

**PROMOTION, RELATIVE PERFORMANCE INFORMATION,
AND THE PETER PRINCIPLE**

by

Eric Wai Yeung Chan

B.S. Accounting and Finance,

University of Maryland, College Park, 2006

Submitted to the Graduate Faculty of
Joseph M. Katz Graduate School of Business in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy in Business Administration

University of Pittsburgh

2015

UNIVERSITY OF PITTSBURGH
JOSEPH M. KATZ GRADUATE SCHOOL OF BUSINESS

This dissertation was presented

by

Eric Wai Yeung Chan

It was defended on

June 9, 2015

and approved by

Jongwoon (Willie) Choi
Assistant Professor of Business Administration
University of Pittsburgh

Shane S. Dikolli
Associate Dean, Faculty Engagement and Associate Professor of the Practice
Duke University

John H. Evans III
Katz Alumni Professor of Accounting and Area Director for Accounting
University of Pittsburgh

Adam Presslee
Assistant Professor of Business Administration
University of Pittsburgh

Dhinu Srinivasan
Associate Professor of Business Administration
University of Pittsburgh

Dissertation Advisor: Donald V. Moser
Professor of Business Administration and Dean's Excellence Faculty Fellow
University of Pittsburgh

Copyright © by Eric Wai Yeung Chan

2015

ABSTRACT

PROMOTION, RELATIVE PERFORMANCE INFORMATION, AND THE PETER PRINCIPLE

Eric Wai Yeung Chan, PhD

University of Pittsburgh, 2015

My dissertation consists of two main parts. In the first part, I present my experimental study in which I examine the effects of providing workers with relative performance information (RPI) on employers' promotion decisions and the impact of those decisions on worker performance. In my experimental setting, the job after promotion requires higher-level abilities than the current job. I find that workers increase their effort to improve their current job performance after a promotion opportunity is announced because they expect this to increase their chances of promotion even though the new task requires higher-level abilities. Moreover, because employers anticipate that workers who have RPI will react negatively if they see that the best current job performer is not promoted, employers promote the best current job performer rather than the worker best suited for the next job more often when workers have RPI than when they do not. Finally, consistent with the Peter Principle, I find that the promoted worker's performance after promotion is lower when workers have RPI than when they do not because the

promoted worker lacks the ability to perform the new job well. My results suggest that providing workers with RPI can have a downside that offsets the benefits documented in previous studies.

In the second part of my dissertation, I discuss how three common firm practices could potentially help firms mitigate the sorting problem I identified in my experimental study. Specifically, I discuss how the communication of the promotion criteria, group incentive plans, and subjective performance evaluation could help employers improve the sorting of workers through promotions. However, while each approach has advantages, each also has drawbacks. It appears that there is no promotion-based incentive system that provides optimal effort incentives to workers while simultaneously achieving the best match between workers and their jobs. This conflict between the incentive and sorting roles of promotion is difficult to resolve and firms must ultimately trade-off between them.

TABLE OF CONTENTS

PREFACE.....	xi
1.0 INTRODUCTION.....	1
2.0 BACKGROUND	7
2.1 OVERVIEW OF CHAPTER.....	7
2.2 PROMOTIONS.....	7
2.3 THE PETER PRINCIPLE.....	9
2.4 RELATIVE PERFORMANCE INFORMATION	11
3.0 DEVELOPMENT OF HYPOTHESES AND RESEARCH QUESTION	13
3.1 OVERVIEW OF CHAPTER.....	13
3.2 KEY ASSUMPTIONS OF SETTING.....	13
3.3 WORKER PERFORMANCE AFTER PROMOTION ANNOUNCEMENT BUT BEFORE PROMOTION	15
3.4 THE EFFECT OF RPI ON EMPLOYERS' PROMOTION DECISIONS... 	16
3.5 THE EFFECT OF RPI ON PROMOTED WORKERS' PERFORMANCE AFTER PROMOTION	18
3.6 THE EFFECT OF RPI ON NON-PROMOTED WORKERS' PERFORMANCE AFTER PROMOTION.....	18
3.7 THE EFFECT OF RPI ON EMPLOYER PROFIT.....	19
4.0 RESEARCH DESIGN.....	21

4.1	OVERVIEW OF CHAPTER.....	21
4.2	OVERVIEW OF EXPERIMENT	21
4.3	DETAILED PROCEDURES	27
4.4	PAYOFFS	30
	4.4.1 Worker Payoffs	30
	4.4.2 Employer Payoffs	31
4.5	PARTICIPANTS.....	32
5.0	EXPERIMENTAL RESULTS.....	33
5.1	OVERVIEW OF CHAPTER.....	33
5.2	VALIDATION OF KEY ASSUMPTIONS	33
5.3	OVERVIEW OF TESTS OF HYPOTHESES	34
5.4	TESTS OF HYPOTHESIS 1.....	35
5.5	TESTS OF HYPOTHESIS 2.....	38
5.6	TESTS OF HYPOTHESIS 3.....	43
5.7	TESTS OF HYPOTHESIS 4A AND 4B	45
5.8	TESTS OF RESEARCH QUESTION 1	49
5.9	SUMMARY OF RESULTS	52
6.0	DISCUSSION OF WAYS TO MITIGATE THE SORTING PROBLEM.....	54
6.1	OVERVIEW OF CHAPTER.....	54
6.2	COMMUNICATION OF THE PROMOTION CRITERIA	55
	6.2.1 Overview	55
	6.2.2 Advantages.....	55
	6.2.3 Disadvantages.....	59

6.3	GROUP INCENTIVE PLANS	62
	6.3.1 Overview	62
	6.3.2 Advantages.....	63
	6.3.3 Disadvantages.....	65
6.4	SUBJECTIVE PERFORMANCE EVALUATION.....	67
	6.4.1 Overview	67
	6.4.2 Advantages.....	68
	6.4.3 Disadvantages.....	72
7.0	DISCUSSION AND CONCLUSION	75
	7.1 OVERVIEW OF CHAPTER.....	75
	7.2 DISCUSSION OF FINDINGS.....	75
	7.3 CONTRIBUTIONS.....	76
	7.4 LIMITATIONS AND FUTURE RESEARCH.....	78
	BIBLIOGRAPHY	82

LIST OF TABLES

Table 1: Workers’ performance after promotion announcement (Tests of H1).....	36
Table 2: Promoted workers’ post-promotion performance (Tests of H3).....	44
Panel A: No RPI vs. RPI.....	44
Panel B: RPI condition – Best current job performer promoted vs. Best candidate for the next job promoted.....	44
Table 3: Non-promoted workers’ post-promotion performance (Tests of H4a).....	46

LIST OF FIGURES

Figure 1: Screenshot examples of slider tasks	23
Panel A: Basic slider task	23
Panel B: Advanced slider task	23
Figure 2: Experimental timeline	25
Panel A: Timeline of key events	25
Panel B: Timeline of <i>RPI</i> condition	26
Panel C: Timeline of <i>No RPI</i> condition	26
Figure 3: Effect of RPI on employers' promotion decision (Tests of H2)	39
Figure 4: Non-promoted workers' response to promotion decision (Tests of H2)	42
Panel A: Non-promoted workers' deservingness judgments	42
Panel B: Non-promoted workers' promotion reactions	42
Figure 5: Interaction effect of RPI and employer's promotion decision on non-promoted workers' post-promotion performance (Tests of H4b)	48
Figure 6: Effect of RPI on employers' profit (Tests of RQ1)	51

PREFACE

I thank my dissertation committee members Willie Choi, Harry Evans, Shane Dikolli, Adam Presslee, and Dhinu Srinivasan for their feedback and advice. I especially thank my dissertation chair Don Moser for his guidance and support over the past five years. I could not have asked for a better mentor and role model. I thank Harry Evans for the unique opportunity to work as an editorial research assistant for *The Accounting Review*. I thank other professors at the University of Pittsburgh, especially Elise Boyas, Mei Feng, Vicky Hoffman, and Nandu Nagarajan for helping me develop as a researcher and teacher. I thank my fellow PhD students, especially Jordan Bable, Yangyang Fan, Michele Frank, Duanping Hong, Patrick Martin, Bryant Stikeleather, and Yinuo Tang for their encouragement and friendship. I thank Chris Fedor, Chris Gursky, and Carrie Woods for their administrative support. I thank the Institute of Management Accountants Research Foundation Doctoral Student Grant, the Fryrear Doctoral Fellowship, and the Mervis Doctoral Fellowship for their financial support. Finally, I thank my wife, Tina, my family, friends, and church for keeping me grounded and reminding me to always focus on the most important things in life.

1.0 INTRODUCTION

Promotions are integral to most organizations and serve two important and distinct roles (Milgrom and Roberts 1992, p.364). The first is to motivate and reward workers (incentive role) and the second is to sort workers into jobs that best match their abilities (sorting role). However, the incentive and sorting roles may conflict when the nature of the tasks changes between hierarchical job levels such that workers' performance in the current job is not fully informative about their performance in the next job. In such cases, the best performer at one level of the hierarchy is often not the best candidate for the next higher-level job.

The situation described above occurs quite often, such as when salespersons are promoted to sales managers, engineers are promoted to project managers, and teachers are promoted to school administrators. In such cases, employers must decide whether to promote the worker with the best current job performance versus the worker with the highest ability to perform the next job. Such promotion decisions are important because they have significant long-term effects on workers' wages and firm productivity.

Prior empirical evidence suggests that employers often promote workers to job positions for which they are not well-suited – a phenomenon known as the Peter Principle (Peter and Hull 1969). Although prior analytical studies offer various economic-based explanations for why it can be optimal for employers to intentionally promote workers to jobs that do not best match their abilities, these studies generally lack supporting empirical evidence (e.g., Waldman 1984; Bernhardt 1995; Fairburn and Malcomson 2001). Further, none of the prior studies explain why

the Peter Principle is observed in some firms (Lazear 1992; Acosta 2010) but not in others (Grabner and Moers 2013).

In this study, I posit that whether workers are provided with relative performance information (hereafter, RPI) can influence employers' promotion decisions and lead to the Peter Principle. Employers often provide their workers with RPI to motivate improved performance (e.g., Nordstrom et al. 1990; Gino and Staats 2011). Prior studies find that RPI generally has a positive effect on worker performance in their current job (e.g., Blanes i Vidal and Nossol 2011; Tafkov 2013). However, an overlooked consequence of providing workers with RPI is that this could influence workers' expectations about who should be promoted and in turn influence employers' promotion decisions. The purpose of my experimental study is to examine these potential effects of RPI on employers' promotion decisions and the impact of those decisions on worker performance.

My dissertation consists of two parts. Chapters 2 to 5 cover the first part of my dissertation in which I present the background, hypotheses, design, and findings of my experimental study. Chapter 6 covers the second part of my dissertation in which I discuss how three specific firm practices could potentially help firms mitigate the sorting problem I identified in my experimental study.

Chapter 2 presents background information for my study. I first discuss the importance of promotions in organizations and the potential conflict between the incentive and sorting roles of promotions. I next discuss the Peter Principle and review prior research that examines possible causes for this phenomenon. Finally, I discuss prior research on the effects of RPI on worker performance.

Chapter 3 develops the hypotheses and research questions. My main prediction is that whether workers are provided with RPI can influence the employer's promotion decisions. Prior research finds that individuals often judge fairness based on the consistency between an action and the related outcome. That is, a positive (negative) outcome that follows from a positive (negative) action is perceived to be deserved, while a negative (positive) outcome that follows from a positive (negative) action is undeserved. Similarly, I expect workers to evaluate the fairness of the employer's promotion decision based on whether they believe the promoted worker deserved the promotion.

When workers are provided with RPI, their relative current job performance is very salient. Therefore, I predict that workers will use their relative current job performance to evaluate the fairness of their employer's promotion decision. If workers who are not promoted see that the promoted worker was *not* the best current job performer, they are likely to perceive the employer's promotion decision as unfair and react negatively. I predict that some employers will anticipate this and promote the best current job performer to avoid such negative worker reactions even when the promoted worker's abilities do not best match those required for the new job.

In contrast, when workers are not provided with RPI, they have no clear basis on which to evaluate the fairness of the employer's promotion decision. Because workers cannot identify the best current job performer, they are less likely to react negatively regardless of which worker the employer chooses to promote. I predict that employers will understand this and therefore promote the worker with the highest ability to perform the next job even when that worker is not the best current job performer. In summary, I expect more behavior consistent with the Peter Principle when workers have RPI than when they do not.

Chapter 4 describes the experimental design and procedures. I conduct an experiment using a setting in which participants are assigned the role of either an employer or worker and interact over multiple work periods. Workers perform tasks that generate revenue for their employer, who in turn pays workers a fixed salary plus an output-based bonus. There are two levels of workers: lower-level and higher-level. Higher-level workers perform a more complex task and earn higher pay than lower-level workers.

I manipulate whether workers receive RPI using a 1 x 2 between-participants design with two conditions: *RPI* vs. *No RPI*. In the *RPI* condition, lower-level workers receive information about their own current job performance rank and the ranks of other workers in the same firm. In the *No RPI* condition, lower-level workers receive no performance rank information.

During the experiment, workers receive an announcement about an upcoming promotion, and the employer subsequently promotes one lower-level worker to a new higher-level position. I examine workers' behavior before and after the promotion announcement and after the promotion itself as well as employers' promotion decisions and profit across the two conditions.

Chapter 5 reports the experimental results. I find that after the announcement of an upcoming promotion, workers in both the *RPI* and *No RPI* conditions increase their effort to improve their current job performance because they expect that this will increase their chances of promotion. Moreover, because employers anticipate that workers who have RPI will react negatively if they see that the best current job performer is not promoted, employers promote the best current job performer rather than the worker best suited for the next job more often when workers have RPI than when they do not. As a result, the promoted workers' performance is lower when workers have RPI because of the mismatch between the promoted workers' abilities and the higher-level task. After being passed over for promotion, non-promoted workers'

performance decreases, and this decrease is greater for workers who have RPI and see that the best current job performer was not promoted because they believe this is unfair. Finally, consistent with the Peter Principle, I find that the employer's profit after promotion is lower when workers have RPI because the promoted worker lacks the ability to perform the new job well.

Chapter 6 discusses three possible ways that firms could mitigate the sorting problem identified in my experimental study and their drawbacks. I first discuss how the communication of the promotion criteria could improve the sorting of workers and increase the transparency of the promotion process by adjusting workers' expectations about the employer's promotion role. However, such communication could reduce the perceived strength of the link between workers' current job performance and the reward of promotion, which could lower workers' motivation to perform their current jobs well. In addition, workers may redirect their time and effort away from their current jobs towards developing their abilities for the next job, resulting in lower job performance. Finally, workers may perceive the promotion process to be unfair if they believe that they have limited control over their abilities to perform the next job.

I next discuss how the use of group incentive plans such as profit-sharing plans could improve the sorting of workers by aligning the interest of the workers with those of the employer. Prior research also finds that group incentive plans encourage mutual monitoring among workers and could lead to higher overall productivity. However, group incentive plans might not improve sorting effectively if workers value their personal incentives over the group's incentives. Further, such plans impose greater risk on workers and could cause the best current job performers to shirk if they do not believe they will be properly rewarded by the employers.

Finally, I discuss how employers could exploit their information advantage and use subjective performance evaluation to improve the sorting of workers by strategically assigning workers' performance ratings. By creating a contrast in ratings between the worker whom the employer wants to promote and all other workers, the employer can justifiably promote the worker best suited for the next job rather than the best current job performer. However, the effectiveness of subjective performance evaluation depends on the information environment and the level of trust between the employer and workers. Further, subjective evaluations could decrease workers' motivation if workers perceive their performance ratings to be inaccurate and negatively biased.

Chapter 7 concludes my dissertation. I discuss the contributions and limitations of my study and potential future research. The findings of my study provide new insights into how employers make promotion decisions in a setting in which the incentive and sorting roles of promotions are in conflict. My study is the first to provide evidence that providing workers with RPI can lead employers to promote workers to job positions for which they are not well-suited as suggested by the Peter Principle. Because of this mismatch between the promoted workers' abilities and their new jobs, the promoted workers' performance after promotion is lower when workers have RPI. As such, I identify an unanticipated cost of providing workers with RPI that can offset the benefits documented in previous studies. My results can help firms better weigh the costs and benefits of providing workers with RPI and thus to design more effective information policies that maximize firm productivity.

2.0 BACKGROUND

2.1 OVERVIEW OF CHAPTER

This chapter provides background information and motivation for my study. Section 2.2 discusses the importance of promotions in organizations and the potential conflict between the incentive and sorting roles of promotions. Section 2.3 introduces the Peter Principle and presents prior analytical research that examines possible causes for the phenomenon. Finally, Section 2.4 discusses the use of relative performance information in practice and prior research that examines its effect on worker performance.

2.2 PROMOTIONS

Promotions are arguably the most important form of incentives in many organizations (Gibbs 1996). In contrast to explicit incentives such as bonuses, promotions are implicit because employers do not typically pre-commit to a specific promotion rule and therefore have *ex post* discretion in determining which workers to promote (Prendergast 1999). Such implicit incentives are particularly strong for those at lower levels of the organizational hierarchy (Gibbons and Murphy 1992; Ederhof 2011). Further, while many explicit incentives such as bonuses are

recurring in nature, most firms have a fixed number of positions to which workers can be promoted, making promotions particularly valuable to workers (DeVaro 2006).

Promotions are often associated with large wage increases. For example, Gibbs (1995) found that a promotion is associated with up to 55% larger total lifetime earnings for workers at a large U.S. firm. Kale et al. (2009) found that the average pay gap between the CEO and the next highest paid executives at companies covered by ExecuComp was \$2.77 million. In addition to substantial wage growth, promotions are often accompanied by higher status, authority, job influence, and work autonomy.

It is well-established in the economics literature that promotions serve two important yet distinct roles: 1) to motivate workers in their current jobs (incentive role), and 2) to sort workers into jobs that best match their abilities (sorting role) (Milgrom and Roberts 1992; Gibbs 1995).¹ However, prior research on promotion focuses primarily on the incentive role of promotions but largely ignores the sorting role (e.g., Gibbs 1995; Campbell 2008; Cichello et al. 2009; Ederhof 2011). For example, some prior analytical studies model promotion as an incentive that is equivalent to, and a substitute of, workers' explicit incentives such as bonuses (Gibbons and Murphy 1992). Other promotion studies that rely on tournament theory also assumes that promotion is used for incentive purposes and that the promotion 'prize' is always awarded to the best current job performer (e.g. Lazear and Rosen 1979, Nalebuff and Stiglitz 1983).

However, in settings in which the nature of the job differs between hierarchical levels such that the best performer at one level of the hierarchy might not be the best candidate for the job at the next higher level, the sorting role of promotions becomes important (Grabner and

¹ In this study, I use the term "ability" to represent an individual trait that directly affects how effectively an individual performs a certain task (Bonner and Sprinkle 2002). Unlike skill which can be improved through training and learning over time, I assume that ability is a relatively fixed trait. Because workers' abilities cannot be easily improved, it is important for firms to match their abilities to their jobs to maximize worker productivity.

Moers 2013). Promotion-based incentive systems that always promote the best current job performer will lead to ineffective sorting because the promoted workers may not be well-suited for their new jobs, resulting in the oft-cited Peter Principle effect described more fully below (Baker et al. 1988).

2.3 THE PETER PRINCIPLE

The Peter Principle refers to the observation that workers in a hierarchy are often promoted to positions for which they are incompetent and therefore cannot attain further promotions (Peter and Hull 1969). Support for the Peter Principle comes from field evidence that worker performance appears to decline after promotion.² For example, Acosta (2010) found that after several promotions, workers at a large U.S. firm were less likely to be promoted than new hires. Medoff and Abraham (1980) reported that workers' performance scores declined after promotion. Finally, Lazear (1992), Baker et al. (1994a, 1994b), and Gibbs and Hendricks (2004) found corroborative evidence across several firms that wage growth, raises, and bonuses fell with tenure, which suggest that workers become less competent at their jobs as they get promoted.

Although there is empirical support for the Peter Principle, there is also debate regarding the underlying cause. Some economic models suggest that the Peter Principle is the result of promoted workers' performance regressing to the mean because they were promoted based on their transitory ability or "luck" (Lazear 2004). Other models show that it is optimal in some settings for the employer to intentionally promote less productive workers. For example,

² Although the field studies described above provide evidence that is consistent with the Peter Principle, these studies do not provide conclusive evidence that the observed decline in workers' performance after promotion is the direct result of the employer's decision to promote workers who are not best suited for their new jobs.

Bernhardt (1995) analyzes a two-period, two-job model in which the employer has private information about workers' abilities, and competing firms use workers' job positions and education as signals of their abilities. Bernhardt shows that the equilibrium outcome is one in which the employer strategically promotes more educated but less productive workers to exploit the employer's information advantage about workers' abilities. Prendergast (1992) argues that the Peter Principle could occur when risk-averse workers seek insurance over the uncertainty about their own abilities. Specifically, Prendergast suggests that when workers do not know their *ex ante* abilities and the employer can infer each worker's ability after one period, risk-averse workers would prefer the employer not to differentiate based on their abilities. In this case, the second best contract is one in which the employer promotes too many workers, resulting in the Peter Principle.

Although these prior analytical studies provide different theoretical explanations for the Peter Principle, they lack empirical support.³ Further, the plausibility of these explanations is questionable. For example, the notion that employers would intentionally not promote less educated but more productive workers to keep their wages low seems unlikely to occur in practice. It also appears likely that workers in actual business settings would prefer their employers to reward high ability workers over low ability workers. In this study, I offer a new explanation and posit that whether workers are provided with RPI could cause the Peter Principle to more likely occur.

³ Currently, there is no empirical evidence showing that employers are intentionally promoting workers who are not suited for the next job because they do not want outside firms to bid up workers' wages as suggested by Bernhardt (1995) or because they believe risk-averse workers want insurance over the uncertainty about their abilities as suggested by Prendergast (1992).

2.4 RELATIVE PERFORMANCE INFORMATION

Employers often provide their workers with RPI to motivate improved performance (e.g., (Anderson et al. 1982; Wikoff et al. 1982; Nordstrom et al. 1990). For example, the women's apparel retailer AnnTaylor uses a daily feedback system that provides its sales staff with relative rank information based various performance metrics such as average sales per hour, units sold, and dollars per transaction (O'Connell 2008). Similarly, retail bank branch managers often provide regular feedback to their tellers on the number of new accounts opened by each teller to motivate them to increase their sales (Gino and Staats 2011).

Economic theory focuses mainly on the effects of RPI when worker pay is tied directly to relative performance, which is known as relative performance evaluation (RPE). Agency theory shows that the when workers face a common "shock" and their outcomes are related, RPE allows the employer to better isolate workers' actions and differentiate workers' performance (Holmstrom 1982). Tournament theory argues that compensating workers based on their ordinal rank could induce worker effort as effectively as piece-rate contracts and could be less costly to implement (Lazear and Rosen 1981; Nalebuff and Stiglitz 1983). Prior field studies provide supporting evidence that RPE is often used to determine CEO pay (Antle and Smith 1986; Gibbons and Murphy 1990).

In addition to research on RPE, there is a growing behavioral literature on the effects of RPI when worker pay is *not* tied to relative performance. Based on social comparison theory, prior studies argue that the presence of RPI motivates individuals to increase their effort because

they want to outperform their peers and avoid negative self-perceptions and shame (Festinger 1964; Beach and Tesser 1995; Suls and Wheeler 2000). Consistent with this argument, prior field and experimental studies find that RPI has a positive effect on workers' performance even when it does not affect their pay (Blanes i Vidal and Nossol 2011; Tafkov 2013). In this study, I argue that a potential downside to providing workers with RPI is that it could lead to the ineffective sorting of workers.

3.0 DEVELOPMENT OF HYPOTHESES AND RESEARCH QUESTION

3.1 OVERVIEW OF CHAPTER

This chapter develops my hypotheses and research questions. Section 3.2 establishes the key assumptions of my experimental setting and the rationale for each assumption. Section 3.3 considers workers' behavior after the announcement of an upcoming promotion. Section 3.4 considers employers' promotion decisions across the two experimental conditions (*RPI* and *No RPI*). Section 3.5 and Section 3.6 consider the post-promotion behaviors of the promoted workers and non-promoted workers, respectively. Finally, Section 3.7 considers the effect of RPI on the employers' profit in my setting.

3.2 KEY ASSUMPTIONS OF SETTING

Before developing my hypotheses, I describe some key assumptions regarding my setting:

- 1) The two jobs differ between hierarchical levels such that the ability necessary to perform the current job is different from the ability necessary to perform the next higher-level job.
- 2) The employer has private information about workers' relative current job performance (i.e. RPI), which can be provided or not provided to workers.

- 3) The employer has private information about workers' relative abilities to perform the next higher-level job, which is not provided to workers.
- 4) The employer's private information about workers' relative abilities to perform the next higher-level job is objective.

Assumption (1) separates workers' current job performance from their abilities to perform the next job, which typically overlap to varying degrees in the real-world. This assumption provides a setting that captures the conflict between the incentive and sorting roles of promotions.⁴ Assumptions (2) and (3) reflect the fact that employers typically have the expertise, resources, and experience to more accurately and objectively assess workers' performance and abilities than can workers themselves (Gibbs 1995; Ray 2007). In addition, as explained in Chapter 6, employers may not want to provide information about workers' relative abilities to perform the next job because doing so could reduce workers' incentives to perform their current jobs well. Finally, by making the assessment of workers' abilities to perform the next job objective in my setting under Assumption (4), I rule out the potential alternative explanation that employers promote the best current job performer because they perceive the assessment of workers' abilities to perform the next job to be too subjective. In other words, because the employers' information about workers' current job performance and their abilities to perform the next job are objective, employers should use both sets of information when making their promotion decisions (Rajan and Reichelstein 2009; Grabner and Moers 2013).

⁴ As the jobs between hierarchical levels become more similar, the conflict between the incentive and sorting roles of promotion is reduced. If the measures of performance for both the lower-level and higher-level tasks are perfectly correlated, there is no longer a conflict between the incentive and sorting roles of promotions because the employer can simply promote the best current job performer whose ability will also best match the next job.

3.3 WORKER PERFORMANCE AFTER PROMOTION ANNOUNCEMENT BUT BEFORE PROMOTION

When an employer announces an upcoming promotion, workers who are eligible for the promotion may respond by providing higher effort. In a field study, Campbell (2008) finds that managers at a fast-food retailer responded to promotion incentives by working harder and improving their scores on their performance measures. However, an important difference between my setting and Campbell's is that in my setting, workers' current job performance is not fully informative about their performance in the next higher-level job. Because workers in my setting are aware that the current job and the next job require different abilities, they may or may not believe that increasing their current job performance will increase their chances of promotion.

Because employers do not typically pre-commit to a specific promotion rule, different workers will likely have different expectations about their employer's promotion rule. If workers expect the employer to make promotion decisions based on their current job performance, they would have strong incentive to increase their effort and current job performance because this increases their chances of being promoted. If workers expect their employer to use an alternative promotion rule such as promoting based on workers' ability to perform the next job, then they have less incentive to increase their current job performance. However, in either case, there is no incentive for workers to intentionally decrease their performance after the promotion announcement. Thus, my first hypothesis is:

H1: Workers' performance *increases* after the employer announces an upcoming promotion.

3.4 THE EFFECT OF RPI ON EMPLOYERS' PROMOTION DECISIONS

When employers make a promotion decision, they likely consider how that decision will affect their workers' subsequent behaviors, particularly workers who were passed over for promotion. I posit that non-promoted workers' reactions to their employer's promotion decision will depend on the perceived fairness of that decision. Prior research on deservingness finds that people make fairness judgments based on the evaluative consistency between an action and the related outcome (Feather 1999a, 1999b). Using a scenario-based experiment, Feather (2008) finds that participants judged a top performer who consistently exceeded performance standards to be *deserving* of a promotion, while a colleague who only had satisfactory performance was judged to be *undeserving* of a promotion. Further, participants were more likely to have negative feelings of disappointment and resentment towards a committee that chose to promote an undeserving job candidate. In a similar manner, I predict that how non-promoted workers judge the deservingness of the promoted worker will affect how they react to their employer's promotion decision.

When workers are provided with RPI, they are able to compare their relative current job performance because it is made very salient. Therefore, I predict that workers will evaluate the fairness of the employer's promotion decision based on the current job performance of the promoted worker. If non-promoted workers see that the promoted worker was *not* the best

current job performer, they will likely perceive the employer's decision to be unfair and react negatively.

Prior literature provides evidence that workers' perceived fairness regarding their employer's actions can affect their organizational commitment (Moorman et al. 1993), organizational citizenship behavior (Moorman 1991), reporting honesty (Zhang 2008), and job performance (Konovsky and Cropanzano 1991). Thus, non-promoted workers who have RPI and see that the best current job performer was not promoted will likely react negatively and thus may lower their effort in future periods. I predict that the employer will anticipate this and promote the best current job performer rather than the best candidate for the next job to avoid such possible negative worker reactions.

In contrast, when workers are not provided with RPI, non-promoted workers have no clear basis on which to evaluate the fairness of the employer's promotion decision. Because workers cannot identify the best current job performer, they are less likely to react negatively, regardless of which worker the employer promotes.⁵ From an economic perspective, it is efficient for the employer to sort workers into jobs that best match their abilities. Thus, I predict that when workers do not have RPI, the employer will promote the worker whose ability best match the next job even if that worker is not the best current job performer. Thus, my second hypothesis is:

H2: Employers are more likely to promote the worker with the best current job performance than the worker with the highest ability to perform the next job when workers have RPI than when workers do not have RPI.

⁵ Because prior studies find that individuals are often overconfident and overoptimistic about their own abilities (e.g. Kruger 1999; Kruger and Dunning 1999), workers with no RPI may still believe they deserved to be promoted and react negatively to being passed over for promotion. However, the non-promoted workers' reactions in this case would likely be less negative than when they have RPI and see that the best current job performer was not promoted.

3.5 THE EFFECT OF RPI ON PROMOTED WORKERS' PERFORMANCE AFTER PROMOTION

The promoted workers' post-promotion performance depends primarily on their ability to perform the new job. Because H2 predicts that the employer is more likely to promote the best current job performer when workers have RPI, the promoted worker will likely be mismatched to the new job. In contrast, the employer is more likely to promote the worker with the highest ability to perform the new job when workers do not have RPI. Thus, the promoted workers' post-promotion performance will be lower when workers have RPI than when they do not because of the mismatch between their ability and the requirements of the new job when they have RPI. Thus, my third hypothesis is:

H3: Promoted workers' performance after promotion is *lower* when workers have RPI than when workers do not have RPI.

3.6 THE EFFECT OF RPI ON NON-PROMOTED WORKERS' PERFORMANCE AFTER PROMOTION

H1 predicts that workers' performance will increase following a promotion announcement because the likelihood of a future promotion increases. Similarly, I predict that non-promoted workers' performance will decrease after a promotion occurs because the likelihood of a future promotion declines.

In addition, H2 predicts that workers will perceive the employer's promotion decision to be unfair if they have RPI and see that the best current job performer was *not* promoted. In this case, I predict that non-promoted workers will reduce their effort as a negative response to the employer's perceived unfair promotion decision. Thus, I predict that the post-promotion performance of such workers will decline more than that of non-promoted workers who perceive the employer's promotion to be relatively fairer either because they have RPI and see that the best current job performer was promoted or because they do not have RPI to evaluate the fairness of the employer's decision. Thus, my fourth hypothesis has two parts:

H4a: Non-promoted workers' performance declines after a promotion occurs.

H4b: Non-promoted workers' performance declines more when workers have RPI and the employer does *not* promote the best current job performer than when either workers have RPI and the employer promotes the best current job performer or workers do not have RPI.

3.7 THE EFFECT OF RPI ON EMPLOYER PROFIT

Finally, I examine the overall effect of RPI on the employer's profit, which depends on workers' performance before and after the promotion occurs. Prior field and experimental studies find that RPI has a positive effect on worker performance even when it is not directly tied to their compensation (Blanes i Vidal and Nossol 2011; Tafkov 2013). Based on social comparison theory, these studies find that when RPI is present, individuals increase their effort and performance to avoid negative self-perceptions and shame (Festinger 1954; Beach and Tesser 1995; Suls and Wheeler 2000). Based on the results of these prior studies, workers' performance *before* promotion is expected to be higher when workers have RPI because of the motivational

effects of RPI. However, H3 predicts that the promoted workers' performance *after* promotion will be lower when workers have RPI because of the potential mismatch between the promoted workers' abilities and their new job. Because the net effect of RPI on the employer's profit of these offsetting performance predictions is unclear, my research question is:

RQ1: What is the effect of RPI on the employer's profit?

4.0 RESEARCH DESIGN

4.1 OVERVIEW OF CHAPTER

This chapter presents the research design and procedures. Section 4.2 provides an overview of the experimental design, including the experimental tasks and timeline. Section 4.3 presents the detailed procedures of the experiment. Section 4.4 presents the workers' and employers' payoff calculations. Finally, Section 4.5 provides demographic information about the participants.

4.2 OVERVIEW OF EXPERIMENT

I conduct an experiment in which one employer and four workers form a single firm. I use a 1 x 2 between-subjects design with two conditions: *RPI* vs. *No RPI*. In the *RPI* condition, workers receive information about their own current job performance rank and the ranks of other workers in their firm at the end of each work period. In the *No RPI* condition, workers receive no performance rank information.

In the experiment, workers perform tasks that generate revenue for their employer who compensates workers for their output. As explained later, workers' compensation consists of a

fixed salary and an output-based bonus per period. The employer is the residual claimant who receives the total revenue generated by the workers less the compensation paid to the workers.⁶

There are two levels of workers: lower-level and higher-level. Higher-level workers earn higher compensation than lower-level workers. The firm produces two products: Product H and Product L. Product H is produced by higher-level workers and generates more revenue per unit than Product L, which is produced by lower-level workers. As explained later, the employer assigns workers to either the lower-level or higher-level position during the experiment.

To produce Product L, lower-level workers perform a computerized ‘basic slider task’ adapted from Gill and Prowse (2011). Lower-level workers see a computer screen that displays a large number of sliders, with each slider initially positioned at 0 along a scrollbar from 0 to 100. The task is to position each slider at exactly 50 using the computer mouse. To produce Product H, higher-level workers perform an ‘advanced slider task’. Unlike the basic slider task in which lower-level workers always adjust the slider to 50, higher-level workers must first solve a math problem in their heads and then adjust the slider to the number position that matches the answer to the problem. Panels A and B of Figure 1 show screenshots of a single slider at its initial position (a) and correct position (b) for the basic and advanced slider tasks, respectively.


⁶ In practice, the superior who makes the promotion decisions may not be the residual claimant. However, it is still often in the superior’s best interest to promote workers in a way that maximizes firm productivity and profit because the success of the firm will likely affect the superior’s own compensation (Fairburn and Malcolmson 2001).

Figure 1: Screenshot examples of slider tasks

Panel A: Basic slider task



Panel B: Advanced slider task



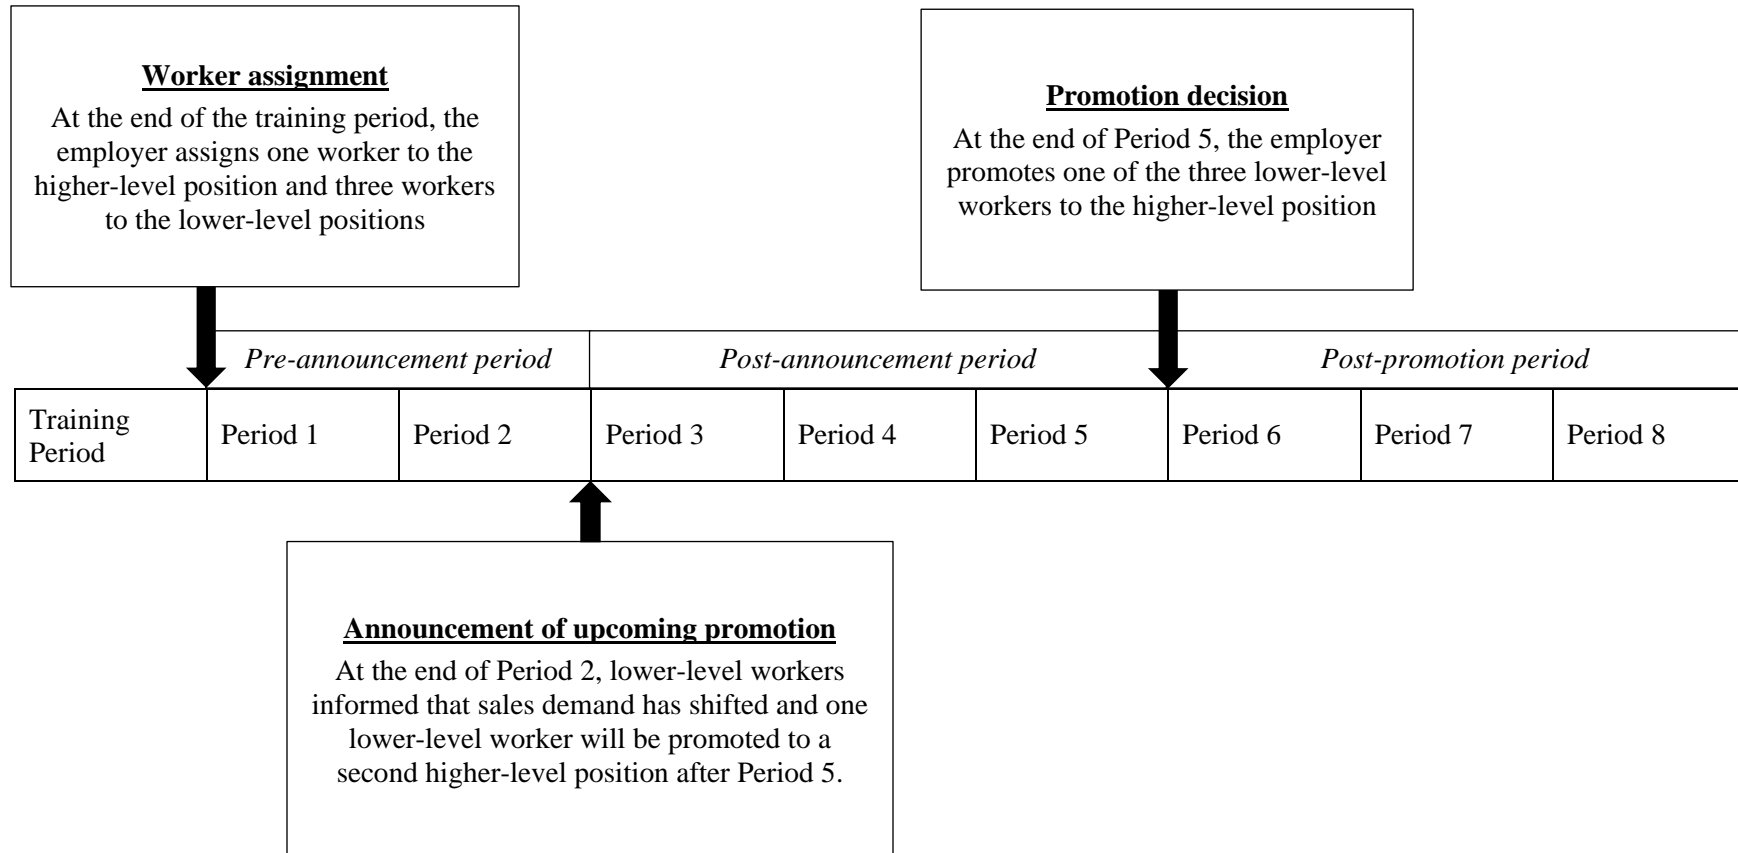
These tasks are designed to require different abilities (see Assumption 1). The basic task requires only mechanical abilities, while the advanced task requires both mechanical and mental arithmetic abilities. Workers' performance in each task is measured by the amount of output they produce, and output is measured by the total number of sliders positioned correctly.

To test my hypotheses, three key events occur during the experiment in both the *RPI* and *No RPI* conditions: 1) worker assignment after an initial training period, 2) promotion announcement after Period 2, and 3) promotion after Period 5.

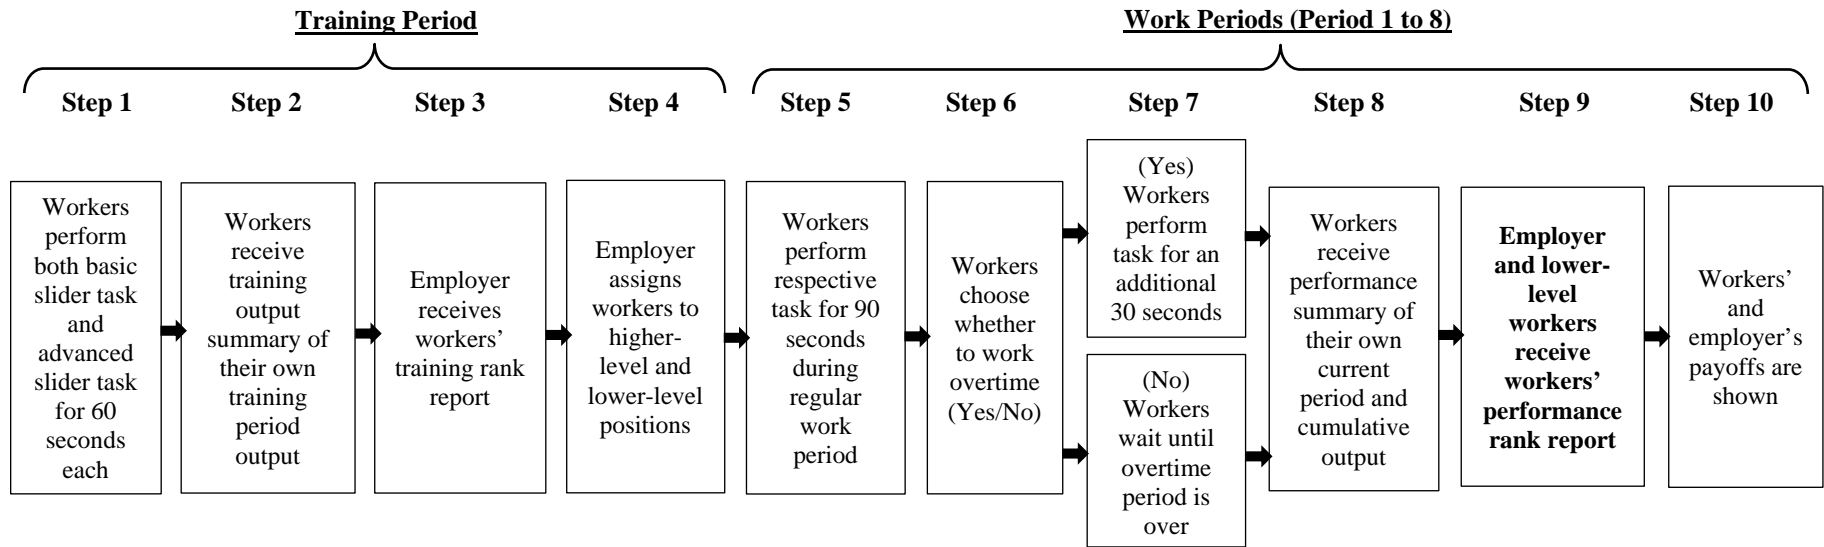
As shown in Panel A of Figure 2, at the end of the training period, the employer assigns one worker to the higher-level position and three workers to the lower-level positions of the firm. Participants are informed that this firm structure is necessary to meet the sales demand for each of the firm's products. This initial worker assignment serves two main purposes. The first is to ensure that workers are aware that their performance during the training period represents their employer's private assessment of their relative abilities to perform each task (see Assumption 3). The second is to reduce lower-level workers' anticipation of a future promotion so that workers' performance during the first two periods represents their baseline performance.

Figure 2: Experimental timeline

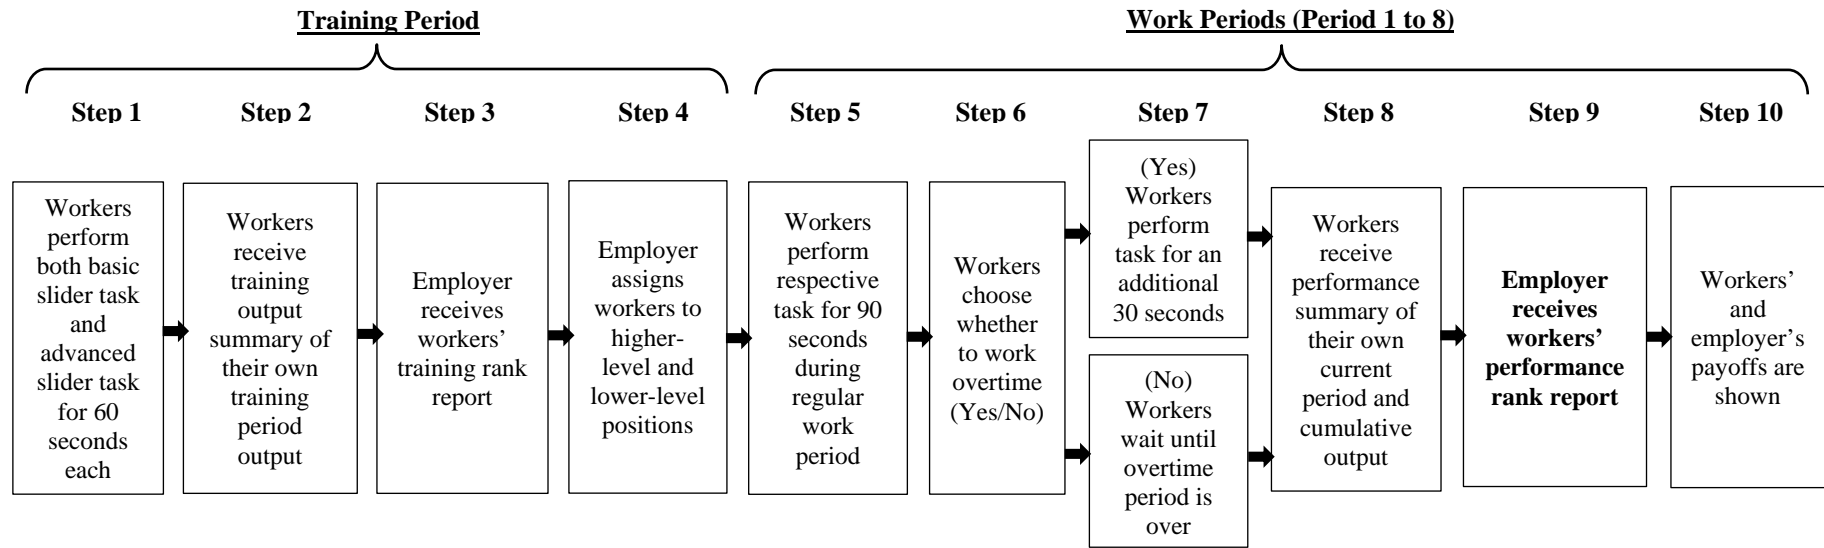
Panel A: Timeline of key events



Panel B: Timeline of RPI condition



Panel C: Timeline of No RPI condition



At the end of Period 2, workers are informed that sales demand is shifting towards Product H, and therefore one of the three lower-level workers will be promoted to a second higher-level position by the employer after Period 5 to meet this demand shift.⁷ Because this promotion announcement is likely unexpected, any change in workers' performance during the next three periods (Period 3 to 5) reflects their response to the promotion announcement (H1).

At the end of Period 5, the employer promotes one of the three lower-level workers to the higher-level position (H2). For the remaining three periods (Period 6 to 8), the promoted worker performs the higher-level task and receives higher pay, and the two non-promoted workers perform the same lower-level task and receive the same pay. This allows me to examine the post-promotion performance of both promoted and non-promoted workers (H3, H4a, and H4b).⁸

4.3 DETAILED PROCEDURES

Participants are randomly assigned to either the role of employer or worker within both the *RPI* and *No RPI* experimental conditions, and they remain in the same roles and in the same firm throughout the session. In total, 220 participants formed 44 unique firms (22 in each condition). Twenty-three sessions were conducted that lasted for approximately 90 minutes each. Participants received a \$5 fee in addition to the payoffs they earn during the experiment.

⁷ Because the focus of my study is on internal promotions, I do not consider situations in which employers can hire workers externally. While precluding this limits the generalizability of my results, there are many settings in which employers find it difficult to hire externally because of a tight labor market and/or prefer to promote internally instead of hire externally (e.g. Parrino 1997; Agrawal et al. 2006).

⁸ In the experiment, workers are not explicitly told whether there will be another promotion in the future to reflect the fact that there is usually some possibility of a future promotion for most workers in the real-world because of some unexpected events (e.g. worker turnover, retirement, change in business environment, etc.).

Participants begin the experiment by reading the instructions and completing a short quiz to ensure that they understand the instructions for both roles. Next, participants learn their randomly assigned roles and proceed to the experiment as described below. At the end of the experiment, participants complete a post-experimental questionnaire and receive their payments. The experiment is conducted on computers using z-Tree software (Fischbacher 2007).

The specific steps for each experimental session are shown in Panel B and Panel C of Figure 2. The procedures for the *RPI* condition are the same as for the *No RPI* condition, except in Step 9 as noted below. I describe the steps for both conditions simultaneously below.

In Step 1, all workers begin with a training period in which they perform both the basic slider task and advanced slider task for 60 seconds each for training and assessment purposes. In Step 2, workers receive a training output summary of the total units of output they produced for each task during the training period. Workers' training output summary includes only their performance and not the performance of any other worker. In Step 3, the employer receives a private training rank report that ranks the performance of the workers from highest to lowest based on their output for each task during the training period (see Assumption 4). In Step 4, the employer assigns one worker to the higher-level position and the other three workers to the lower-level positions. After the worker assignments, the first work period begins in Step 5.

There are a total of eight work periods, and Steps 5 to 10 are repeated for each of these work periods. Each work period consists of a regular period followed by an optional overtime period. In Step 5, lower-level and higher-level workers perform their respective tasks for 90 seconds during the regular period. In Step 6, after the regular period ends, workers can choose whether to work overtime. As shown in Step 7, if workers choose to work overtime, they perform the same task as during the regular period for an additional 30 seconds. As explained in

more detail later, working overtime is personally costly and reduces the workers' compensation for the period, but doing so can increase their output and performance rank. If workers choose not to work overtime, they are free to relax during the overtime period. This overtime option is designed to capture workers' willingness to take costly actions to improve their performance.⁹

In Step 8, workers receive a performance summary of the total units of output, including any overtime period output, they produced in that period and cumulatively for all periods completed so far. In Step 9, in the *RPI* condition, lower-level workers and the employer receive a performance rank report at the end of each period that ranks lower-level workers' performance from highest to lowest based on the *cumulative* output they have produced up to that point (see Assumption 2).¹⁰ This report does not specify whether workers chose to work overtime. However, in the *No RPI* condition, only the employer receives this performance rank report. In Step 10, workers are shown the compensation they earned and their impact on the employer's profit for that period, and the employer is shown the total profit earned from all four workers.

Before making the promotion decision at the end of Period 5, the employer is shown two different sets of information: 1) each worker's training ranking on the *higher-level* task at the end of the training period (see Step 3), and 2) each worker's cumulative performance rank in the *lower-level* task at the end of Period 5 (see Step 9). This design allows me to clearly distinguish whether the employer wants to promote the best current job performer or the worker with the highest ability to perform the next higher-level job. Employers only receive rank information rather than specific performance information to simplify their promotion decisions.

⁹ To capture the notion that work is costly, prior studies have participants either perform a real effort task (e.g., Sprinkle 2000; Taftkov 2013) or choose an effort level from an effort cost schedule (e.g., Hannan et al. 2008; Kuang and Moser 2009). In my experiment, I use a hybrid approach by having participants perform an effort task and also providing them with an option to work overtime at a cost. This way, I capture changes in workers' effort intensity in the task as well as their willingness to incur a personal cost to increase their performance.

¹⁰ The *RPI* condition in my experiment is similar to the *public RPI* condition in Taftkov (2013)'s study, except that participants remained anonymous in my experiment. This is done so that employer and worker participants could not identify each other, which reduces the likelihood that employers will make biased promotion decisions.

One design choice that warrants further discussion is that I do not allow employers to explain their promotion decisions, and this was intentional for several reasons. First, it provides the strongest experimental control.¹¹ Second, because workers know that their employer has information about their relative abilities to perform both tasks, workers with RPI can easily infer the employer's promotion rule, which negates the usefulness of an explanation. For example, if workers with RPI see that the employer did not promote the best current job performer, then the logical alternative is that the employer promoted the worker best suited for the next job.

4.4 PAYOFFS

4.4.1 Worker Payoffs

The monetary amounts in my experiment are expressed in terms of an experimental currency, *Lira*, which is converted to U.S. dollars at a rate of 1,000 *Lira* per U.S. dollar at the end of the experiment.

Worker compensation is accumulated over eight work periods and calculated as follows:

Compensation = Fixed Salary + (Worker Output x Bonus per Output) – Overtime Cost

Lower-level workers receive a fixed salary of 800 *Lira* per period, a variable bonus of 10 *Lira* per unit of output (i.e. each correctly positioned slider), and incur overtime cost of 150 *Lira* per period. Higher-level workers receive a fixed salary of 1,800 *Lira* per period, a variable bonus

¹¹ Allowing employers to explain their promotion decisions complicates my subsequent analyses because employers could give a variety of explanations that could differentially influence workers' behaviors. Some employers may even choose to deceive their workers by giving false explanations. An alternative is to limit the menu of choices available to the employer, but this could cause unwanted demand effects. Thus, not allowing employer explanations across conditions provides the strongest experimental control.

of 30 *Lira* per unit of output, and incur overtime cost of 250 *Lira* per period. These payoff parameters are chosen so that worker participants receive reasonable compensation for their participation. The overtime cost parameters are chosen so that any additional bonus earned by workers during the overtime period will most likely be lower than the personal cost they incurred for working overtime.¹² The increase in pay for the higher-level workers reflects the financial benefits of being promoted. As explained later, each unit of output produced by higher-level workers generates more revenue than those produced by lower-level workers.

Workers' compensation per period includes both fixed and variable pay components, which reflect many compensation schemes in practice.¹³ From a wealth-maximizing perspective, workers under this incentive scheme should maximize their task output without working overtime to maximize their compensation for each period. Thus, any deviations from this wealth-maximizing strategy would indicate a change in workers' actual or perceived incentives.

4.4.2 Employer Payoffs

Employer profit is accumulated over eight work periods and calculated as follows:

$$\text{Profit} = (\text{Worker Output} \times \text{Revenue per Output}) - \text{Worker Compensation}$$

Workers' output includes all output generated during the regular and overtime period. Each unit of output from the lower-level and higher-level task generates 40 *Lira* and 200 *Lira* of revenue for the employer, respectively. The higher revenue per output ratio for the higher-level

¹² My results showed that there were only 17 out of 274 (6%) and 12 out of 90 (13%) instances in which lower-level and higher-level workers earned higher compensation as a result of working overtime, respectively. Thus, the overtime option is almost always costly for workers, and therefore it captures workers' willingness to take costly actions to increase their current job performance and output.

¹³ Workers in actual settings are typically paid a fixed wage plus some form of variable pay such as tips, cash bonuses, recognition awards, and stock options.

task reflects the greater impact higher-level workers typically have on the employer's profit. Although employers are not required to perform any tasks other than deciding the worker assignments after the training period and making the promotion decision at the end of Period 5, they bear greater risk and have a larger range of possible payoffs than worker participants. The employer is shown the same task screens as the workers throughout the experiment.

4.5 PARTICIPANTS

Participants were recruited from the participant pool of an experimental economics laboratory at a large U.S. public university. A total of 220 individuals participated in the study. Participants on average were 21 years old; 56 percent were female; 89 percent were undergraduate students. There are no significant differences across conditions for gender, age, and student status (all p -values > 0.23). Participants earned an average of \$18.50.¹⁴

¹⁴ Employers earned an average of \$23.79; assigned higher-level workers earned \$23.09; promoted higher-level workers earned \$17.36; and non-promoted lower-level workers earned \$14.15. Employers have the highest average earnings but also assumed the highest risk with actual earnings ranging from \$6.80 to \$44.90. As expected, assigned and promoted higher-level workers earned more than lower-level workers.

5.0 EXPERIMENTAL RESULTS

5.1 OVERVIEW OF CHAPTER

This chapter reports the results of the tests of my hypotheses and research question. Section 5.2 presents results that validate a key assumption of my setting. Section 5.3 provides a general overview of the statistical tests of my hypotheses. Sections 5.4, 5.5, 5.6, 5.7, and 5.8 present the formal statistical tests of the hypotheses and research question. Finally, Section 5.9 provides a summary of my experimental results.

5.2 VALIDATION OF KEY ASSUMPTIONS

As described earlier, there are four key assumptions in my setting. Assumption (1) is that the tasks between hierarchical job levels require different worker abilities. To verify that this assumption is satisfied, I performed a Pearson correlation test between workers' output in the lower-level and higher-level tasks during the training period and found a low correlation ($r = 0.14$, $p = 0.04$, one-tailed, untabulated). In addition, participations indicated on a 7-point scale with endpoints of "Not at all different" (1) and "Very different" (7) the extent to which they believe the underlying ability necessary to perform the lower-level and higher-level tasks are

different. The mean response was 5.50, which is at the high end of the scale and indicates that participants believe the two tasks require different abilities. These results indicate that Assumption (1) is satisfied. Assumptions (2), (3), and (4) are satisfied by design as described previously.

5.3 OVERVIEW OF TESTS OF HYPOTHESES

In my statistical tests below, I analyze workers' performance by groups of related work periods based on the key events that occur in the experiment as described earlier. The work period groups I analyze are: 1) Periods 1 and 2, which capture workers' baseline performance before the promotion announcement (*Pre-announcement* period), 2) Periods 3 through 5, which capture performance after the promotion announcement but before the promotion (*Post-announcement* period), and 3) Periods 6 through 8, which capture performance after the promotion (*Post-promotion* period).

As explained earlier, the employer assigns one worker to the higher-level position after the training period. I find that across both experimental conditions, 39 out of 44 employers (18 of 22 in the *No RPI*; 21 of 22 in the *RPI* condition) assigned the worker with the highest rank on the higher-level task during the training period to the higher-level position. This result indicates that employers understand the importance of sorting workers into jobs that best match their abilities. Because the assigned higher-level workers are unaffected by the subsequent promotion, I exclude these workers' performance and their impact on the employer's profit from my analyses below.

5.4 TESTS OF HYPOTHESIS 1

H1 predicts that workers' performance increases after the employer announces an upcoming promotion. Workers' performance is measured by their output produced during each work period. Each work period consists of a regular period followed by an optional overtime period. To test H1, I compare lower-level workers' output during the *Post-announcement* period to their output during the *Pre-announcement* period, pooled across conditions (see Table 1).¹⁵ Consistent with H1, lower-level workers' regular period output, overtime period output, and total output increased significantly (all p -values < 0.01) from the *Pre-announcement* period (33.67, 3.03, and 36.70, respectively) to the *Post-announcement* period (36.20, 5.58, and 41.78, respectively). In addition, the proportion of workers who worked overtime increased significantly ($\chi^2 = 22.94$, $p < 0.01$) from 25% to 44% over the same period.

¹⁵ Lower-level workers' average total output is not significantly different (p 's > 0.63) between the *No RPI* and *RPI* conditions during the *Pre-announcement* period (36.95 vs. 36.47) or *Post-announcement* period (42.13 vs. 41.43). Therefore, I pooled the data across the two conditions in my analyses above.

Table 1: Workers' performance after promotion announcement (Tests of H1)

	<i>Pre- announcement period (1-2) [n = 132]</i>	<i>Post- announcement period (3-5) [n = 132]</i>	(Paired <i>t</i>-test/ Chi-square test)
Lower-level workers' regular period output ^a	33.67 (4.65)	36.20 (4.57)	Diff. = 2.53 (<i>t</i> = 13.02, <i>p</i> < 0.01)
Lower-level workers' overtime period output ^b	3.03 (4.04)	5.58 (5.55)	Diff. = 2.55 (<i>t</i> = 5.34, <i>p</i> < 0.01)
Lower-level workers' total output ^c	36.70 (6.60)	41.78 (8.17)	Diff. = 5.08 (<i>t</i> = 9.39, <i>p</i> < 0.01)
Avg. % of lower-level workers who choose to work overtime	25%	44%	Diff. = 19% ($\chi^2 = 22.94$, <i>p</i> < 0.01)
<i>Increase chances of promotion</i> ^d	2.86 (1.96)	5.49 (1.81)	Diff. = 2.63 (<i>t</i> = 11.98, <i>p</i> < 0.01)
<i>Increase performance rank</i> ^e	4.07 (2.02)	5.61 (1.80)	Diff. = 1.54 (<i>t</i> = 7.95, <i>p</i> < 0.01)
<i>Maximize Compensation</i> ^f	6.31 (1.09)	5.87 (1.47)	Diff. = -0.44 (<i>t</i> = 3.44, <i>p</i> < 0.01)

^a *Workers' regular period output* is the average task output produced by workers during the regular period of each work period.

^b *Workers' overtime period output* is the average task output produced by workers during the overtime period of each work period.

^c *Workers' total output* is the average task output produced by workers during the regular and overtime period of each work period.

^d *Increase chances of promotion* measures the extent to which workers' motivation to perform well on their task is influenced by their desire to increase their chances of promotion in the future on a 7-point scale with endpoints of "No Influence" (1) and "Very high Influence" (7).

^e *Increase performance rank* measures the extent to which workers' motivation to perform well on their task is influenced by their desire to increase their performance rank relative to other workers on a 7-point scale with endpoints of "No Influence" (1) and "Very high Influence" (7).

^f *Maximize Compensation* measures the extent to which workers' motivation to perform well on the slider task is influenced by their desire to maximize their compensation of the current periods on a 7-point scale with endpoints of "No Influence" (1) and "Very high Influence" (7).

To better understand workers' behavior before and after the promotion announcement, I examine their responses to a post-experimental question in which workers indicated on a 7-point scale with endpoints of "No influence" (1) and "Very high influence" (7) the extent to which their motivation to perform well on their task is influenced by their desire to 1) increase their chances of promotion in the future, 2) increase their current performance rank, and 3) maximize their current period compensation. Workers answered the same question for the *Pre-announcement*, *Post-announcement*, and *Post-promotion* periods in the post-experimental questionnaire.

As shown in Table 1, workers' desire to increase their chances of promotion and increase their performance rank increased significantly (all p -values < 0.01) from the *Pre-announcement* period (2.86 and 4.07, respectively) to the *Post-announcement* period (5.49 and 5.61, respectively). Moreover, workers' desire to maximize their current period compensation decreased significantly (6.31 vs. 5.87; $p < 0.01$) over the same period, which is consistent with workers understanding the tradeoff between lower pay in the current periods for possible higher pay from receiving a promotion. Consistent with H1, workers increased their effort and were more willing to work overtime to improve their current job performance after the promotion announcement because they expected this to increase their chances of promotion. These results provide additional support for H1.

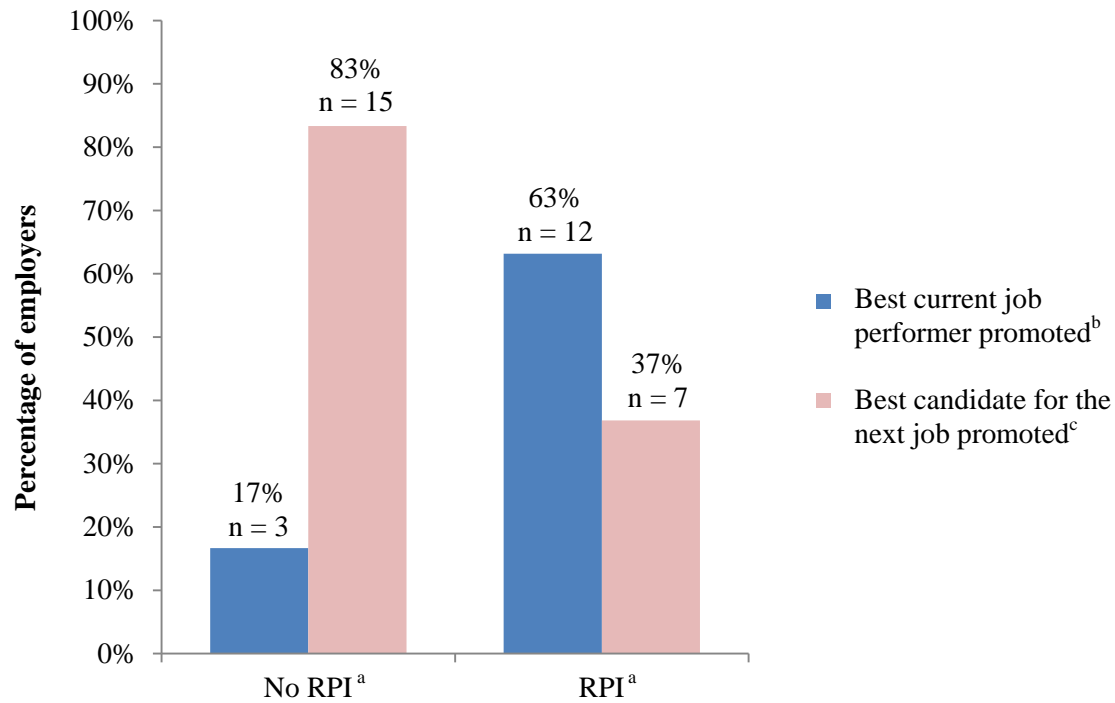
5.5 TESTS OF HYPOTHESIS 2

H2 predicts that employers are more likely to promote the worker with the best current job performance than the worker with the highest ability to perform the next job when workers have RPI than when they do not. To test H2, I compare the employer's promotion decision across the *No RPI* and *RPI* conditions. I define the "best current job performer" as the worker with the highest cumulative performance rank in the lower-level task at the end of Period 5 (i.e. the period before promotion) and the "best candidate for the next job" as the worker with the highest rank in the higher-level task during the training period.

The employer's promotion decision is straightforward if one of the three lower-level workers in a given firm is best at performing both the lower-level and higher-level tasks because it is clear which worker should be promoted. Because my hypotheses are based on the assumption that the best current job performer is not also the best candidate for the next job, I exclude the seven firms (four in the *No RPI*; three in the *RPI* condition) for which this occurs from all subsequent tests below.

Consistent with H2, Figure 3 shows that 12 out of 19 employers (63%) in the *RPI* condition promoted the best current job performer rather than the best candidate for the next job versus 3 out of 18 employers (17%) in the *No RPI* condition. The difference in the employers' promotion decisions across the *RPI* and *No RPI* conditions is significant ($\chi^2 = 8.29, p < 0.01$).

Figure 3: Effect of RPI on employers' promotion decision (Tests of H2)



^a RPI is manipulated as a between-subjects factor at two levels: *No RPI* and *RPI*. Workers in the *No RPI* condition are not provided with any RPI. Workers in the *RPI* condition are provided with their current job performance rank at the end of each work period.

^b *Best current job performer promoted* represents the cases in which the employer chooses to promote the lower-level worker with the highest cumulative performance rank in the at the end of Period 5.

^c *Best candidate for the next job promoted* represents the cases in which the employer chooses to promote the lower-level worker with the highest training rank in the higher-level task during the training period.

I next examine the mechanism by which RPI influences employers' promotion decisions. As discussed in the development of H2, I posit that when employers make a promotion decision, they anticipate non-promoted workers' reactions to their decision. Specifically, I predict that non-promoted workers will likely perceive the promotion decision to be unfair and react negatively if they have RPI and see that the best current job performer was *not* promoted. In contrast, workers who do not have RPI are less likely to perceive their employer's decision to be unfair regardless of which worker was promoted because they cannot identify the best performer.


To test whether my expectations are correct, I examine non-promoted workers' responses to a post-experimental question in which they indicated the extent to which they believe the promoted worker deserved to be promoted (*Deservingness judgment*) on a 7-point scale with endpoints of "Not at all" (1) and "Very much" (7). In addition, they indicated their reactions to their employer's promotion decision (*Promotion reaction*) on a 7-point scale with endpoints of "Very disappointed" (-3) and "Very pleased" (+3) and a midpoint of "Neither pleased nor disappointed" (0).

Based on H2, I expect an ordinal interaction between RPI and the employer's promotion decision, such that non-promoted workers who have RPI and see that the best current job performer was *not* promoted will believe the promoted worker to be *less* deserving of the promotion and will be *more* disappointed than workers who either have RPI and see that the best current job performer was promoted or do not have RPI. Therefore, I performed separate planned contrasts with non-promoted workers' *Deservingness judgment* and *Promotion reaction* as the dependent measures and assigned contrast weights of -3 for the *RPI* condition when the best candidate for the next job was promoted, +1 for the *RPI* condition when the best current job performer was promoted, and +1 for the *No RPI* conditions (Buckless and Ravenscroft 1990).

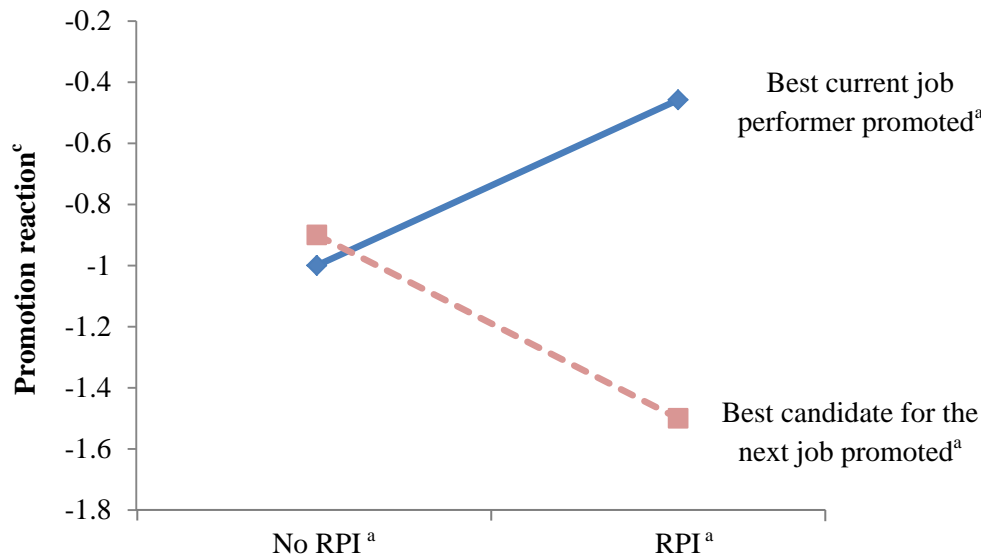
Untabulated planned contrast results show that non-promoted workers' *Deservingness judgment* ($p < 0.01$) and *Promotion reaction* ($p = 0.07$) are consistent with the pattern described above. Panels A and B of Figure 4 depict these results. Untabulated simple effects tests indicate that non-promoted workers believe that the promoted worker was significantly *less* deserving of the promotion (both p 's < 0.02 , one-tailed) when workers have RPI and the best candidate for the next job was promoted (3.29) than when workers either have RPI and the best current job performer was promoted (5.04) or do not have RPI and the best candidate for the next job was promoted (4.47). Similarly, non-promoted workers were significantly *more* disappointed about their employer's promotion decision (both p 's < 0.07 , one-tailed) when workers have RPI and the best candidate for the next job was promoted (-1.50) than when workers either have RPI and the best current job performer was promoted (-0.46) or do not have RPI and the best candidate for the next job was promoted (-0.90).

Figure 4: Non-promoted workers’ response to promotion decision (Tests of H2)

Panel A: Non-promoted workers’ deservingness judgments



Panel B: Non-promoted workers’ promotion reactions



^a See variable descriptions in Figure 3.

^b *Deservingness judgment* measures the extent to which workers believe the lower-level worker their employer chose to promote deserved to be promoted on a 7-point scale with endpoints of “Not at all deserved” (1) and “Very much deserved” (7).

^c *Promotion reaction* measures workers’ reaction to their employer’s promotion decision on a 7-point scale with endpoints of “Very disappointed” (-3) and “Very pleased” (+3).

Consistent with the underlying logic for H2, these results suggest that employers anticipate their non-promoted workers' negative reactions to their promotion decisions in the *RPI* condition and make different decisions depending on whether workers have or do not have RPI.

5.6 TESTS OF HYPOTHESIS 3

H3 predicts that the promoted workers' post-promotion performance is lower when workers have RPI than when they do not. To test H3, I compare workers' output during the *Post-promotion* period across the *No RPI* and *RPI* conditions. Consistent with H3, Panel A of Table 2 shows that the promoted workers' regular period output and total output during the *Post-promotion* period are significantly higher (both p 's < 0.02) in the *No RPI* condition (16.41 and 17.24, respectively) than in the *RPI* condition (12.96 and 13.80, respectively). As expected, I find no significant differences in workers' overtime period output (0.83 vs. 0.84; $p = 0.95$) and an equally low likelihood of working overtime (13% vs. 16%; $\chi^2 = 1.00$, $p = 0.32$) across conditions because there are no future benefits for workers to work overtime after being promoted. These results indicate that it is the difference in the promoted workers' abilities to perform the higher-level job that is driving the results across conditions.

Table 2: Promoted workers' post-promotion performance (Tests of H3)**Panel A: No RPI vs. RPI**

	<i>Post-promotion period (6-8)</i>		<i>(t-test/ Chi-square test)</i>
	No RPI^a [n = 18]	RPI^a [n = 19]	
Promoted higher-level workers' regular period output ^a	16.41 (3.37)	12.96 (3.09)	Diff. = -3.45 ($t = 3.24, p < 0.01$)
Promoted higher-level workers' overtime period output ^a	0.83 (1.96)	0.84 (1.35)	Diff. = 0.01 ($t = 0.02, p = 0.99$)
Promoted higher-level workers' total output ^a	17.24 (4.71)	13.80 (3.60)	Diff. = -3.44 ($t = 2.50, p = 0.02$)
Avg. % of promoted workers who choose to work overtime	13%	16%	Diff. = 3% ($\chi^2 = 0.18, p = 0.67$)
<i>Maximize Compensation^e</i>	6.56 (1.15)	6.32 (1.06)	Diff. = 0.24 ($t = 0.66, p = 0.51$)

Panel B: RPI condition – Best current job performer promoted vs. Best candidate for the next job promoted

	<i>Post-promotion period (6-8)</i>		<i>(t-test/ Chi-square test)</i>
	RPI^a		
	Best current job performer promoted^a [n = 12]	Best candidate for the next job promoted^a [n = 7]	
Promoted higher-level workers' regular period output ^a	11.53 (1.98)	15.43 (3.20)	Diff. = 3.90 ($t = 3.31, p < 0.01$)
Promoted higher-level workers' overtime period output ^a	0.58 (0.88)	1.29 (1.93)	Diff. = 0.71 ($t = 1.10, p = 0.29$)
Promoted higher-level workers' total output ^a	12.11 (2.29)	16.71 (3.69)	Diff. = 4.60 ($t = 3.38, p < 0.01$)
Avg. % of promoted workers who choose to work overtime	11%	24%	Diff. = 13% ($\chi^2 = 1.61, p = 0.20$)
<i>Maximize Compensation^e</i>	6.42 (1.00)	6.14 (1.22)	Diff. = -0.28 ($t = 0.53, p = 0.60$)

^a See variable descriptions in Table 1 and Figure 3.

To provide further evidence that it is the employer's promotion decision that caused the promoted workers' post-promotion performance to be lower in the *RPI* condition than in the *No RPI* condition, I compare their output *within* the *RPI* condition when the best current job performer was promoted versus when the best candidate for the next job was promoted. As shown in Panel B of Table 2, the promoted workers' regular period output and total output during the *Post-promotion* period are significantly lower (both p 's < 0.01) when the best current job performer was promoted (11.53 and 12.11, respectively) than when the best candidate for the next job was promoted (15.43 and 16.71, respectively). Consistent with H3, these results indicate that the difference in the promoted workers' post-promotion performance across the *No RPI* and *RPI* conditions is a consequence of the employer's promotion decision.

5.7 TESTS OF HYPOTHESIS 4A AND 4B

H4a predicts that non-promoted workers' performance declines after a promotion occurs. To test H4a, I compare non-promoted workers' output during the *Post-promotion* period to their output during the *Post-announcement* period, pooled across conditions.¹⁶ Consistent with H4a, Table 3 shows that their overtime period output and total output decreased significantly (all p -values < 0.01) from the *Post-announcement* period (5.18 and 40.84, respectively) to the *Post-promotion* period (1.69 and 37.73, respectively), mainly because fewer workers chose to work overtime (41% vs. 13%; $\chi^2 = 8.97, p < 0.01$).

¹⁶ Non-promoted workers' average regular period, overtime period, and total output are not significantly different (all p -values > 0.13, untabulated) between the *RPI* and *No RPI* conditions during the *Post-announcement* period (35.12, 5.18, and 40.30, respectively vs. 36.21, 5.19, and 41.40, respectively) and *Post-promotion* period (34.90, 2.10, and 37.00, respectively vs. 37.24, 1.27, and 38.51, respectively). Therefore, I pooled the data across the two conditions in my analyses above.

Table 3: Non-promoted workers' post-promotion performance (Tests of H4a)

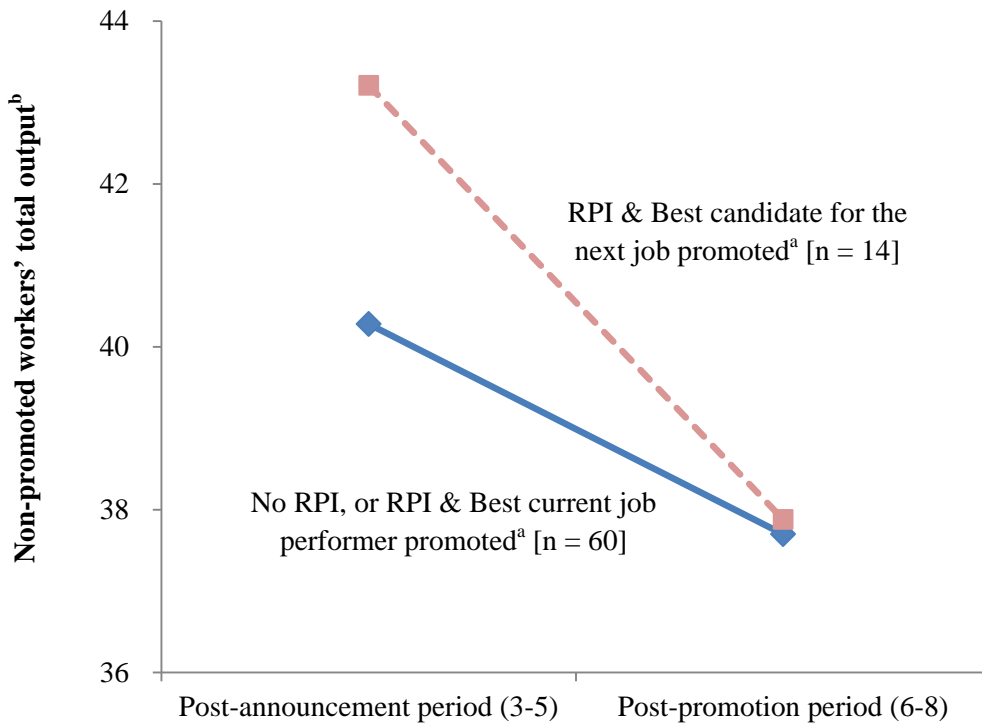
	<i>Post- announcement period (3-5)</i> [n = 74]	<i>Post- promotion period (6-8)</i> [n = 74]	(Paired <i>t</i>-test/ Chi-square test)
Non-promoted lower-level workers' <i>regular period</i> output ^a	35.65 (4.33)	36.04 (6.58)	Diff. = 0.39 (<i>t</i> = 0.61, <i>p</i> = 0.54)
Non-promoted lower-level workers' <i>overtime period</i> output ^a	5.18 (5.52)	1.69 (3.91)	Diff. = -3.49 (<i>t</i> = 5.59, <i>p</i> < 0.01)
Non-promoted lower-level workers' <i>total</i> output ^a	40.84 (8.05)	37.73 (8.40)	Diff. = 3.10 (<i>t</i> = 3.35, <i>p</i> < 0.01)
Avg. % of non-promoted lower-level workers who choose to work overtime	41%	13%	Diff. = -28% ($\chi^2 = 44.37$, <i>p</i> < 0.01)
<i>Increase chances of promotion</i> ^a	5.39 (1.74)	1.91 (1.42)	Diff. = -3.48 (<i>t</i> = 14.01, <i>p</i> < 0.01)
<i>Increase performance rank</i> ^a	5.65 (1.53)	3.11 (2.03)	Diff. = 2.54 (<i>t</i> = 9.50, <i>p</i> < 0.01)
<i>Maximize Compensation</i> ^a	6.03 (1.24)	6.07 (1.48)	Diff. = 0.04 (<i>t</i> = 0.20, <i>p</i> = 0.84)

^a See variable descriptions in Table 1.

To better understand the decrease in the non-promoted workers' performance, I examine their responses to the same post-experimental questions regarding their motivation to perform well on their task used to analyze H1. As shown in Table 3, non-promoted workers' desire to increase their chances of promotion and increase their performance rank decreased significantly (both p 's < 0.01) from the *Post-announcement* period (5.39 and 5.65, respectively) to the *Post-promotion* period (1.91 and 3.11, respectively). Consistent with H4a, these results suggest that non-promoted workers have less incentive to work overtime because future promotion is less likely. However, workers' regular period output and desire to maximize their current period compensation are relatively stable (both p 's > 0.54) from the *Post-announcement* period (35.65 and 6.03, respectively) to the *Post-promotion* period (36.04 and 6.07, respectively) as workers remain motivated to maximize their output-based bonus during the regular period.

H4b predicts that non-promoted workers' performance declines more when they have RPI and see that the best current job performer was *not* promoted than when they either have RPI and see that the best current job performer was promoted or do not have RPI. Although as shown in Figure 5, the decrease in non-promoted workers' total output from the *Post-announcement* period to the *Post-promotion* period is directionally greater when workers have RPI and the best current job performer was *not* promoted (i.e. the best candidate for the next job was promoted) ($37.88 - 43.21 = -5.33$) than when workers either have RPI and the best current job performer was promoted or have no RPI ($37.70 - 40.28 = -2.58$), the decrease is not statistically significant at conventional levels ($t = 1.17$; $p = 0.12$, one-tailed, untabulated).

Figure 5: Interaction effect of RPI and employer's promotion decision on non-promoted workers' post-promotion performance (Tests of H4b)



^a See variable descriptions in Figure 3.

^b *Workers' total output* is the total average task output produced by workers during the regular and overtime period of each work period.

The test above includes all non-promoted workers in both conditions. However, the subset of workers expected to react most negatively are the best current job performers in the *RPI* condition who know that they were passed over for promotion despite the fact that they were the best performers.¹⁷ Therefore, I repeated the test described above with only the best current job performers in the group of workers who have *RPI* and were not promoted. As expected, the decrease in the non-promoted workers' total output from the *Post-announcement* period to the *Post-promotion* period is significantly greater ($t = 1.69$; $p = 0.05$, one-tailed, untabulated) when the best current job performers have *RPI* and see that they were not promoted ($39.19 - 47.29 = -8.10$) than when non-promoted workers either have *RPI* and the best current job performer was promoted or have no *RPI* ($37.70 - 40.28 = -2.58$). Taken together, the results reported above are consistent with H4b. However, these results should be interpreted with caution because I cannot rule out the possibility that they could also partially reflect a decrease in non-promoted workers' willingness to work overtime because there is no opportunity for a future promotion in my setting.¹⁸

5.8 TESTS OF RESEARCH QUESTION 1

RQ1 asks how providing workers with *RPI* affects the employer's profit. Recall that the employer's profit is the sum of the revenue generated by workers' output less the compensation

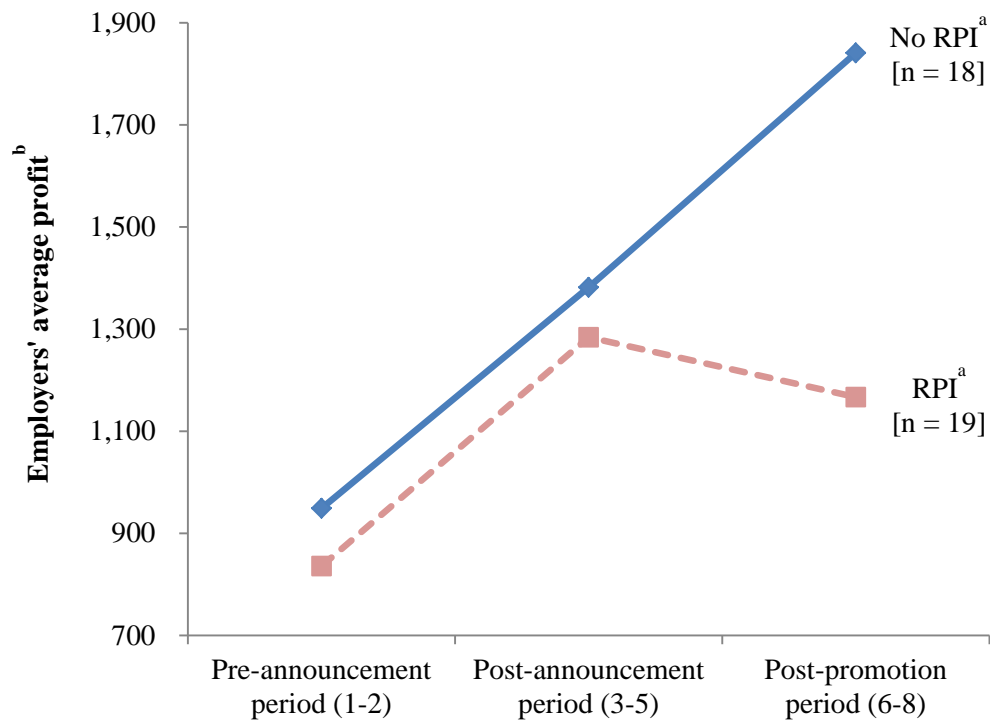
¹⁷ Non-promoted workers were asked to indicate the extent to which they believe their employer should have promoted them instead of another worker on a 7-point scale with endpoints of "Not at all" (1) and "Very much" (7). The best current job performers who had *RPI* and were not promoted (6.00) believed that they should have been promoted more ($t = 1.92$; $p < 0.04$, one-tailed, untabulated) than those who were not the best performers (4.29).

¹⁸ Consistent with this possibility, my tests of H4a show that the decrease in total output in the post-promotion period is mainly due to a decrease in workers' overtime period output.

paid to them. I test RQ1 by comparing the employer's profit in the *No RPI* condition to that in the *RPI* condition. I find that the employer's total profit is significantly higher ($t = 2.18$; $p = 0.04$, two-tailed, untabulated) in the *No RPI* condition (11,568) than in the *RPI* condition (9,026). As shown in Figure 6, the employer's average profit is very similar in the two conditions in the periods before promotion but diverges after the promotion occurs. Specifically, the employer's average profit is not significantly different (both p 's > 0.30 , untabulated) between conditions during the *Pre-announcement* period (949 vs. 836) or the *Post-announcement* period (1,382 vs. 1,284).¹⁹ However, the employer's average profit during the *Post-promotion* period is significantly higher ($t = 2.37$, $p = 0.02$, untabulated) in the *No RPI* condition (1,841) than in the *RPI* condition (1,167). This difference in profit reflects the difference in the promoted workers' post-promotion performance between the *No RPI* and *RPI* conditions (see H3). Overall, these results show that in my setting, the employer's profit after promotion is lower when workers have RPI than when they do not.

¹⁹ Because prior research suggests that RPI increases workers' performance, this result initially appears surprising (Hannan et al. 2008; Tafkov 2013). However, an important difference between my setting and the settings in the prior studies is that in the prior settings there were no employers and workers' pay was not tied to their peer performance. In contrast, workers in both the *No RPI* and *RPI* conditions in my setting know that their employer knows their relative current job performance rank and could use this information for future evaluative purposes. Thus, even workers with no RPI are motivated to increase their performance rank because it could affect their future pay. Consistent with this argument, when asked the extent to which they were concerned about how well they were performing relative to other workers on a 7-point scale with endpoints of "Not at all" (1) and "Very concerned" (7), I find that workers were equally concerned ($t = 1.09$, $p = 0.28$, untabulated) about their current job performance rank across the *No RPI* (4.09) and *RPI* (4.39) conditions.

Figure 6: Effect of RPI on employers' profit (Tests of RQ1)



^a See variable descriptions in Figure 3.

^b *Employer's average profit* is the average profit earned by the employer in each work period.

An interesting question is whether employers in the *RPI* condition are actually better off when they promote the best current job performer rather than the best candidate for the next job. Additional analysis indicates that *within* the *RPI* condition, the employer's average profit during the *Post-promotion* period and total profit are significantly lower (both p 's < 0.05 , untabulated) when the employer promotes the best current job performer (848 and 7,624, respectively) rather than the best candidate for the next job (1,714 and 11,430, respectively). The question then is why many employers chose to promote the best current job performer when workers have *RPI* despite having lower profit as a result. There are two possibilities. First, employers could overestimate the negative impact of *not* promoting the best current job performer on non-promoted workers' subsequent performance. Such employers could continue to promote the best current job performer because the employers would not know the impact of taking alternative actions. That is, employers cannot directly compare the outcomes of different promotion decisions because they can only make one promotion decision at a time. Second, prior research finds that individuals are willing to bear a cost to appear fair to others (Pillutla and Murnighan 1995; Dana et al. 2006). Similarly, employers in my experiment could be concerned about appearing fair to their workers and therefore promote the best current job performer when workers have *RPI* even if it hurts their profit.

5.9 SUMMARY OF RESULTS

My experiment examines whether providing workers with relative performance information influences employers' promotion decisions. I find that after the announcement of an upcoming promotion, workers increase their effort to improve their current job performance

because they expect that this will increase their chances of promotion. Moreover, because employers anticipate that workers who have RPI will react negatively if they see that the best current job performer is not promoted, employers promote the best current job performer rather than the worker best suited for the next job more often when workers have RPI than when they do not. As a result, the promoted workers' performance is lower when workers have RPI because of the mismatch between the promoted workers' abilities and the higher-level task. After being passed over for promotion, non-promoted workers' performance decreases, and this decrease is greater for workers who have RPI and see that the best current job performer was not promoted because they believe this is unfair. Finally, consistent with the Peter Principle, I find that the employer's profit after promotion is lower when workers have RPI because the promoted worker lacks the ability to perform the new job well.

6.0 DISCUSSION OF WAYS TO MITIGATE THE SORTING PROBLEM

6.1 OVERVIEW OF CHAPTER

The results of my experiment reported in the prior chapters provide evidence that providing workers with RPI on their current job performance could lead employers to promote the best current job performer rather than the worker best suited for the next job. As a consequence, the performance of the promoted workers declines because they are not well-suited for their new job, consistent with the Peter Principle.

This chapter discusses how three commonly used firm practices could potentially help firms mitigate the sorting problem I identified in my experimental study. Specifically, Sections 6.2, 6.3, and 6.4 discuss how 1) communication of the promotion criteria, 2) group incentive plans, and 3) subjective performance evaluation could help employers improve the sorting of workers, respectively. Within each section, I also present possible drawbacks that could reduce the effectiveness of each firm practice.

6.2 COMMUNICATION OF THE PROMOTION CRITERIA

6.2.1 Overview

In the setting examined in my experiment, employers did not communicate the promotion criteria to workers. This design choice reflects the fact that employers often do not provide the explicit criteria for promotion because this gives them greater flexibility when making promotion decisions (Campbell 2008). However, employers in practice could choose to communicate the promotion criteria to workers at their discretion. As discussed below, such communication could help employers adjust workers' expectations about the employer's promotion rule and lead to more effective sorting of workers into their jobs. However, such communication could also lower workers' motivation to perform their current jobs.

6.2.2 Advantages

The main finding from my experiment described in the previous chapters is that employers are more likely to promote the best current job performer than the worker best suited for the next job when workers have RPI than when they do not. This leads to the sorting problem known as the Peter Principle in which workers are promoted to jobs for which they are not well-suited. The key driver for my finding is that employers anticipate that their workers expect the best current job performer to be promoted and therefore choose to promote the best current job performer to

avoid possible negative worker reactions. In other words, despite the current job being different from the next higher-level job, employers anticipate that workers expect the best current job performer to be promoted when workers are *not* provided with any information regarding the criteria for promotion.

One possible way to mitigate the sorting problem described above is to adjust workers' expectations by informing them of the promotion criteria. Employers could communicate to their workers that their promotion prospect will be based on both their current job performance and also their abilities to perform the next job. If workers have a clear understanding of the promotion criteria, they may be less likely to expect the best current job performer to be promoted. Consequently, when the employer promotes the worker best suited for the next job rather than the best current job performer, workers who are passed over for promotion are less likely to react negatively because the promotion decision is consistent with the established promotion criteria.

The employer could communicate the promotion criteria either explicitly, implicitly, or both. Explicit communication could occur when workers begin their jobs and at different intervals throughout their job tenure. Repeated communication could help workers internalize the promotion criteria and set their expectations about the employer's promotion rule before the promotion decision is made.

Alternatively, the employer could communicate the promotion criteria implicitly by formally measuring and providing feedback on workers' abilities to perform the next job as part of their individual balanced scorecard (Kaplan and Norton 2001; Lipe and Salterio 2000). The main purpose of the balanced scorecard is to help workers implement the firm's strategy by translating it into a broad, diverse set of performance measures (Kaplan and Norton 1996). The

purpose of measuring and evaluating workers' performance using multiple measures is to motivate workers to improve their performance and abilities across different dimensions that are important to the firm's success.

To highlight the importance of workers' abilities to perform the next job, employers could include it as a separate performance measure in the balanced scorecard. One advantage of this approach over the explicit communication of the promotion criteria is that the employer is not committed to a defined set of criteria for promotion, and therefore this approach provides the employer with greater flexibility and discretion in the promotion decision. For example, in Ittner, Larcker, and Meyer's (2003) field study, they indicate that the retail bank's individual balanced scorecard includes a measure of leadership skills even for lower-level workers, presumably because the firm views such skills as important for workers to advance to higher-level positions.

When workers receive feedback on their abilities to perform the next job as part of their formal performance evaluation, workers will likely expect their employer to use this information when making their promotion decision. Because workers understand the implicit criteria for promotion, the employer can more comfortably promote the worker best suited for the next job rather than the best current job performer, and workers are less likely to react negatively because the promotion decision is not a surprise.

Another advantage of communicating the promotion criteria to workers is that it could increase their perceived procedural fairness of the promotion process. The organizational justice literature (e.g., Greenberg 1987; Colquitt et al. 2001) distinguishes three different types of fairness: 1) fairness of the distribution of outcomes (distributive justice), 2) fairness of the policies and procedures used to determine the outcome distributions (procedural justice), and 3) fairness of the treatment of the individuals receiving the outcomes (interactional justice). While

all three types of organizational justice have important implications for worker behavior, procedural justice is most relevant in the promotion setting because employers' concerns regarding workers' perceived fairness of the promotion rule influences the employer's promotion decision.

Prior research finds consistent empirical evidence that when workers perceive their performance evaluation as unfair, workers have more negative attitudes toward their superiors, lower motivation to work, and lower job performance (e.g., Colquitt et al. 2001; Cohen-Charash and Spector 2001). In addition, prior studies find that communication of the evaluation criteria increases transparency of the evaluation process, which in turn increases workers' perceptions of procedural fairness.

In a survey study by Folger and Konovsky (1989), first-line workers at a manufacturing plant were asked about their views regarding the practices their supervisors used to determine their most recent salary increase and the fairness of those practices. They found that workers perceive the procedural fairness of their performance evaluation to be high when their supervisors inform them of the evaluation criteria in advance. In another survey study by Greenberg (1986), middle managers were asked to recall a past performance evaluation and describe the most important factor that made the evaluation fair or unfair. The results indicate that workers believe the most important determinant of procedural fairness is their supervisor's consistent application of the evaluation criteria. Collectively, these studies show that workers perceive the procedural fairness of the performance evaluation process to be high when they are informed of the evaluation criteria and when the criteria are used consistently.

In summary, communication of the promotion criteria that highlights the importance of workers' abilities to perform the next job could adjust workers' expectations about the

employer's promotion rule and reduce the likelihood that workers will react negatively when the best current job performer is *not* promoted. This allows employers to promote the worker best suited for the next job, avoiding the Peter Principle problem. In addition, communication of the promotion criteria could increase workers' perceptions of procedural fairness by improving the transparency of the promotion process.

6.2.3 Disadvantages

There are several potential drawbacks to communicating the promotion criteria to workers. One major drawback is that it could weaken workers' motivation to perform their current jobs well. Expectancy theory suggests that workers will have lower motivation to work if the link between their current job performance and promotion is weak. Expectancy theory has been widely used to explain worker behavior within organizations (e.g. Vroom 1964, Brownell and McInnes 1986; Kren 1990; Kominis and Emmanuel 2007). It models workers' motivation as a function of three separate components: 1) the perceived link between effort and job performance (expectancy), 2) the perceived link between job performance and reward (instrumentality), and 3) the perceived value of the reward (valence). Workers' perceived strength of these three components is predicted to affect their work motivation. In the promotion setting that I examine, the expectancy component is strong because workers' current job performance is closely tied to their level of effort. The valence component is also strong because promotions are associated with a significant increase in pay in addition to higher status and reputation. However, the strength of the instrumentality component is in question because it depends on workers' expectations about how their performance affects their chances of promotion.

Prior studies find that instrumentality is an important factor that affects workers' job satisfaction, organizational commitment, and job performance (Heneman and Schwab 1972; Lawler 1973; Sims, Szilagyi, and McKemey 1976; Fletcher and Williams 1996). If workers perceive the link between their job performance and the reward to be weak, then their motivation to work is greatly diminished (Vroom 1964; Lawler and Porter 1967).

When workers are informed that the promotion criteria includes both their current job performance *and* their abilities to perform the next job, workers are likely to perceive the link between their current job performance and the reward of promotion to be weaker than when promotion is based only on their current job performance. In particular, in a setting in which the nature of the jobs are substantially different between hierarchical levels, workers could believe that their employer will promote workers based primarily on their abilities to perform the next job. In this case, workers are likely to believe that that increasing their current job performance will have a limited effect on their chance of promotion, and consequently they will be less motivated to perform their current jobs well.

To strengthen the instrumentality component, employers could provide alternative ways to reward workers other than with promotion. For example, employers could reward workers with cash bonuses and other recognition awards for strong current job performance. However, it is unlikely that such alternative incentives could fully replace the financial and non-financial benefits of receiving a promotion (Gibbs 1996; DeVaro 2006).

Besides reducing workers' motivation to perform their current jobs, communication of the promotion criteria could also cause workers to focus more on developing their abilities to perform the next job rather than improving their current job performance. This could have both positive and negative consequences for the firm.

On the positive side, it could encourage workers to develop and acquire skills that will enable them to succeed at higher-level positions. Prendergast (1993) argues that the opportunity for promotion could provide strong incentives for workers to invest in human capital. Consistent with this, a recent field study by Grabner and Moers (2015) finds that bank employees with promotion opportunities invest significantly more time in training to improve their skills for higher-level jobs than those without promotion opportunities. Further, workers who undertake more training are more likely to be promoted. Thus, emphasizing the importance of workers' abilities to perform the next job could lead to better trained workers who are more prepared to perform higher-level jobs.

On the negative side, because promotion opportunities within a firm are typically limited, it may not be in the firm's best interest for all workers to invest in developing their abilities for the next job when only a few will be promoted. With workers' time being a scarce resource, the firm likely prefers that most workers focus on increasing their current job performance rather than developing their higher-level abilities. Further, to the extent that certain abilities are fixed traits that cannot be easily improved through learning, workers' attempts to develop such abilities may not be productive for them or for the firm (Bonner and Sprinkle 2002). Thus, communication of the promotion criteria could redirect workers' time and effort away from their current job and result in lower current job performance.

Another consequence of communicating the promotion criteria is that this could reduce workers' perceived procedural fairness of the promotion process if they believe that they have limited control over their abilities to perform the next job. Prior research finds that workers perceive higher procedural fairness in the performance evaluation process when they have more control over the attainment of the process outcomes (e.g., Greenberg and Folger 1983; Landy,

Barnes, and Murphy 1978; Greenberg 1990). In my setting, workers have greater control over their current job performance than over their abilities to perform the next job because their job performance is tied directly to their level of effort. Thus, if workers believe that their employer will base the promotion decision heavily on their abilities to perform the next job over which they have limited control, then workers could view the process as unfair.

6.3 GROUP INCENTIVE PLANS

6.3.1 Overview

As second firm practice that could help employers improve the sorting of workers is the use of a group incentive plan. Group incentive plans are compensation plans in which workers are assigned to a group (e.g., team, unit, department, firm) and each group member's compensation is determined by the group's performance (Guthrie and Hollensbe 2000; Bohlander and Snell 2007). The most widely adopted group incentive plan is profit-sharing, in which workers receive a share of the firm's profits. It is estimated that approximately 70 percent of large organizations in the U.S. use profit-sharing plans as an incentive to motivate workers (Lawler and Mohrman 2003).

Because profit-sharing plans are the most common type of group incentive plan, I focus my discussion below mainly on profit-sharing plans. However, the discussion could also generalize to other types of group incentive plans such as firm-wide stock option and bonus plans. As discussed below, group incentive plans could improve the sorting of workers by

aligning the interests of the employer and workers. However, such plans may not be as effective as intended if workers' value their own incentives over the group's incentives.

6.3.2 Advantages

Standard agency theory predicts that group incentive plans such as profit-sharing plans could be ineffective because they lead to the free-riding problem (Holmstrom 1982). Specifically, because each worker's reward depends on other workers' effort, the benefit associated with the marginal effort for any single worker is diluted by a factor of $1/n$, where n is the group size. As such, all workers have incentive to shirk, and the equilibrium is one in which workers free-ride off each other. This free-riding problem becomes more severe as the size of the group increases, and thus profit-sharing plans are theoretically ineffective for any reasonably large-sized firms (Weitzman and Kruse 1990).

However, contrary to economic predictions, many prior studies across different companies and industries find that group incentive plans increased worker productivity and firm performance (e.g., Ehrenberg and Milkovich 1987; Jones and Kato 1995; Kruse 1993; Knez and Simester 2001). Some theorists have shown that group incentive plans could promote mutual monitoring and teamwork, which in turn could overcome the moral hazard problems in multi-agent settings (e.g., Ma 1988; Itoh 1991). Other researchers attribute the effectiveness of group incentive plans primarily to behavioral factors, including peer pressure, social norms, and loyalty (Kandel and Lazear 1992).

Prior research finds that workers under group incentive plans often monitor each other's performance. (Kruse 1992). For example, a field study by Knez and Simester (2001) finds that

worker performance at an airline company improved after the introduction of a group-based bonus plan because workers closely monitored their peers within each autonomous work group. Other studies find that group incentive plans also have positive effects on workers' organizational commitment, cooperative behavior, and trust in their colleagues and upper-level management (Coyle-Shapiro et al. 2002; Bayo-Moriones and Larraza-Kintana 2009; Towry 2003).

A primary purpose of group incentive plans is to align workers' interest with those of the employer. When workers have a direct stake in the employer's profit, they are more able to consider the employer's perspective and understand how their own actions and decisions as well as those of their peers affect profit. In other words, workers can more readily place themselves in the positions of their employers under group incentive plans. This could have an important impact on how employers make their promotion decisions.

Recall that in my setting workers know that the abilities necessary to perform the next higher-level job differ from those necessary to perform the current job. Workers also know that the employer can maximize profit by promoting the worker best suited for the next job so that all workers are matched to jobs for which they are best suited. Despite this, I find that workers believe the best current job performer deserves to be promoted and react negatively when they see that the best performer is *not* promoted. One reason for this result may be that because workers work under an individual pay scheme, they focus primarily on their own compensation rather than the employer's profit.

In contrast, when there is a group incentive plan, workers might focus more on the employer's profit because the amount of profit directly affects the workers' compensation. In this case, workers as a group could prefer that the employer take actions that maximizes profit

because this also maximizes the profit available to be shared among the workers. Thus, workers working under a group incentive plan could prefer that the employer promote the worker best suited for the next job because this leads to the highest profit because it effectively sorts workers. Moreover, workers could also more easily understand that it is in the employer's and their own best interests to retain the best current job performers at their current job because this maximizes profit. Therefore, workers are less likely to react negatively when they see that the best current job performer is *not* promoted. In summary, group incentive plans align the interests of the employer and workers, which could improve the sorting of workers because better sorting increases both the employer's and the workers' payoffs.

6.3.3 Disadvantages

Despite the potential advantages described above, there are a number of potential disadvantages of group incentive plans. Although workers as a group could prefer that the employer take actions that maximize profit, there could be a conflict between workers' group incentives and their own personal incentives to get promoted. In any hierarchical firm, workers at different levels will perform jobs with varying degrees of responsibilities. Higher-level workers who perform jobs that require higher-level abilities will inevitably earn higher pay than lower-level workers. Therefore, workers' have personal incentives to be promoted because a promotion is associated with higher pay and other non-financial rewards such as higher reputation and status. As a consequence, self-interested workers may prioritize their personal incentives over the group's incentives, which could have negative consequences for the firm.

For example, some workers could intentionally perform their current jobs poorly and focus on demonstrating their abilities to perform the next job with the expectation that doing so will increase their chances of promotion. This would negatively affect worker productivity, the employer's profit, and workers' collective pay. Further, if enough workers engage in such dysfunctional behavior, then there could be an unraveling of the mutual monitoring mechanism and workers could free-ride on the efforts of other workers as predicted by economic theory.

Another potential disadvantage of group incentive plans is that workers are exposed to greater risk than those under an individual incentive plan (Kruse 1996). In the standard principal-agent framework, an efficient contract should tie the agents' payoff closely to their output because this induces agents to put forth effort. However, when the agents' payoff depends on factors other than their effort, risk-averse agents bear greater risk because there is greater uncertainty surrounding their payoff. Therefore, the optimal contract seeks to balance maximizing agents' incentives to work and minimizing their exposure to risk.

A group incentive plan such as a profit-sharing plan exposes workers to greater risk because firm profit is determined by the actions of other workers and a multitude of external factors over which workers have limited control. For example, an unexpected decline in oil prices could negatively affect a firm's profits and all workers under a profit-sharing plan will suffer because of the change in economic conditions, regardless of their individual effort and performance. When risk-averse workers bear greater risk, the employer must pay workers a risk premium to compensate for the volatility in their payoffs. Therefore, group incentive plans are likely more costly to the employer than individual incentive plans, and this additional cost increases with workers' level of risk-aversion and the volatility in the firm's profit (Weitzman 1980).

Finally, although there is empirical evidence showing a positive correlation between the use of profit-sharing plans and firm performance, these results do not indicate that such plans have no negative effects on workers' performance, but rather that the benefits could outweigh the costs. For example, in a field study at a large financial services firm, Hansen (1997) finds that average worker productivity increased by 17% after the implementation of a group bonus plan. Although the study finds an overall improvement in worker performance, the results indicate a decline in performance for the firm's most productive workers. Similarly, Weiss (1987) finds that after the introduction of a profit-sharing plan at different manufacturing plants, there is a reduction in output over time for workers with above-average performance. These results suggest that group incentive plans could discourage the more able workers from exerting high effort presumably because their marginal benefit of effort is lower.

6.4 SUBJECTIVE PERFORMANCE EVALUATION

6.4.1 Overview

The third firm practice that could potentially help firms improve the sorting of workers through promotions is subjective performance evaluation, which is common and used at all levels of organizations. From rank-and-file workers to high-level executives, workers' performance evaluation and compensation are often somewhat subjectively determined (Murphy and Oyer 2003; Bushman, Indjejikian, and Smith 1996; Moers 2005; Ittner et al. 2003). Evaluators (i.e. employers, managers, or direct supervisors) often have discretion regarding performance

measures and can make *ex post* adjustments to workers' performance evaluation based on factors other than the performance measures specified *ex ante* (Gibbs, Merchant, Van der Stede, and Vargus 2004; Bol 2008). As discussed below, subjective performance evaluation could help employers improve the sorting of workers by allowing them to strategically assign workers' performance ratings. However, such strategic use of performance ratings could also decrease workers' motivation if they perceive their ratings to be inaccurate and/or negatively biased.

6.4.2 Advantages

Subjectivity in the performance evaluation process can potentially improve compensation contracting in many work settings. For example, workers' output could be too noisy to measure accurately, workers' performance could be difficult to quantify, and the available objective measures may not sufficiently capture all dimensions of a worker's job responsibilities (Holmstrom and Milgrom 1991; Hayes and Schaefer 2000; Murphy and Oyer 2003). In such cases, there could be benefits to granting the evaluator discretion to assess workers' performance subjectively.

Economic theory suggests that the use of subjective measures could improve contracting if those measures provide additional, valuable information about workers' performance (Feltham and Xie 1994; Baker, Gibbons, and Murphy 1994). Using only objective measures could cause workers to ignore important, value-enhancing aspects of their job that are unmeasured (Bol 2008). Further, because objective measures are backward-looking in nature, subjective measures could encourage workers to adopt a longer-term view of their actions. Consistent with these arguments, prior studies find that firms are more likely to use subjective performance measures

when the available objective measures are less complete (Murphy and Oyer 2003) and when the firm has greater growth opportunities and longer product time horizons (Bushman et al. 1996).

Another potential benefit of subjective performance evaluation is that the evaluator could reduce workers' downside risk by making *ex post* discretionary adjustments to filter out uncontrollable events that affect workers' performance (Banker and Datar 1989). In an extensive survey of 250 U.S. car dealerships, Gibbs et al. (2004) provide evidence that subjectivity is used to reduce workers' downside risk such as when workers face targets that are difficult to attain or when the firm is in a loss position. Similarly, Bol and Smith (2011) provide experimental evidence that evaluators use their discretion to adjust workers' performance ratings upwards when an uncontrollable factor decreases workers' performance, but make no adjustments when the same factor increases workers' performance. This result suggests that evaluators tend to compensate workers for bad luck, but do not punish workers for good luck.

Although prior studies show that the use of subjective performance evaluation could improve workers' incentives and reduce their risk, it could also introduce biases to the evaluation process (Prendergast 1999). Prior research finds that evaluators are influenced by two main biases when they make subjective performance evaluations: centrality bias and leniency bias. Centrality bias refers to the tendency for evaluators not to differentiate between high and low performers by compressing workers' performance ratings (Motowidlo and Borman 1977). Leniency bias refers to the tendency for evaluators to overstate the performance of low performers by assigning them higher ratings than is warranted (Saal and Landy 1977).

Prior studies argue that these biases stem from evaluators' desire to maintain positive work relationships with their colleagues and to avoid potential confrontations (Harris 1994; Murphy and Cleveland 1991). Because evaluators often work closely with their workers,

providing negative feedback is unpleasant and could damage their relationships (Varma, Denisi, and Peters 1996). Further, when the evaluator gives more compressed and lenient performance ratings, most workers will be satisfied with receiving high ratings that are similar to their peers. Bol (2011) also shows that centrality bias and leniency bias could help reduce the evaluator's costs of gathering information about workers' actual performance.

There is strong empirical evidence supporting both the centrality bias and leniency bias. For example, Bretz, Milkovich, and Read (1992) find that even though most organizations' evaluation systems comprise of five performance levels, evaluators typically rate most workers in the top two levels. In a field study at a large manufacturing firm, Moers (2005) finds that the greater use of subjective performance measures leads to more compressed and higher performance ratings.

Although evaluator biases could lead to inaccurate performance ratings, they are not necessarily detrimental to the firm and could even be beneficial. Providing workers with more favorable performance ratings could increase workers' motivation to work. Consistent with this argument, Bol (2011) finds that leniency bias has a positive effect on workers' future performance because workers perceive the incentive system to be fairer when they receive higher performance ratings.

Subjectivity could also play an important role in a promotion setting. In my experiment, there was a single objective measure of workers' current job performance. Because workers performed a single task and the task output was closely tied to their level of effort, there was no subjectivity in workers' performance evaluation. However, as described earlier, workers in practice often perform multiple tasks and their performance could be affected by many uncontrollable factors and events. Thus, evaluators in actual business settings often have some

discretion over their workers' performance evaluation. Further, the same individuals who evaluate workers' performance may also be responsible for deciding which worker to promote.

When evaluators make promotion decisions, they must consider how the performance ratings they assign could affect workers' reactions to their promotion decisions. Thus, centrality bias and leniency bias could be problematic for the evaluator. Specifically, if the evaluator compresses workers' performance ratings and give equally high ratings to all workers, then all workers are likely to believe that they deserve to be promoted and could react negatively if they are not. That is, giving more compressed and lenient performance ratings could lead to higher perceived unfairness and possible negative worker reactions to the promotion decision.

Alternatively, evaluators could benefit from inflating the performance rating of the worker whom the evaluator wants to promote while deflating the ratings of other workers. By creating a contrast in performance ratings between the promoted worker and all other workers, which I term "contrast bias," the evaluator can more easily justify why the promoted worker deserves the promotion. Contrast bias lowers the promotion expectations of the non-promoted workers who receive lower performance ratings, and it increases the perceived fairness of the promotion process because workers see that the highest rated worker was promoted. Therefore, subjectivity allows the evaluator to justifiably promote the worker best suited for the next job rather than the best current job performer, resulting in more effective sorting of workers.

Prior studies provide evidence that evaluators are often strategic when they make subjective performance evaluations. In a meta-analysis study by Jawahar and Williams (1997), they found that evaluators give more lenient performance ratings when the ratings are used for compensation and promotion purposes, but give more accurate and harsher feedback for developmental purposes to help workers identify and correct their performance deficiencies. In

addition, Bol and Leiby (2015) provide experimental evidence that managers evaluate the same actions taken by their subordinates differently when assessing their job performance versus their promotion prospects. These findings support the notion that evaluators assess workers' performance differently depending on the purpose of the performance evaluation. In summary, subjectivity in the performance evaluation process allows evaluators to exercise discretion and assess workers' performance strategically in a way that could improve their sorting of workers through promotions.

6.4.3 Disadvantages

There are several potential factors that could reduce the effectiveness of subjective performance evaluation as a means of improving the sorting of workers. One is that the information environment could affect whether evaluators can strategically assign workers' performance ratings. In some environments, the objective measures of workers' relative current job performance are transparent. For example, workers may be able to observe their peers' performance when they work in close proximity of each other. The evaluator has less flexibility to manipulate workers' performance ratings when there is such low information asymmetry between the evaluator and workers regarding workers' relative current job performance. In such cases, the evaluator must provide more accurate performance ratings to avoid being perceived as unfair. Thus, the transparency of the objective measures of workers' current job performance could prevent evaluators from strategically assigning performance ratings to improve the sorting of workers.

Another crucial factor is the level of trust between the evaluator and workers (Gibbs et al. 2004). If workers do not trust that the evaluator will make fair performance assessments, they could perceive the evaluation process as unfair. Trust is also required to overcome workers' rational belief that their evaluators are influenced by their own self-interests (Prendergast and Topel 1993; Ferris and Judge 1991; Bol 2011). If workers believe their evaluator has incentives to negatively bias their performance ratings, workers could be less motivated to work because their effort could go unrewarded.

As described earlier, contrast bias could cause evaluators to deflate the performance ratings of the best current job performers in order to keep them from being promoted. However, continuous deflation of workers' performance ratings could discourage them because their best efforts would not lead to higher ratings and better rewards. Consequently, the best performers could lose self-confidence and motivation, resulting in lower performance (Benabou and Tirole 2002). This would defeat the intended purpose of retaining the best performers at their current jobs to maximize performance. Thus, contrast bias could cause workers to respond to inaccurate negative feedback in a negative manner.

Finally, prior research finds that evaluator discretion could introduce favoritism. Prendergast and Topel (1993) argue that when the evaluators are not the residual claimants, they are more likely to be affected by their own self-serving interests and preferences. For example, evaluators could give positively biased ratings to reward workers for reasons other than their job performance, such as their loyalty and job tenure (Longenecker, Sims, and Gioia 1987). Conversely, evaluators could give negatively biased ratings to punish workers whom they dislike despite their high job performance. As a result, there are incentives for workers to pursue

influence activities to obtain higher performance ratings that are unproductive for the firm (Ferris and Judge 1991; Milgrom 1988).

7.0 DISCUSSION AND CONCLUSION

7.1 OVERVIEW OF CHAPTER

Section 7.2 summarizes the main findings of my experiment and the practices that could help firms mitigate the sorting problem identified in my experiment. Section 7.3 discusses the contribution my dissertation to the academic literature and to practice. Section 7.4 concludes this chapter by discussing the limitations of my experiment and describes potential future projects related to my findings

7.2 DISCUSSION OF FINDINGS

My experimental study investigates the effects of providing workers with relative performance information (RPI) on employers' promotion decisions and the impact of those decisions on worker performance over time. In my experimental setting, the job after promotion requires higher-level abilities than the current job. I find that workers increase their effort and are more willing to work overtime to increase their current job performance after a promotion opportunity is announced because they expect this to increase their chances of being promoted. Because employers anticipate that workers who have RPI will react negatively if they see that the best

current job performer is not promoted, employers promote the best current job performer rather than the worker best suited for the next job more often when workers have RPI than when they do not. After being passed over for promotion, non-promoted workers' performance decreases, and this decrease is greater for workers who have RPI and see that the best current job performer was not promoted because they view this as unfair. Finally, consistent with the Peter Principle, I find that the promoted workers' performance and employer's profit after promotion are lower when workers have RPI because the promoted worker lacks the ability to perform the new job well.

There are several ways for firms to potentially mitigate the sorting problem described above. Employers could 1) communicate the promotion criteria to workers to adjust workers' expectations about the promote rule, 2) implement group incentive plans to align the workers' interests with those of their own, and 3) strategically assign workers' subjective performance ratings to help justify their promotion decisions. However, while each approach has advantages, each also has drawbacks. One common drawback among these approaches is that they could cause workers to have lower incentive and motivation to perform their current jobs well. Thus, it appears that there is no promotion-based incentive system that provides optimal effort incentives to workers while simultaneously achieving the best match between workers and their jobs.

7.3 CONTRIBUTIONS

My study contributes to theory and practice in several ways. First, I extend the promotion literature. Data limitations have led archival and field researchers to focus primarily on the incentive role of promotions rather than the sorting role (e.g., Gibbs 1995; Campbell 2008;

Cichello et al. 2009; Ederhof 2011). Promotion studies that rely on tournament theory assume that promotions are always awarded to the best current job performer, without considering whether the promoted worker's abilities match those required for the new job (e.g. Lazear and Rosen 1979, Nalebuff and Stiglitz 1983). Because the prior literature largely ignores the sorting role, it provides an incomplete picture of why and how promotions are used by firms. My study examines how employers make promotion decisions in a setting in which the sorting of workers is important. I show that it is important for employers to balance both roles of promotions by considering how their promotion decisions affect both the promoted workers and the non-promoted workers. As described in Chapter 6, the conflict between the sorting and incentives roles of promotions is difficult to resolve and firms must ultimately make a tradeoff between the two.

Second, my study advances our understanding of the possible causes of the Peter Principle. Although there is empirical support for the Peter Principle's prediction that workers' performance declines after promotion, there is debate regarding the underlying cause. Prior analytical studies offer various economic-based explanations for the Peter Principle but they lack supporting empirical evidence (e.g. Bernhardt 1995; Lazear 2004; Fairburn and Malcolmson 2001). Moreover, these prior studies suggest that the Peter Principle is a consequence of the employers' optimal actions.

My study identifies a new factor that leads to behavior that is consistent with the Peter Principle. Specifically, I find evidence that providing workers with RPI causes employers to promote the best current job performer instead of the worker best suited for the next job. As a result, the promoted workers are not well-suited for their new jobs, leading to lower productivity. Further, my results indicate that employers are financially worse off when they promoted the

best current job performer when workers have RPI, which suggests that the Peter Principle is a consequence of employers making inefficient decisions.

Finally, my study provides new insights on the effects of RPI. Prior research has focused primarily on the direct effects of RPI on worker performance (e.g. Hannan et al. 2008; Tafkov 2013). Prior field and experimental studies find that RPI generally has a positive effect on worker performance because workers make social comparisons and are motivated to outperform their peers and avoid shame (e.g., Blanes i Vidal and Nossol 2011; Tafkov 2013). My study is the first to suggest that providing workers with RPI can lead to ineffective sorting of workers and lower worker performance after promotion. Firms can use my findings to better weigh the costs and benefits of RPI when designing their information policies. Although there are possible approaches that firms can use to mitigate the sorting problem I identify in my study, each of these approaches appears to have drawbacks that make them imperfect solutions.

7.4 LIMITATIONS AND FUTURE RESEARCH

My experiment is subject to several limitations that could limit the generalizability of my results, but these limitations also provide opportunities for future research. First, in the setting examined in my experiment, employers had no direct interactions with their workers. This design choice was made to maintain strong experimental control because such interactions could influence the employers' promotion decisions. However, it is not clear whether allowing employers to interact with their workers would weaken or strengthen my results.

On one hand, as described in Section 6.2, communication of the promotion criteria to workers could improve the sorting of workers because it could adjust workers' expectations and reduce their perceived unfairness of the promotion decision. However, there are possible negative consequences to such communication, including lower worker motivation in their current jobs. On the other hand, interaction between employers and workers builds trust, and employers might be more inclined to act in a reciprocal manner by rewarding and promoting the best current job performer, which exacerbates the sorting problem. Future research could examine whether allowing employers to interact and communicate with their workers affects workers' and employers' behavior.

Second, only a single promotion occurs in my experiment, precluding possible reputation formation. If multiple promotions were possible, workers could learn their employer's promotion rule over time and change their work behavior accordingly. In a multi-period setting, the first promotion decision that the employer makes is important because it establishes workers' expectations about the criteria for future promotions. If workers have RPI and see that the best current job performer is promoted, then workers will likely expect the employer to continue using this same promotion rule in the future. In this case, workers could be more motivated to perform their current jobs well because they believe this will likely increase their chances of promotion. In contrast, if workers see that the best current performer is *not* promoted, workers could be less motivated to perform their current jobs well when future promotion opportunities are announced because they are likely to believe that their current job performance will have a limited effect on their chances of promotion. Thus, my results could be strengthened in a multi-period setting. Future research could extend my study to a setting in which multiple promotions

occur over a longer time horizon to examine whether and how this affects workers' and employers' behavior.

Third, my study takes a firm policy perspective and focuses on the employer's provision of RPI to workers. However, firms do not always have control over the dissemination of workers' RPI. For example, workers could obtain informal RPI by observing peers who work in close proximity (Hannan et al. 2013). In such cases, the generalizability of my results will likely depend on the level of information symmetry between the employer and the workers. If the employer and workers share the same informal RPI as they do the more formal RPI in my experiment, I would expect my results to generalize to such settings. However, if the workers have informal RPI and the employer either does not have RPI or has a different type of formal RPI, the degree to which my results will generalize is an open empirical question because workers' expectations about the employer's promotion rule are not obvious. Future research could examine the extent to which my results generalize to settings with informal RPI.

Finally, future research could attempt to replicate my findings in the field. Although no prior empirical study has directly tested the research questions I examine in my experiment, several prior studies provide indirect evidence consistent with my results. For example, Anderson et al. (1999) find that retail firms often promote their top salespersons to sales managers even though they are not well-suited for their new managerial role. Although their study is silent about whether the salespersons in their setting have RPI, we know from other studies that sales contests are prevalent in the retail industry and sales staff often have RPI (Kalra and Shi 2001; Murphy and Sohi 1995). Together these prior findings suggest that, consistent with my results, the Peter Principle can occur when workers have RPI. In contrast, in a field setting in which workers did not have RPI, Grabner and Moers (2013) find that bank managers

place less weight on their subordinates' current job performance and more weight on their ability to perform the next job when making promotion decisions.²⁰ Thus, consistent with my results, their findings suggest that the Peter Principle may not occur when workers do not have RPI. Future research could further examine the relation between workers' information environment and the occurrence of the Peter Principle in field settings.

²⁰ The authors confirmed in private correspondence that the bank workers in their field setting did not receive RPI that would allow them to compare their current job performance with that of their peers.

BIBLIOGRAPHY

- Acosta, P. 2010. Promotion dynamics the Peter Principle: Incumbents vs. external hires. *Labour Economics* 17(6): 975-986.
- Agrawal, A., C. R. Knoeber, and T. Tsoulouhas. 2006. Are outsiders handicapped in CEO successions? *Journal of Corporate Finance* (12): 619–644.
- Anderson, D., C. Crowell, S. Sponsel, M. Clarke, and J. Brence. 1982. Behavior management in the public accommodations industry: A three-project demonstration. *Journal of Organizational Behavior Management* 4: 33–66.
- Anderson, R. E., A. J. Dubinsky, and R. Mehta. 1999. Sales Managers: Marketing's Best Example of the Peter Principle? *Business Horizons* 4, 2, 1, 19: 19-26.
- Antle, R. and A. Smith. 1986. An Empirical Investigation of the Relative Performance Evaluation of Corporate Executives. *Journal of Accounting Research* 24(1): 1-39.
- Baker, G., M. Gibbs, and B. Holmstrom. 1994a. The Wage Policy of a Firm. *Quarterly Journal of Economics* 109: 881–919.
- Baker, G., M. Gibbs, and B. Holmstrom. 1994b. The Internal Economics of the Firm: Evidence from Personnel Data. *Quarterly Journal of Economics* 109: 921–955.
- Baker, G. P., R. Gibbons and K. J. Murphy. 1994. Subjective performance measures in optimal incentive contracts. *Quarterly Journal of Economics* 109(4): 1125-1156.
- Baker, G. P.; M. C. Jensen, and K. J. Murphy. 1988. Compensation and Incentives: Practice vs. Theory. *Journal of Finance* 43: 593–616.
- Banker, R. D. and S. M. Datar. 1989. Sensitivity, precision, and linear aggregation of signals for performance evaluation. *Journal of Accounting Research* 27(1): 21-39.
- Bayo-Moriones, A., and M. Larraza-Kintana. 2009. Profit-sharing plans and affective commitment: Does the context matter? *Human Resource Management* 48(2): 207-226.
- Beach, S., and A. Tesser. 1995. Self-esteem and the extended self-evaluation maintenance model: The self in social context. In *Efficacy, Agency, and Self-Esteem* 145–170. New York, NY: Plenum Press.

- Benabou, R. and J. Tirole. 2002. Self-confidence and personal motivation. *The Quarterly Journal of Economics* 117(3): 871-915.
- Bernhardt, D. 1995. Strategic promotion and compensation. *Review of Economic Studies* 62: 315–339.
- Blanes i Vidal, J., and M. Nossol. 2011. Tournaments without Prizes: Evidence from Personnel Records. *Management Science* 57 (10): 1721-1736.
- Bohlander, G., and S. Snell. 2007. Managing human resources. *Australia: Thomson South-Western*.
- Bol, J. C. 2008. Subjectivity in compensation contracting. *Journal of Accounting Literature* 27: 1–27.
- Bol, J. 2011. The Determinants and Performance Effects of Managers' Performance Evaluation Biases. *The Accounting Review* 86: 1549-1575.
- Bol, J., and J. Leiby. 2015. Subjectivity in Professionals' Incentive Systems: Differences between Promotion- and Performance-Based Assessments. Working Paper. *Tulane University and University of Florida*.
- Bol, J., and S. Smith. 2011. Spillover Effects in Subjective Performance Evaluation: Bias, Fairness, and Controllability. *The Accounting Review* 86: 1213-1230.
- Bonner, S. and G. Sprinkle. 2002. The effects of monetary incentives on effort and task performance: Theories, evidence, and a framework for research. *Accounting, Organizations and Society* 27: 303-345.
- Bretz, R. D., G. T. Milkovich, and W. Read. 1992. The current state of performance appraisal research and practice: Concerns, directions, and implications. *Journal of Management* 18 (2): 321–352.
- Brownell, P. and P. McInnes. 1986. Budgetary participation, motivation and management performance. *The Accounting Review* 61: 587–600.
- Buckless, F. A., and S. P. Ravenscroft. 1990. Contrast coding: A refinement of ANOVA in behavioral analysis. *The Accounting Review* 65: 933–945.
- Bushman, R., R. Indjejikian, and A. Smith. 1996. CEO compensation: The role of individual performance evaluation. *Journal of Accounting and Economics* 21 (2): 161–193.
- Campbell, D. 2008. Nonfinancial Performance Measures and Promotion-Based Incentives. *Journal of Accounting Research* 46: 297–332.

- Cichello, M. S., C. E. Fee, C. J. Hadlock, and R. Sonti. 2009. Promotions, Turnover, and Performance Evaluation: Evidence from the Careers Division Managers. *The Accounting Review* 84: 1119–43.
- Cohen-Charash, Y., and P. E. Spector. 2001. The role of justice in organizations: A meta-analysis. *Organizational Behavior and Human Decision Processes* 86: 278–321.
- Colquitt, J., E. Conlon, M. Wesson, C. Porter, and K. Ng. 2001. Justice at the millennium: A meta-analytic review of 25 years of organizational justice research. *Journal of Applied Psychology* 86: 425–445.
- Coyle-Shapiro, J. A-M., P.C. Morrow, R. Richardson, and S. R. Dunn. 2002. Using profit sharing to enhance employee attitudes: A longitudinal examination of the effects on trust and commitment. *Human Resource Management* 41(4): 423-439.
- Dana, J., D. M. Cain, and R. M. Dawes. 2006. What You Don't Know Won't Hurt Me: Costly (But Quiet) Exit in Dictator Games. *Organizational Behavior and Human Decision Processes* 100(2): 193–201.
- DeVaro, J. 2006. Strategic Promotion Tournaments and Worker Performance. *Strategic Management Journal* 27: 721–740.
- Ederhof, M. 2011. Incentive Compensation and Promotion-Based Incentives of Mid-Level Managers: Evidence from a Multinational Corporation. *The Accounting Review* 86: 131–53.
- Ehrenberg, R. and G. T. Milkovich. 1987 Compensation and Firm Performance. In Kleiner M. et al. (eds.) *Human Resource Management and the Performance of the Firm*, IRRA, Madison.
- Evans, J. H., Y. Hwang, N. Nagarajan, and K. Shastri. 1997. Involuntary benchmarking and quality improvement: The effect of mandated public disclosure on hospitals. *Journal of Accounting, Auditing & Finance* 12 (3): 315–352.
- Fairburn, J. A., and J. M. Malcomson. 2001. Performance, Promotion, and the Peter Principle. *Review of Economic Studies* 68: 45–66.
- Feather, N. T. 1999a. Judgments of deservingness: Studies in the psychology of justice and achievement. *Personality and Social Psychology Review* 3: 86–107.
- Feather, N. T. 1999b. Values, achievement, and justice: Studies in the psychology of deservingness. *New York: Kluwer Academic/Plenum*.
- Feather, N. T. 2008. Perceived Legitimacy of a Promotion Decision in Relation to Deservingness, Entitlement, and Resentment in the Context of Affirmation Action and Performance. *Journal of Applied Social Psychology* 38(5): 1230–1254.

- Fehr, E., L. Goette and C. Zehnder. 2009. A behavioral account of the labor market: The role of fairness concerns. *Annual Review of Economics* 1(1): 335-384.
- Fehr, E. and K. Schmidt. 1999. A Theory of Fairness, Competition, and Cooperation. *Quarterly Journal of Economics* 114 (3): 817-68.
- Feltham, G. A. and J. Xie. 1994. Performance measure congruity and diversity in multi-task principal/agent relations. *The Accounting Review* 69(3): 429-453.
- Ferris, G. R. and T. A. Judge. 1991. Personnel/human resources management: A political influence perspective. *Journal of Management* 17(2): 447-488.
- Festinger, L. 1954. A theory of social comparison. *Human Relations* 7: 117–140.
- Fischbacher, U. 2007. z-Tree: Zurich toolbox for ready-made economics experiments. *Experimental Economics* 10 (2): 171-178.
- Fletcher, C., and R. Williams. 1996. Performance Management, Job Satisfaction and Organizational Commitment. *British Academy of Management* 7 (2): 169–79.
- Folger, R., and M. A. Konovsky. 1989. Effects of procedural and distributive justice on reactions to pay raise decisions. *Academy of Management Journal* 32: 115–130.
- Gibbs, M., K. A. Merchant, W. A. Van der Stede and M. E. Vargus. 2004. Determinants and effects of subjectivity in incentives. *The Accounting Review* 79(2): 409-436.
- Gibbons, R. and K. J. Murphy. 1990. Relative Performance Evaluation for Chief Executive Officers. *Industrial and Labor Relations Review* 43(3): 30-51.
- Gibbons, R. and K. J. Murphy. 1992. Optimal Incentive Schemes in the Presence of Career Concerns: Theory and Evidence. *Journal of Political Economy* 100: 468-505.
- Gibbs, M. 1995. Incentive Compensation in a Corporate Hierarchy. *Journal of Accounting and Economics* 19: 247–77.
- Gibbs, M. 1996. Promotions and Incentives. Working paper. *University of Chicago*.
- Gibbs, M., and W. Hendricks. 2004. Do Formal Salary Systems Really Matter? *Industrial and Labor Relations Review* 58 (1): 71–93.
- Gill, D. and V. Prowse. 2011. A novel computerized real effort task based on sliders. *IZA Discussion Paper No. 5801*.
- Gino, F., and B. Staats. 2011. Driven by Social Comparisons: How Feedback about Coworkers' Effort Influences Individual Productivity. Working paper. *Harvard Business School*.

- Grabner, I., and F. Moers. 2013. Managers' Choices of Performance Measures in Promotion Decisions: An Analysis of Alternative Job Assignments. *Journal of Accounting Research* 51(5): 1187-1220.
- Grabner, I., and F. Moers. 2015. Implicit Incentives for Human Capital Acquisition. Working paper. *Maastricht University*.
- Greenberg, J. 1986. Determinants of perceived fairness of performance evaluations. *Journal of Applied Psychology* 71: 340-342.
- Greenberg, J. 1987. Reactions to procedural injustice in payment distributions: Do the means justify the ends? *Journal of Applied Psychology* 72: 55-61.
- Greenberg, J. 1990. Looking fair vs. being fair: managing impression of organizational justice. *Research in Organizational Behavior* 12: 111-157.
- Greenberg, J., and Folger, R. 1983. Procedural justice, participation and the fair process effect in groups and organizations. In *Basic Group Process edited by P. Paulus: 235-236. New York: Springer-Verlag*.
- Guthrie, J.P. and E.C. Hollensbe. 2004. Group incentives and performance: A study of spontaneous goal setting, goal choice and commitment. *Journal of Management* 30: 263-84.
- Hannan, R. L., R. Krishnan, and A. H. Newman. 2008. The effects of disseminating relative performance feedback in tournament versus individual performance compensation plans. *The Accounting Review* 83: 893-913.
- Hannan, R. L., G. P. McPhee, A. H. Newman, and I. D. Tafkov. 2013. The effect of relative performance information on performance and effort allocation in a multi-task environment. *The Accounting Review* 88 (2): 553-75.
- Hansen, D. 1997. Worker Performance and Group Incentives: A Case Study. *Industrial & Labor Relations Review* 51: 37-49.
- Harris, M. M. 1994. Rater motivation in the performance appraisal context: A theoretical framework. *Journal of Management* 20 (4): 737-756.
- Hayes, R. M. and S. Schaefer. 2000. Implicit contracts and the explanatory power of top executive compensation for future performance. *Rand Journal of Economics* 31(2): 273-293.
- Heneman, H. G. III, and D. P. Schwab. 1972. Evaluation of research on expectancy theory predictions of employee performance. *Psychological Bulletin* 78 (1): 1-9
- Holmstrom, B. 1982. Moral hazard in teams. *Bell Journal of Economics* 13: 324-340.

- Holmstrom, B. and P. Milgrom. 1991. Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics and Organization* 7(2): 24-53.
- Itoh, H. 1991. Incentives to Help in Multi-Agent Situations. *Econometrica* 59(3): 611-636.
- Ittner, D.C., F. D. Larcker, W. M. Meyer. 2003. Subjectivity and the weighting of performance measures: evidence from a balanced scorecard. *The Accounting Review* 78(3): 725-759.
- Jawahar, I. M. and C. R. Williams. 1997. Where all the children are above average: The performance appraisal purpose effect. *Personnel Psychology* 50(4): 905-926.
- Jones, D. C. and T. Kato. 1995. The productivity effects of employee stock-ownership plans and bonuses: Evidence from Japanese panel data. *American Economic Review* 85(3): 391-414.
- Kale, J., E. Reis, and A. Venkateswaran, 2009. Rank-order tournaments and incentive alignment: The effect on firm performance. *Journal of Finance* 64: 1479-1512.
- Kalra, A. and M. Shi. 2001. Designing Optimal Sales Contests: A Theoretical Perspective. *Marketing Science* 20: 170-193.
- Kandel, E., and E. P. Lazear. 1992. Peer pressure and partnerships. *Journal of Political Economy* 100(4): 801-817.
- Kaplan, R.S., and D.P. Norton. 1996. *The Balanced Scorecard-Translating Strategy into Action*, Boston, *Harvard Business School Press*.
- Kaplan, R.S., and D.P. Norton. 2001. Transforming the Balanced Scorecard from Performance Measurement to Strategic Management: Part II. *Accounting Horizons* 15: 147-160.
- Knez, M., and D. Simester. 2001. Firm-wide Incentives and Mutual Monitoring at Continental Airlines. *Journal of Labor Economics* 19(4): 743-772.
- Komins, G., and C. R. Emmanuel. 2007. The expectancy-valence theory revisited: Developing an extended model of managerial motivation. *Management Accounting Research* 18 (1): 49-75.
- Konovsky, M. A. and R. Cropanzano. 1991. Perceived Fairness of Employee Drug Testing as a Predictor of Employee Attitudes and Job Performance. *Journal of Applied Psychology* 76 (5): 698-707.
- Kren, L. 1990. Performance in a budget-based control system: an extended expectancy theory model approach. *Journal of Management Accounting Research* 2: 100-112.

- Kruger, J. 1999. Lake Wobegon be gone! The "below-average effect" and the egocentric nature of comparative ability judgments. *Journal of Personality and Social Psychology* 77(2): 221-232.
- Kruger, J., and D. Dunning. 1999. Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology* 77(6): 1121-1134.
- Kruse, D. L. 1992. Profit sharing and productivity: Microeconomic evidence from the United States. *Economic Journal* 102 (410): 24-36.
- Kruse, D. L. 1993. Profit sharing: Does it make a difference? : The productivity and stability effects of employee profit-sharing plans. *Kalamazoo, Michigan: W.E. Upjohn Institute for Employment Research.*
- Landy, F.L., J.L. Barnes, and K.R. Murphy. 1978. Correlates of perceived fairness and accuracy of performance evaluation. *Journal of Applied Psychology* 63: 751-754.
- Lawler, E. E. 1973. *Motivation in Work Organizations*. Monterey, CA: Brooks/Cole.
- Lawler, E. E., and S. A. Mohrman. 2003. *Creating a strategic human resources organization: An assessment of trends and new directions*. Palo Alto: Stanford University Press.
- Lawler, E. E. and L. W. Porter. 1967. The effect of performance on job satisfaction. *Industrial Relations* 7: 20-28.
- Lazear, E. P. 1992. The Job as a Concept. *Performance Measurement, Evaluations, and Incentives* W. Bruns. Boston, MA: Harvard University Press: 183–215.
- Lazear, E. P. 2004. The Peter Principle: A Theory of Decline. *Journal of Political Economy* 112: 141–63.
- Lazear, E. P. and S. Rosen. 1981. Rank-Order Tournaments as Optimal Labor Contracts. *Journal of Political Economy* 89: 841-864.
- Lipe, M. S. and S. E. Salterio. 2000. The balanced scorecard: judgmental effects of common and unique performance measures, *The Accounting Review* 75: 283–298.
- Longenecker, C. O., H. P. Sims and D. A. Gioia. 1987. Behind the mask: The politic of employee appraisal. *The Academy of Management Executive* 1 (3): 183-193.
- Ma, C. 1988. Unique Implementation of Incentive Contracts with Many Agents. *Review of Economic Studies* 55(4): 555-572.

- Matsumura, E.M., and J.Y. Shin. 2006. An Empirical Analysis of an Incentive Plan with Relative Performance Measures: Evidence from a Postal Service. *The Accounting Review* 81 (3): 533-566.
- Medoff, J. L., and K. G. Abraham. 1980. Experience, Performance, and Earnings. *Quarterly Journal of Economics* 95: 703–36.
- Milgrom, P. 1988. Employment contracts, influence activities, and efficient organization design. *Journal of Political Economy* 96(1): 42-60.
- Milgrom, P. and J. Roberts. 1992. *Economics, Organization and Management*. Englewood Cliffs, NJ: Prentice-Hall International.
- Moers, F., 2005. Discretion and bias in performance evaluation: the impact of diversity and subjectivity. *Accounting, Organizations and Society* 30: 67–80.
- Moorman, R. H. 1991. Relationship between organizational justice and organizational citizenship behaviors: Do fairness perceptions influence employee citizenship? *Journal of Applied Psychology* 76: 845–855.
- Moorman, R. H., B. P. Niehoff, and D. W. Organ. 1993. Treating employees fairly and organizational citizenship behavior: Sorting the effects of job satisfaction, organizational commitment, and procedural justice. *Employee Responsibilities and Rights Journal* 6: 209–225.
- Motowidlo, S. J., and W. C. Borman. 1977. Behaviorally anchored scales for measuring morale in military units. *Journal of Applied Psychology* 62 (2): 177–183.
- Murphy, K., and P. Oyer. 2001. Discretion in executive incentive contracts: Theory and evidence. Working paper. *University of Southern California and Stanford University*.
- Murphy, K. R. and J. Cleveland. 1995. *Understanding Performance Appraisal: Social, Organizational, and Goal-based Perspectives*: Sage Publications.
- Murphy, W. H. and R. S. Sohi. 1995. Salespersons perceptions about sales contests: Towards a greater understanding. *European Journal of Marketing* 29 (13): 42-66.
- Nalebuff, B., and J. Stiglitz. 1983. Prizes and incentives: Towards a general theory of Compensation and Competition. *Bell Journal of Economics* 14 (1): 21–43.
- Nordstrom, R., P. Lorenzi, and R. Hall. 1990. A Review of public posting of performance feedback. *Journal of Organizational Behavior Management* 11: 101–123.
- O’Connell, V. O. 2008. Retailers reprogram workers in efficiency push. *Wall Street Journal* (September 10): B1.

- Parrino, R. 1997. CEO turnover and outside succession: A cross-sectional analysis. *Journal of Financial Economics* 46: 165–197.
- Peter, L. J., and R. Hull. 1969. *The Peter Principle: Why Things Always Go Wrong*. New York: Morrow.
- Pillutla, M. M., and J. K. Murnighan. 1995. Being fair or appearing fair: strategic behavior in ultimatum bargaining. *Academy of Management Journal* 38: 1408-1426.
- Prendergast, C. 1999. The Provision of Incentives within Firms. *Journal of Economic Literature* 31: 7-63.
- Prendergast, C. 1993. The role of promotion in inducing specific human capital acquisition. *Quarterly Journal of Economics* 108(2): 523.
- Prendergast, C. and R. H. Topel. 1993. Discretion and bias in performance evaluation. *European Economic Review* 37(2-3): 355-365.
- Prendergast, C. and R. H. Topel. 1990. Favoritism in organizations. *Journal of Political Economy* 104(5): 958-978.
- Rajan, M.V. and S. Reichelstein. 2009. Objective versus subjective indicators of managerial performance. *The Accounting Review* 84: 209-237.
- Ray, K. 2007. The Retention Effect of Withholding Performance Information. *The Accounting Review* 82(2): 389-425.
- Saal, F.E., and F.J. Landy. 1977. The Mixed Standard Rating Scale: An Evaluation. *Organizational Behavior and Human Performance* 18: 19-35.
- Sims, H. P., A. D. Szilagyi, and R. D. McKemey. 1976. Antecedents of Work-Related Expectancies. *Academy of Management Journal* 19 (4): 547–59.
- Suls, J., and L. Wheeler. 2000. *Handbook of Social Comparison: Theory and Research*. New York, NY: Kluwer Academic/Plenum Publishers.
- Tafkov, I. 2013. Private and public relative performance information under different compensation contracts. *The Accounting Review* 88 (1): 327–50.
- Towry, K. L. 2003. Control in a teamwork environment: The impact of social ties on the effectiveness of mutual monitoring contracts. *The Accounting Review* 78 (4): 1069-1095.
- Varma, A., A. S. DeNisi and L. H. Peters. 1996. Interpersonal affect and performance appraisals: A field study. *Personnel Psychology* 49(2): 341-360.
- Vroom, V. H. 1964. *Work and motivation*. San Francisco, CA: Jossey-Bass.

- Waldman M. 1984. Job assignments, signaling and efficiency. *Rand Journal of Economics* 15: 255–270.
- Weiss, A. 1987. Incentives and Worker Behavior. In *Information, Incentives and Risk Sharing*. Haig Nalbantian, ed. Towota, N.J.: Rowan and Littlefield.
- Weitzman, M. L. 1980. Efficient Incentive Contracts. *Quarterly Journal of Economics* 94: 719-730.
- Weitzman, M. L., and D. L. Kruse. 1990. Profit sharing and productivity. In *Paying for Productivity*, ed. AS Blinder, pp. 95–140. Washington, DC: Brookings Inst.
- Wikoff, M., D. Anderson, and C. Crowell. 1982. Behavior management in a factory setting: Increasing work efficiency. *Journal of Organizational Behavior Management* 4: 97–127.
- Zhang, M. 2008. The effect of perceived fairness and communication on honesty and collusion in a multi-agent setting. *The Accounting Review* 83(4): 1069-1095.