Wright, P.M., & DeShon, R.P. (2001). The assessment of goal commitment: A measurement model meta-analysis. *Organizational Behavior and the Human Decision Processes*, 85, 32-55.
- Kozlowski, S.W.J & Bell, B.S. (2006). Disentangling achievement orientation and goal setting: Effects on self-regulatory processes. *Journal of Applied Psychology*, 91, 900-916
- Latham, G.P. & Kinne, S.B., (1974). Improving job performance through training in goal setting. *Journal of Applied Psychology*, 59, 187-191.
- Latham, G.P & Seijts, G.H. (1999). The effects of proximal and distal goals on performance on a moderately complex task. *Journal of Organizational Behavior*, 20, 421-429.

- Litman, J., Hutchins, T., & Russon, R., (2005). Epistemic curiosity, feeling-of-knowing, and exploratory behavior. *Cognition & Emotion, 19*, 559-582.
- Litman, J.A. & Spielberger, C.D. (2003). Measuring epistemic curiosity and its diverse and specific components. *Journal of Personality Assessment, 80*, 75-86.
- Locke, E.A. & Latham, G.P., (1990). *A theory of goal-setting and task performance*. Englewood Cliffs, NJ: Prentice-Hall.
- Locke, E.A. & Latham, G.P. (2002). Building a practically useful theory of goal setting and task motivation. *American Psychologist, 57*, 705-717.
- Locke, E.A. (1968). Toward a theory of task motivation and incentives. *Organizational Behavior and Human Performance, 3*, 157-189.
- Locke, E.A., Shaw, K.N., Saari, L.M., & Latham, G.P. (1981). Goal-setting and task performance: 1969-1980. *Psychological Bulletin, 90*, 125-152.
- Loewenstein, G. (1994). The psychology of curiosity: a review and reinterpretation. *Psychological Bulletin, 116*, 75-98.
- March, J.G. (1991). Exploration and exploitation in organizational learning. *Organizational Science, 2*, 71-87.
- Miller, C.S., Lehman, J.F., & Koedinger, K.R. (1999). Goals and learning in microworlds. *Cognitive Science, 23*, 305-336.
- Nooteboom, B. (1999). Discovery and organization: priorities in the theory of innovation. *Paper for the 1999 DRUID conference*.
- Pritchard, D., S.D.Jones, P.L.Roth, K.K.Stuebing, & S.E.Ekeberg (1988). Effects of group feedback, goal setting, and incentives on organizational productivity. *Journal of Applied Psychology, 73*, 337-358.
- Pyke, G.H. (1984). Optimal foraging: a critical review. *Annual Review of Ecology and Systematics, 15*, 523-575.
- Random House, Inc. (2006). *Random House Unabridged Dictionary*. Random House.
- Reio, T.G., Petrosko, J.M., Wiswell, A.K., & Thongsukmag, J. (2006). The measurement and conceptualization of curiosity. *The Journal of Genetic Psychology, 167*, 117-135.
- Rode, C., Cosmides, L, Hell, W. & Tooby, J. (2001). When and why do people avoid



unknown probabilities in decisions under uncertainty? Testing some predictions from optimal foraging theory. *Cognition*, 72, 269-304

Rothaermel, F.T. & Deeds, D.L. (2004). Exploration and exploitation alliances in biotechnology: a system of new product development. *Strategic Management Journal* 25, 201–221.

Seijts, G.H. & Latham, G.P. (2001). The effect of distal learning, outcome, and proximal goals on a moderately complex task. *Journal of Organizational Behavior*, 22, 291-307.

Seijts, G.H. & Latham, G.P. (2005). Learning versus performance goals: When should each be used? *Academy of Management Executive*, 19, 124-131.

Seijts, G.H., Latham, G.P., Tasa, K., & Latham, B.W. (2004). Goal setting and goal orientation: an integration of two different yet related literatures. *Academy of Management Journal*, 47, 227-239.

Sweller, J. & Levine, M. (1982). Effects of goal specificity on means-ends analysis and learning. *Journal of Experimental Psychology*, 88, 463-474.

VandeWalle, D. (1997). Development and validation of a work domain goal orientation instrument. *Educational and Psychological Measurement*, 57, 995–1015.

Vollmeyer, R., Burns, B., & Holyoak, K. (1996). The impact of goal specificity on strategy use and the acquisition of problem structure. *Cognitive Science*, 20, 75-100.

Wood, R. (1986). Task complexity: definition of the construct. *Organizational Behavior and Human Decision Processes*, 37, 60-82.

Wood, R., Mento, A., & Locke, E.A. (1987). Task complexity as a moderator of goal effects: a meta-analysis. *Journal of Applied Psychology*, 72, 416-425.

MICHIGAN STATE UNIVERSITY LIBRARIES



3 1293 03062 7990