Job Allocation Rules and Sorting Efficiency:
Experimental Outcomes in a Peter Principle Environment

David L. Dickinson

Marie Claire Villeval

ABSTRACT

An important issue in personnel economics is the design of efficient job allocation rules. Firms often use promotions both to sort workers across jobs and to provide them with incentives. However, the Peter Principle states that employees' output tends to fall after a promotion. Lazear (2004) suggests that self-selection may improve job allocation efficiency while preserving incentive effects. We reproduce this Peter Principle in the laboratory and compare the efficiency of a promotion standard with subjects self-selecting their task. We find no evidence of effort distortion, as predicted by theory. Furthermore, we find that when the Peter Principle is not severe, promotion rules often dominate self-selection efficiency of task assignment. Results are consistent with imperfect appraisal of transitory ability and a lack of strategic behavior.

Dickinson--Corresponding author: Appalachian State University, Dept. of Economics, Boone, NC 28608, U.S.A.. Email: dickinsondl@appstate.edu

Villeval: University of Lyon, Lyon, F-69007, France; CNRS, GATE, 93 chemin des Mouilles, Ecully, F-69130, France; IZA, Bonn (Germany); CCP, Aarhus (Denmark). E-mail: villeval@gate.cnrs.fr

1. Introduction

The productivity of any firm is affected by the efficiency of its job allocation rules. Promotion policies allocate jobs while also rewarding workers, and this dual role may generate conflicts that reduce the efficiency of organizations (Milgrom and Roberts, 1992). Specifically, promotions may allocate workers to jobs for which they are not necessarily the best suited. This gives rise to the Peter Principle that states that in any hierarchical organization, employees are (inefficiently) promoted to their level of incompetence (Peter and Hull, 1969).

Lazear (2004) offers a statistically-based explanation of the Peter Principle. He theoretically shows that the Peter Principle is a necessary consequence of any promotion rule when observable outcomes are a function of both a permanent and transitory (i.e., random) ability component.[1] When job-specific transitory ability has expected value equal to zero, those promoted most likely had a favorable "random draw" prior to promotion. A regression-to-the-mean effect therefore causes the Peter Principle in the Lazear framework. Unless employers take this effect into account, the result will be job allocation (and organizational) inefficiency.[2] This suggests a type of adverse selection, except that with the Peter Principle even the agents themselves (i.e., workers) can be worse off due to job misallocation if compensation remains performance-based. In this paper, we focus on the question of worker or agent behavior in the face of this potential job sorting inefficiency. Job misallocation can reduce worker payoffs, and so we explore whether workers are more efficient at self-selecting their task compared to when a promotion rule is used—efficient self selection may, however, be limited if workers are overconfident or underconfident in their own skills. We also explore, as suggested in Lazear (2004), whether low (high) ability employees may strategically distort their effort downward (upward) pre-promotion in order to ensure efficient task assignment.

Peter and Hull (1968) and Lazear (2004) provide a variety of real-world examples of the Peter Principle: the decline of faculty productivity after tenure, the lower quality of movie sequels, a relatively less satisfying second experience at a new restaurant—each of these examples are more likely if a random or luck component contributes to the first stage outcome of interest. Empirical tests of the Peter Principle are scarce, however. Using data from a large hierarchical British financial sector firm, Barmby et al. (2010) find some evidence that performance declines after promotion. Two-thirds of the fall in post-promotion performance is attributed to the Peter Principle and one-third to reduced incentives. Using job transition data from the personnel records of a large U.S. corporation, Acosta (2010) identifies the existence of the Peter Principle by analyzing the difference in promotion dynamics between incumbents and new hires, namely the fact that after several promotions incumbents are less likely to be promoted in comparison with new hires.

The relative lack of empirical tests of the Peter Principle is largely due to difficulties with selection bias and measurement of "luck".[3] In naturally occurring environments, lower-level positions may have higher turnover (e.g., non-promoted workers may have higher quit rates), promotions may not be based solely on performance (e.g., discrimination), and jobs may not always produce easily measurable individual outcomes (e.g., work team production). Any of these will imply confounds in interpreting output data across hierarchical levels, resulting in potentially biased estimates of the Peter Principle. The laboratory offers a distinct advantage in this instance because all of these concerns are removed by virtue of the experimental environment. The aim of this paper is to generate a Peter Principle environment in the laboratory using a real-effort task and then to examine the robustness of the Peter Principle, the existence of

effort distortion when agents have an incentive to be efficiently assigned to a job, and the relative efficiency of promotions versus task self-selection.

Our laboratory environment successfully generates the Peter Principle, though only when the variance of the transitory component is quite high. This is consistent with Lazear (2004), and it also validates our experimental design. Allowing employees to self-select their task is predicted to be efficient, but this assumes perfect knowledge of one's abilities. We find some evidence of overconfidence regarding one's task outcome, which would induce agents to inefficiently self-select into the hard task (see Camerer and Lovallo, 1999, on overconfidence inducing excess entry). As a result of such overconfidence, promotion rules may dominate self-selection in terms of efficiency as a job allocation device, at least in the absence of a severe Peter Principle. Lastly, we find no evidence of effort distortion, contrary to theoretical predictions. This is consistent with the hypothesis that agents are naïve rather than strategic in this environment.

2. Theory

In this section, we present briefly the main characteristics and predictions of Lazear's (2004) model. The firm's productive activity consists of an easy (E) and a hard (H) task. Every agent's ability consists of a permanent component, A, and a transitory or random component in period t, ε_t, that is *i.i.d.* over time. The principal cannot observe an agent's permanent ability, but only $\hat{A} = A + \varepsilon_t$.[4] The agent's productivity, which determines her pay, on the easy task is: $Y^E = \alpha + \beta(A + \varepsilon_t)$ and on the hard task is: $Y^H = \gamma + \delta(A + \varepsilon_t)$, with $\alpha > \gamma$ and $\beta < \delta$. Thus, a less (more) able agent has a comparative advantage at the easy (hard) task. Where these two productivity profiles cross corresponds to the ability level at which the employer and worker are

both indifferent between being promoted or not, assuming incentive compatible piece-rate pay for workers (see Figure 1).

The model proceeds in two stages. In the first stage of the game, all agents are assigned to the easy task. At the end of the first stage, the principal only observes $\hat{A}$ and must decide on each agent's task allocation for the second stage of the game. The principal determines a standard, A^*, such that if $\hat{A} \geq A^*$ the agent is promoted to the hard task, and if $\hat{A} < A^*$ he again performs the easy task in the second stage. Note that within this framework, a "promotion" to the hard task is considered desirable only by high ability workers, as they have higher expected earnings in the hard task compared to the easy task. The model does not consider many other factors connected to real world promotions (e.g., status, power, responsibility, etc), but rather it focuses just on how worker ability determines how promotions assignments should be made.

Lazear's model generates the prediction that, on average, the agents who are promoted have a lower observable outcome after promotion—the Peter Principle. The prediction results entirely from the fact that observable outcomes are, in part, a result of a random error draw, ε_t. In more formal terms, the conditional expected value of ε_1 for those agents who have been promoted is given by: $E\left(\varepsilon_1 | A + \varepsilon_1 > A^*\right) = \int_{-\infty}^{\infty} E\left(\varepsilon | \varepsilon > A^* - A\right) f(A) dA$, which is greater than zero. In contrast, the expected value of ε_2 is zero. As a consequence, there is a predicted decrease in the expected observed ability of the promoted agents:

$A + E\left(\varepsilon_1 | A + \varepsilon_1 > A^*\right) > A + E\left(\varepsilon_2 | A + \varepsilon_1 > A^*\right)$ (see details in Lazear (2004)). An intuitive prediction that follows from the model is that the Peter Principle will be stronger when the transitory ability component is high relative to one's permanent ability component.[5] Thus, a larger variance to ε_t will give rise to a stronger Peter Principle.

An extension of the model introduces strategic agent behavior. When agents have perfect information on their permanent ability, they can manipulate their effort in stage 1 in order to influence their assignment to a specific task. Let e_E and e_H denote the effort levels chosen by an agent on the easy and hard tasks, respectively, with $c(e)$ the associated cost incurred. Since the expected value of the transitory component is 0, agents prefer the hard task if the following holds: $\alpha + \beta(A + e_E) - c(e_E) < \gamma + \delta(A + e_H) - c(e_H)$. The testable implication is that, under a promotion rule, an agent will strategically distort his level of effort downward (upward) in stage 1 if his permanent ability is low (high). Intuitively, the agent distorts effort to increase the likelihood of assignment to the task in which he has the comparative advantage, because mis-assignment is costly to the worker when pay is based on output. The incentive for such effort distortion does not exist in stage 2 after promotions assignments have been made.

To sum up, Lazear (2004) delivers several testable implications that we examine:

a) The larger the variability of the random component relative to the permanent one, the stronger the Peter Principle (e.g., the larger the decrease in average post-promotion output).
b) If agents know their ability, allowing agents to self-select their second-stage task will be more efficient than a promotion rule.
c) If agents know their permanent ability, then effort will be distorted downward (upward) by lower (higher) ability agents prior to a promotion stage (i.e., in stage 1).

3. Experimental Design

Our experiment is designed to generate the Peter Principal as a statistical effect, as in Lazear (2004). An initial objective is to successfully produce the Peter Principle in the lab (i.e., validate the design) so that self-selection and effort distortion hypotheses can be meaningfully examined. We simplify the lab environment by focusing only on agent (i.e., worker) behavior,

and so subjects are all assigned as agents.[6] In promotion rule treatments, we use an exogenous standard for promotions.

The Task

The cognitive task utilized is a modified version of the task in Baddeley (1968). The easy task involves reading a short sentence that describes the order of three letters. The subjects have to type in these three letters in the order implied by the sentence. For example, "*S follows B, is not preceded by G*" implies the order "*BSG*".[7] The hard task is similar except that each sentence contains five letters. For example, "*S follows B, is not preceded by G, precedes H, is not followed by J*" implies the ordering "*JBSGH*". Sentences are randomly generated from all letters of the alphabet. Once typed, the sentence is auto-checked against the correct answer, so that subjects can learn and improve more rapidly. Though Baddeley (1968) indicates only minor learning with repetition of the task, we consider learning as a potential confound in the data and implement a practice stage of both easy and hard tasks prior to the two main treatment stages (i.e., the pre- and post-promotion stages). The advantage of this task is that no specific knowledge is needed, difficulty can be increased by increasing the number of letters without altering the cognitive focus of the task, and the outcome is easily measurable. Also, effort is distinguished from ability by use of paid practice stages that contain no incentives to distort effort, because promotions are not a part of the practice stages.

Game Stages and Parameter Values

A treatment consists of two stages. In Stage 1, the subject performs the 3-letter task for 7 minutes. The number of correctly solved sentences in Stage 1 is our proxy for the subject's

permanent ability, A. At the end of Stage 1, the subject draws an *i.i.d* random number from a uniform distribution, the variance of which will depend on the treatment as described below. This random number corresponds to the task-specific transitory component in the theoretical model. The subject's "observable" outcome in Stage 1—the outcome on which payment is based—is the sum of the number of all correctly solved sentences and the random draw number.

In Stage 2, some subjects are assigned the easy task and others the hard task. Assignments for Stage 2 are made in one of three ways, depending on the treatment. In the *Calibration* treatment, the Stage 2 task assignment is random: half of the subjects are assigned to each task. In the *Promotion* treatments, subjects whose observable outcome in Stage 1 surpasses an exogenous standard are assigned to the hard task in Stage 2, the other subjects perform the easy task in Stage 2. The *Promotion* treatments will allow us to recreate the Peter Principle as a statistical effect that varies with the strength (i.e., the variance) of the transitory term, and they will also provide one way to examine effort distortion. Thus, the *Promotion* treatments are used to examine testable implications (a) and (c). In the *Self-Selection* treatment, subjects are allowed to choose their Stage 2 task after having completed the easy task in Stage 1. Comparing outcomes in the *Self-Selection* and *Promotion* treatments will allow us to compare relative efficiency of job allocation rules to examine testable implication (b).

The specifics of the error-term distributions used for the transitory component of the observable outcomes were as follows. In the *Calibration, Self-Selection,* and the baseline *Promotion* treatment (called *Promotion-σ_{low}*), a random draw from a uniform $U[-12,+12]$ distribution is added to the number of correct sentences solved in Stage 1 to arrive at the observable outcome. Two additional *Promotion* treatments, *Promotion-σ_{Medium}* and *Promotion-σ_{High}*, increase the variance of the random term distribution for the easy task to $U[-24,+24]$ and

$U[-75,+75]$, respectively. Subjects who perform the hard task in Stage 2 solve 5-letter sentences

during a 7-minute period and then draw an *i.i.d.* random number from a uniform $U[-4,+4]$ for the

Calibration, *Self-Selection*, and *Promotion-σ_{low}* treatments. The distributions used for

Promotion-σ_{Medium} and *Promotion-σ_{High}* for the hard task are $U[-8,+8]$, and $U[-25,+25]$,

respectively.

The reason the support of the easy task error term distributions is 3 times as wide as that

of the hard task within the same treatment is that we experimentally-induce a ratio of one hard

task to three easy tasks for compensation equivalence. That is, one hard sentence solved is paid

a piece rate three times higher than the piece rate for each easy task solved, which corresponds to

$\delta = 3\beta$ for the agent production functions in the Lazear framework. Accordingly, the variance

of the random component in the easy task is 3 times larger than in the hard task. By

manipulating the piece-rates so that the hard task is rewarded at a rate *greater* than the subjects'

observed opportunity cost of the hard task, which was determined to be 2 units of the easy task

from pilot sessions (n=12 subjects), we induce wage profiles for each task that differ in slope.

Figure 1 is adapted from Lazear (2004), and highlights the final parameterization for our

treatments.[8] The vertical axis measures payoff points (see Instructions in Appendix), which are

directly related to the observable output. Fixed payments are added to the piece rate

compensation for both easy and hard tasks, with a higher fixed payment for the easy task (108

points) than for the hard task (9 points), so that we induced the desired overlapping wage profiles

shown in the Figure. Results from Stages 1 and 2 of the pilot sessions and *Calibration* treatment

generated the average ability estimate of approximately 60-70 units of the easy task and 30-35

units of the hard task—thus, we consider the hard task opportunity cost to be about 2 units of the

easy task, as noted above. The final fixed payments we utilized with each wage profile create a

crossing point at ability=68 units of easy task (equivalent to 34 units of hard task output), as shown in Figure 1. Thus, we use the standard of easy task output=68 for the *Promotion* treatments. Though it is not critical in our design to have 50% promotions in the *Promotion* treatments, 68 units of the easy task was likely a bit above average ability, and so we anticipated a promotion rate of just under 50% .[9, 10]

Finally, we note that in the *Calibration* and *Promotion* treatments, a hypothetical response of preferred Stage 2 task is elicited at the beginning of this stage. It was made clear to subjects that their task had already been determined and the hypothetical response would have no bearing on their Stage 2 task. Nevertheless, this response provides an indication of whether subjects correctly assessed their abilities, or whether they are overconfident. Though Lazear (2004) does not consider it, others have shown that overconfidence with respect to one's abilities may generate over-entry in hierarchical organizations (Camerer and Lovallo, 1999), which is possible in the *Self-Selection* treatment. Because optimism is uncontrolled in our design, this hypothetical self-assessment of abilities will at least allow us to further evaluate self-selection efficiency. Important treatment details are shown in Table 1.

Experimental Procedures

Ten sessions were conducted at the Groupe d'Analyse et de Théorie Economique (GATE), Lyon, France, with 2 sessions per treatment. 190 undergraduate subjects were recruited from local Engineering and Business schools using the ORSEE software (Greiner, 2004). 81.05% of the subjects had never participated in an experiment. On average, a session lasted 60 minutes. The experiment was computerized using the REGATE program (Zeiliger, 2000). Upon arrival, the subjects were randomly assigned to a computer terminal, and experimental

instructions (see Appendix) were distributed and read aloud. Questions were answered privately.

In order to limit the learning effect confounds across the two main stages of the experiment, subjects had first to practice both the easy and the hard tasks prior to commencement of real experiment Stages 1 and 2.[11] Subjects practiced first the easy task with 3 sentences without payment until they were correctly solved, after which point a 7-minute paid easy task started (i.e., easy task practice). A similar hard task practice stage followed distribution of instructions for that task. No random component was added to one's real output in the practice stages, but subjects were compensated in the practice rounds similar to in the main experiment.[12] Following easy and hard task practice, the instructions for both Stages 1 and 2 were distributed together and read aloud. Knowing the rules for task assignment in Stage 2 before starting Stage 1 allowed subjects to behave strategically in Stage 1, if desired, in the *Promotion* treatments (i.e. distorting their effort upwards (downwards) if believing they were (not) able enough to make more money performing the hard task in Stage 2). After Stage 1, subjects were given a brief 3-minute break, with no communication allowed. We used a conversion rate of 100 points = €1 for payments. At the end of each session, payoffs were totaled for the practice and main stages. In addition, the participants received €4 as a show-up fee. The average payoff was €15.40, and subjects were each paid in private in a separate room.

4. Results

The Peter Principle

Table 2 shows average ability levels across all treatments (including the pilot sessions) in practice stages as well as Stages 1 and 2. Recall that these data do *not* include the transitory term. They indicate that some learning occurs between the practice stages and Stages 1 and 2, which was the purpose of utilizing practices stages. A comparison of Stage 1 outcomes with

those subjects performing the easy task again in Stage 2 indicates that such learning is not a

concern after the practice stages (though effort distortion is a distinct concern we address below).

Comparing the Stage 1 and 2 final outcome levels—these outcomes *include* the error

term component in the outcome—in the *Promotion* treatments in Table 3 will indicate whether

we successfully produce the Peter Principle. In order to compare outcomes across easy and hard

tasks, outcome levels in the hard task are normalized to the easy task metric in the following

way: we multiply the hard task score by 2 (using the conversion rate from our pilot sessions) and

the random component by 3 (using the piece-rate conversion).[13]

Table 3 shows that in *Promotion-σ_{High}*, the post-promotion average outcome of the

promoted subjects is 43.5% lower than the pre-promotion level (*p*=.061 Wilcoxon signed-rank

test). This treatment provides the cleanest evidence of the Peter Principle. In the other

promotion treatments, outcome differences are not statistically significant (*p*>0.10) pre- and

post-promotion. The Peter Principle also implies that observable output of *non*-promoted

subjects should increase in Stage 2, as our data indicate. The increase is 11.90%, 16.99%, and

161.27%, respectively, in *Promotion-σ_{Low}, Promotion-σ_{Medium}*, and *Promotion-σ_{High}* . These

differences are all statistically significant (*p*=0.091, *p*=0.073, and *p*=0.001, respectively). Figure

2 shows that luck plays the intended role in our experimental environment, as those promoted to

the hard task had error term draws significantly higher in Stage 1. In *Promotion-σ_{Low}*, 72% of

the promoted subjects had a positive transitory component in Stage 1, compared to only 42% of

the non-promoted subjects (Mann-Whitney, *p*=.079). In *Promotion-σ_{Medium}* , the respective

percentages are 74% and 33% (*p*=0.014) and in *Promotion-σ_{High}* , 80% and 9% (*p*<0.001).

Figure 2 confirms that in most *Promotion* treatments, the promoted subjects (labeled

'Hard 2') have been luckier than the non-promoted ones, whereas in the *Calibration* treatment

the difference is not significant. The average Stage 1 transitory ability of promoted agents is

positive and it increases in its variance. According to Mann-Whitney tests, the differences in

average transitory components between promoted and non-promoted subjects are significant in

Promotion-σ_{Low} (*p*=0.011) and *Promotion-*σ_{High} (*p*<0.001), while marginally insignificant in

Promotion-σ_{Medium} (*p*=0.120). Table 4 presents Stage 1 abilities (i.e., no transitory term added)

and highlights mistaken promotions of those with Stage 1 ability less than the breakeven ability

level of 68 (i.e., Fig. 2 cutoff point). Note that correct (or efficient) assignment to the Stage 2

task in the far-right column in Table 4 implies that a subject is performing the task with highest

expected earnings given the subject's revealed ability in Stage 1. Another source of inefficiency

shown in Table 4 is mistaken non-promotions of those with Stage 1 ability greater than 68.

As a whole, the *Promotion* treatments data (Tables 3, 4, Figure 2) result from a

framework created to generate the principle, so these results are an important validation of a

rather complex experimental design. We can therefore explore comparison institutions and their

predicted effects on effort distortion and self-selection—two key outcome variables not

amenable to simulation.

Self-selection vs. Promotions

Our *Self-Selection* treatment allows us to test the hypothesis that one can avoid the Peter

Principle by letting the individuals select their task. In *Self-Selection*, subjects who promoted

themselves to the hard task in Stage 2 (see "promoted subjects" columns in Table 3) *increased*

normalized outcomes by a significant 18.44% relative to Stage 1 (*p*<0.001), due partly to a

higher rate of negative transitory components in Stage 1 (Fig. 2). In contrast, those who chose

the easy task only increased their outcome by an insignificant 1.40% (Table 3, *p*>0.10). Though

this result contrasts with the clearly observed Peter Principle in *Promotion-σ_{High}*, recall that the error term distribution in *Self-Selection* is comparable to the distribution in *Promotion-σ_{Low}*, which did not show robust evidence for the Peter Principle. Thus, further analysis is required to shed light on whether we find evidence of any significant behavioral difference between these *Promotion* and *Self-Selection* allocation rules.

The lack of a Peter Principle in the *Self-Selection* treatment does not imply that decisions in that treatment are efficient. Table 4 indicates that, in *Self-Selection*, 22.50% promoted themselves in Stage 2 though their ability was below 68—mistake 1—and 15.00% of the subjects made the opposite mistake 2. The respective proportions of mistakes were 5.41% and 2.70% in *Promotion-σ_{Low}*, which has a comparable transitory variance, and the differences are significant (Mann-Whitney tests, *p*=0.031 and *p*=0.017, respectively).

In our *Promotion* treatments, we asked subjects to make a hypothetical choice at the end of Stage 1 about which task they would prefer in Stage 2. Comparing these choices with the actual scores in Stage 1 allows us to determine whether the subjects were biased in their self-assessment or not, which may help explain inefficiencies in *Self-Selection*. In Table 5, "optimism" ("pessimism") indicates the proportion of subjects who would choose the hard (easy) task although their score in Stage 1 was below (equal or above) 68; "accurate" corresponds to the proportion of subjects who were unbiased and would have efficiently self-selected. In this Table, the data from the *Self-Selection* treatment correspond to the actual subjects' choices.

Table 5 shows that a relatively high proportion of subjects overestimate their ability to succeed at the hard task in Stage 2, consistent with non-hypothetical choices in *Self-Selection*. Though Table 5 shows evidence that some job matching errors are due to pessimism, the data indicate that optimism with respect to the transitory component draw is more pervasive. Because

hypothetical accuracy in *Promotion-σ$_{Low}$* is greater than actual accuracy in *Self-Selection* (but with same transitory component variance), we do not believe that mistakes are due to lack of financial incentives for the hypothetical question in the *Promotion* treatments.[14]

Another possibility is that subjects overestimate the extent to which they may improve due to learning from Stages 1 to 2. If learning occurs with the letter-sorting task, then a subject might score less than 68 in Stage 1 but *still* prefer the Hard task if anticipating improved productivity due to learning, ceteris paribus. However, data from the Calibration treatment indicate that, as we expected, learning is not a serious concern with this task, and it accounts for score improvements of less than 10% (see Table 2). Additionally, the data show optimism grows with the variance of the transitory component (Table 5), which also seems to indicate that optimism with respect to the random component has a bigger influence on task assignment inefficiency than does overconfidence with respect to one's ability.

A parametric approach to investigate the hypothesis that optimism influences self-selection and hypothetical choice errors is also explored. We estimated two Probit models in which the dependent variable equals 1 if the actual or hypothetical choice for Stage 2 is the hard task. In both models, the explanatory variables include the subject's score and transitory ability in Stage 1. In the hypothetical choice model (i.e., *Promotion* treatments), we control for the variance of the transitory ability with a variable equal to 0, 1, or 2 for *Promotion-σ$_{Low}$*, *Promotion-σ$_{Medium}$*, and *Promotion-σ$_{High}$*, respectively. Table 6 reports the results of these estimations.

Not surprisingly, Table 6 indicates that both the actual and hypothetical choices of the hard task increase as a function of one's Stage 1 score. It also shows that bad luck in Stage 1 increases the likelihood of choosing the hard task in Stage 2, even though it is common

knowledge that the transitory component draws are independent across stages. These results may indicate that the unlucky subjects in Stage 1 expect to be lucky in Stage 2. Though subject seem to take into account the mean-reversion of the error term, the *Promotion* model confirms that optimism is independently related to the transitory ability variance—a larger variance breeds optimism.[15]

When is self-selection likely to be more efficient than using a promotion standard, given that the theoretical assumption of perfect knowledge of one's own abilities may be violated? If we compare the values in Table 5 to the efficiency of actual task assignments from Table 4, we find that the promotion standard performs better than the hypothetical choices when the variance of the transitory ability is low (92% vs. 73%) or medium (82% vs. 55%). In contrast, when this variance is high, the hypothetical choices of the subjects are accurate in 59% of the cases whereas only 43% of the actual promotion decisions are accurate (Wilcoxon test: *p*=0.010); the 54% efficiency rate with simulated assignment was also lower than with hypothetical choices. These data lead us to conclude that promotions are more efficient than task self-selection when the Peter Principal is quite small.

Effort Distortion

Lazear (2004) hypothesizes that promotions may induce effort distortion by the agents in anticipation of the promotion. This hypothesis requires the assumption that agents know (but principals do not know) their permanent ability and have unbiased expectations regarding their transitory ability component. With these assumptions, agents would know their preferred Stage 2 task assignment. Because the transitory component of ability affects promotion decisions, high (low) ability agents have an incentive to distort effort upward (downward) to ensure efficient assignment to the hard (easy) task.

From the experimental data, we take several approaches to examining this hypothesis. If effort is distorted in Stage 1, then Stage 2 should reflect a non-distorted effort. Low ability agents should have lower productivity scores (i.e., without considering the error term) in Stage 1 of the easy task than Stage 2, while high ability agents who are promoted should have Stage 1 easy task scores that are higher than their normalized hard task scores in Stage 2. For the subset of non-promoted subjects, the score is on average 9.16 % higher in Stage 2 in *Promotion-σ_{Low}* (Wilcoxon test: $p=0.028$), 8.47 % higher in *Promotion-σ_{Medium}* ($p=0.011$), and 7.79 % in *Promotion-σ_{high}* ($p=0.004$). These differences are likely a small amount of residual learning, however, because in *Self-Selection* the non-promoted subjects still performed 8.27 % better in Stage 2 (Wilcoxon test: $p=0.001$) in the absence of any incentive to distort effort. Thus, the entire increase in output in Stage 2 of low ability subjects is attributable to residual learning. For high ability subjects, similar Stage 1 to 2 trends occur in the *Promotion* and *Self-Selection* treatments, so no evidence exists for effort distortion in high ability agents either.

An alternative approach to examining effort distortion is to compare the variance of scores (i.e., non-transitory ability) in Stage 1 in the *Promotion* and *Self-Selection* treatments. Effort distortion should imply a higher variance of scores in Stage 1 as low ability subjects distort effort down but high ability subjects distort upward. We find that the standard deviation of scores is 15.83 in the *Self-Selection* treatment and 13.55 in the *Promotion-σ_{Low}* treatment. Contrary to what should be expected from effort distortion, the variance of effort is therefore is smaller, and not larger, when a promotion standard is used. Yet another approach would be to examine within-treatment effects utilizing the practice stages. Effort distortion in Stage 1 of the *Promotion* treatments should lead to a higher variance in scores than in the practice stage of the easy task of that treatment, where strategic motivation is not present. We examine these

practice-to-real stage differences in standard deviations and fail to reject the null hypothesis that the variance differences are zero in any of the *Promotion* treatments ($p>.10$, variance ratio tests).[16]

In sum, we do not find support for the hypothesis of strategic effort distortion. It is possible that effort distortion may not have the same expected payoff as in Lazear (2004) if expectations are biased with respect to the transitory ability component. That is, if low ability agents are pessimistic, thinking they will receive a negative transitory ability component, they do need not distort effort downwards to insure assignment to their optimal easy task in the last stage. This interpretation also implies that if high ability subjects are optimistic with respect to the transitory component (i.e., expecting a positive draw), then the expected marginal benefit of distorting effort upwards in *Promotion* treatments is smaller. A comparison of mean abilities in Stage 1 for *Promotion* treatment subjects who hypothetically chose the hard over the easy task finds significant ability differences (71 versus 58 respectively—$p<.001$ for the t-test of mean differences). This indicates that subjects who would prefer to be assigned the hard (easy) task were, on average, higher (lower) ability subjects, and not merely optimistic (pessimistic) with regards to the transitory component draw. Thus, subjects may have some bias with respect to their transitory ability component, but they mostly appear to utilize the information they possess with respect to their own abilities in a rational way. Nevertheless, because effort distortion requires some level of strategic thinking, and subjects do not appear to distort effort, our results indicate that agents are mostly rational but perhaps naïve. A naïve agent may be capable of understanding that a higher ability increases the odds that the hard task is a more efficient task assignment, but may not think strategically enough to distort effort to guarantee assignment to the task with higher expected payoff.

5. Conclusion

When transitory ability components affect observable outcomes, the Peter Principle can occur, and post-promotion performance declines. Lazear (2004) provides a theoretical framework in which the Peter Principle is interpreted as a regression-to-the-mean phenomenon, and the relative size of the transitory component determines the severity of the Peter Principal. Two important implications follow with respect to full-information unbiased agent (employee) behavior. First, promotion decisions based on an exogenous cutoff point are less efficient than if agents self-select their jobs. Secondly, agents will distort effort in anticipation of promotions to ensure efficient assignment to the Stage 2 task.

Our experiment recreates the essential features of Lazear's (2004) framework. In particular, the permanent ability component is captured by subject effort, while a random error term simulates transitory ability. Our design is validated by generation of a laboratory Peter Principal. Contrary to Lazear's hypothesis, we do not find that self-selection of one's task generates fully efficient job assignment. This is likely due to biased beliefs related to the transitory ability component. Van den Steen (2004) highlights how a random component to outcomes may give rise to systematic biases that could also affect effort. While we find some evidence of optimism/pessimism with respect to the random component (and therefore some evidence of task self-selection errors), we do not find that these biased beliefs are systematically related to effort. In addition, when the variance of the transitory ability component is large enough, the analysis of hypothetical choices suggests that self-selection would be efficiency-improving compared to the use of an exogenous standard, assuming incentive-compatible wages.

We also fail to find support for the effort distortion hypothesis. Again, this may result from biased beliefs with regards to transitory ability components, or a lack of strategic behavior

(i.e., naïve agents). Throughout, we have implicitly assumed risk neutrality of agents. If subjects are risk averse this would only increase the incentives to distort effort in the *Promotion* treatment. The possibility of inefficient task assignment in Stage 2 is the relevant risk that could influence the agent choice of effort. Our lack of evidence for effort distortion does not necessarily imply agents are not risk averse, but if they are then agent optimism apparently counteracts such risk aversion in subjective expected payoff assessments.

Our main result suggests that the Peter Principle is only of concern when the transitory component of one's outcome is relatively severe, such as in highly volatile environments or in activities where the role of luck may be large compared to talent. The managerial implications are that simple promotion rules based on worker outcomes can be an efficient means of job sorting in most cases. However, as noted in Lazear (2004), relatively large transitory components to outcomes should be taken into account by employers. In such instances, promotions should be based on more comprehensive measurements or require longer term observation to more accurately assess worker ability. Such comprehensive assessments or longer probationary periods would be costly, but employers would benefit by reducing job misallocation (i.e., inefficiency).

This research offers avenues for further study. In Lazear (2004), employers adjust promotion standards in anticipation of Peter Principle effects. Future research could therefore analyze the principals' behavior regarding the optimal adjustment of the promotion standards. It is also the case that different tasks could be subject to a more or less severe optimism bias in one's self-assessed abilities, which would affect the efficiency of self-selection. These qualifications highlight reasons for caution in how one interprets our results. Nevertheless, our

results are an important first step in utilizing laboratory methods to address some important

questions regarding the Peter Principle.

REFERENCES

Acosta, Pablo. 2010. Promotions dynamics the Peter Principle: Incumbents vs. external hires. *Labour Economics*, 17:975-86.

Baddeley, Alan D. 1968. A 3 Min reasoning test based on grammatical transformation. *Psychonomic Science*, 10:341-42.

Barmby, Tim., Barbara Eberth, and Ada Ma. 2010. Incentives, learning, task difficulty, and the Peter Principle: Interpreting individual output changes in an organisational hierarchy. Unpublished paper. University of Aberdeen.

Bernhardt, Dan. 1995. Strategic promotion and compensation. *Review of Economic Studies*, 62:315-39.

Bernhardt, Dan., and David Scoones. 1993. Promotion, Turnover, and Preemptive Wage Offers. *American Economic Review*, 83:771-91.

Camerer, Colin F., and Dav Lovallo. 1999. Overconfidence and excess entry: An experimental approach. *American Economic Review*, 89 :306-18.

Gibbons, Robert., and Michael Waldman. 1999. Careers in Organizations: Theory and Evidence, in *Handbook of Labor Economics*, Volume 3B, edited by Orley Ashenfelter and David Card. Amsterdam: Elsevier Science, pp. 2373-2437.

Greiner, Ben. 2004. An Online Recruitment System for Economic Experiments. *Forschung und wissenschaftliches Rechnen* 2003, edited by Kurt Kremer and Volker Macho. GWDG Bericht 63, Göttingen : Ges. für Wiss. Datenverarbeitung, 79-93.

Koch, Alexander K., and Julia Nafziger. 2007. Job assignments under moral hazard: The Peter Principle revisited. IZA Discussion Paper, 2973.

Kruger, Justin., and David Dunning. 1999. Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessment. *Journal of Personality and Social Psychology*, 77:1121-34.

Lazear, Edward P. 2004. The Peter Principle: A theory of decline. *Journal of Political Economy*, 112:S141-S163.

Meyer, Margaret A. 1991. Learning from coarse information: Biased contests and career profiles. *Review of Economic Studies*, 58:15-41.

Milgrom, Paul., and John Roberts. 1992. *Economics, Organization and Management.* Englewood Cliffs, Prenctice Hall International.

Peter, Laurence J., and Raymond Hull. 1969. *The Peter Principle: why things always go wrong.* William Morrow & Company: New-York.

Prendergast, Canice., and Robert H. Topel. 1996. Favoritism in Organizations. *Journal of Political Economy*, 104:958-78.

Van den Steen, Eric. 2004. Rational overoptimism (and other biases). *American Economic Review*, 94:1141-51.

Waldman, Michael. 1984. Job assignments, signalling, and efficiency. *Rand Journal of Economics*, 15:255-67.

Zeiliger, Romain. 2000. A presentation of Regate, Internet based software for experimental economics. *http://www.gate.cnrs.fr/~zeiliger/regate/RegateIntro.ppt*, GATE. Lyon: GATE.

Table 1. Main treatment characteristics

Treatment	Assignment to Hard task	Variance On Easy task	Variance on Hard task	# of sessions	# of subjects
Calibration	Random	U[-12,+12]	U[-4,+4]	2	38
Promotion-σ_{Low}	Standard	U[-12,+12]	U[-4,+4]	2	37
Promotion-σ_{Medium}	Standard	U[-24,+24]	U[-8,+8]	2	38
Promotion-σ_{High}	Standard	U[-75,+75]	U[-25,+25]	2	37
Self-Selection	Self-Selection	U[-12,+12]	U[-4,+4]	2	40

Table 2. Average Ability (raw outcomes—no transitory component) across treatments

	Practice Stages		STAGE 1	STAGE 2	
	Easy task	**Hard task**	**Easy task**	**Easy task**	**Hard task**
Pilot Experiments	42.67	30.00	65.33	71.50	34.83
Calibration	37.00	25.39	59.32	63.84	30.26
Self-Selection	44.20	29.33	67.03	66.80	30.26
Promotion-σ_{Low}	48.22	31.38	70.84	62.58	39.72
Promotion-σ_{Medium}	44.64	28.87	68.74	57.13	40.39
Promotion-σ_{High}	40.11	26.24	60.65	65.41	32.86

Table 3. Stage 1 and Stage 2 average outcomes (*includes* transitory component)

	Non-promoted subjects			Promoted subjects		
		Stage 1	Stage 2		Stage 1	Stage 2
Treatment	# obs.	Outcome	Outcome	# obs.	Outcome	Outcome
Calibration	19	62.47	64.47	19	59.53	63.78
Self-Selection	20	64.25	65.15	20	70.65	83.68
Promotion-σ_{Low}	12	55.92	62.58	25	81.04	83.95
Promotion-σ_{Medium}	15	51.80	60.60	23	84.22	85.41
Promotion-σ_{High}	22	27.14	70.91	15	92.13	52.08

Table 4. Stage 1 ability of promoted/non-promoted subjects and assignment mistakes

	Stage 1-average ability (no transitory component added)		Number of mistakes (percentages)		
Treatment	Non-promoted	Promoted	Mistake 1 (promotion)	Mistake 2 (no-promotion)	Correct decisions
Calibration	59.95	58.68	12 (31.58)	5 (13.16)	21 (55.26)
Self-Selection	61.7	72.35	9 (22.50)	6 (15.00)	25 (62.50)
Promotion-σ_{Low}	57.33	77.32	2 (5.41)	1 (2.7)	34 (91.89)
Promotion-σ_{Medium}	52.67	79.17	4 (10.53)	3 (7.89)	31 (81.58)
Promotion-σ_{High}	60.68	60.60	11 (29.73)	10 (27.03)	16 (43.24)

Table 5. *Self-Selection* task choice and hypothetical choice in *Promotion* treatments

Treatment	Optimistic	Pessimistic	Accurate	Total
Self-Selection	22.50 %	15.00 %	62.50 %	100 %
Promotion-σ_{Low}	13.51 %	13.51 %	72.98 %	100 %
Promotion-σ_{Medium}	26.32 %	18.42 %	55.26 %	100 %
Promotion-σ_{High}	37.84 %	2.70 %	59.46 %	100 %

**Table 6. Actual and hypothetical choices of tasks in *Self-Selection* and *Promotion* treatments (**Probit Models)**

	Actual choice of the hard task		Hypothetical choice of the hard task	
	Self-Selection		*Promotion*	
Dependent variables	Coeff.	S.E.	Coeff.	S.E.
Score in Stage 1	0.035**	0.017	0.029***	0.007
Transitory ability in Stage 1	-0.062*	0.032	-0.010*	0.006
Var. of transitory ability	---	---	0.367***	0.144
Constant	-2.345**	1.160	-1.815***	0.482
Nb of observations	40		150	
Log likelihood	-23.113		-84.867	
LR χ^2	9.23		30.44	
Prob>χ^2	0.010		0.000	
Pseudo R^2	0.166		0.152	

Note: *, **, and *** indicate statistical significance at the 0.10, 0.05, and 0.01 level, respectively.

Figure Captions

Figure 1 : Pay Profiles and Promotion Standard

Figure 2 : Average Stage 1 transitory ability sorted by task in Stage 2
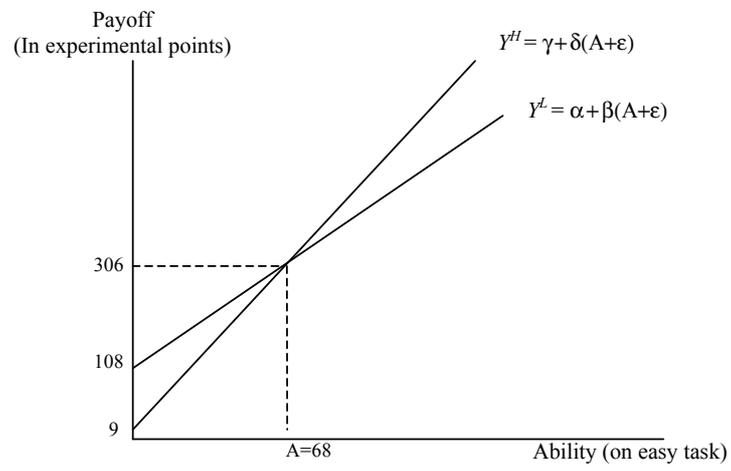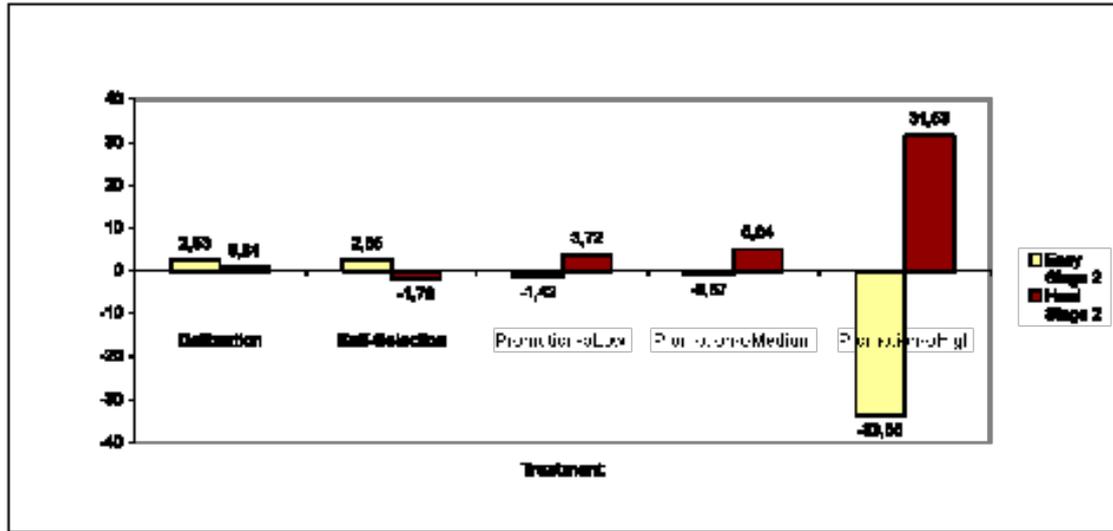
Figure 1

Figure 2

APPENDIX 1– Instructions of the *Promotion-σ$_{Low}$* treatment
(Other instructions available upon request)

This is an experiment in decision-making. This experiment is separated into different parts. During each part, you will perform the following task. The task involves reading a short sentence that describes the order of letters. Based on the sentence, you must type in the order of the letters that the sentence implies. This task is described in details below. You will earn a specific amount of money for each accurately solved sentence. The earnings in points that you will get in each part will be added up and converted into Euros at the end of the session. The conversion rate is the following:

100 points = 1 Euro

In addition, you will receive 4 € for participating in this experiment. Your earnings will be paid in cash at the end of the session in a separate room. Earnings are private and you will never be informed of anyone else's outcomes or earnings in the experiment.

The instructions related to the first part are described below. The instructions relative to the next parts will be distributed later on.

Part 1

The task involves reading a short sentence that describes the order of 3 letters. Based on the sentence, you must type in the order of the letters that the sentence implies. You <u>must</u> determine the proper order of each 3-letter sentence by doing the following. First, determine the order of the first two letters. Then, keep those two letters together and determine where the third letter should go in relation to <u>the first pair</u> of letters. Using this rule, there is only one possible ordering of the 3 letters for each sentence. You can use indifferently lower case letters or capital letters and you can return. You must validate your answer by clicking the enter key.

Here are some examples of 3-letter sentences followed by the correct letter-order.
Short sentence	Correct order
1) B is followed by A, precedes H	BAH
2) X does not precede Y, is not followed by M	MYX

Once you have submitted an answer, a new sentence will automatically appear for you to repeat this task on a different sentence. This will continue for 7 minutes. You will be told on your screen whether you were correct or not after each sentence. If your answer is not correct, you cannot change it; you automatically proceed to the next sentence. You screen indicates the current number of correct letter-orderings and the remaining time.

Your result in units in part 1 is given by the number of correct letter-orderings in this part. For each letter-order that you get correct, you are paid 3 points. In addition, you receive a fixed payoff of 108 points.

Your earnings in points in part 1 = (your result in units in part 1 * 3 points)
+ 108 points

To ensure that you have understood the instructions, you are kindly requested to practice this task by attempting to solve the 3 sentences that will appear on your screen once you have cliked the « practice » button. During this practice, the next sentence will appear only after the current one has been accurately solved. These answers are not taken into account in your result or your earnings.

Please raise your hand if you have any questions. The 7-minute task-period will start after each participant has completed the practice and once we have answered your questions. Communication between participants is forbidden throughout the session.

Part 2

The task involves reading a short sentence that describes the order of 5 letters. Based on the sentence, you must type in the order of the letters that the sentence implies. You <u>must</u> determine the proper order of each 5-letter sentence by doing the following. First, determine the order of the first two letters. Then, keep those two letters together and determine where the third letter should go in relation to <u>the first pair</u> of letters. Then, keep those three letters together and determine where the fourth letter should go in relation to the first three letters. Then, keep those four letters together and determine where the fifth letter should go in relation to the first four letters. Using this rule, there is only one possible ordering of the letters for each sentence.

Here are two examples of 5-letter sentences followed by the correct letter-order.

<u>Short sentence</u>   <u>Correct order</u>
1) A is followed by D, is followed by B, is preceded by C, follows H
HCADB
2) Z does not precede X, does not follow Y, does not follow W, precedes B   XZYWB

Once you have submitted an answer, a new sentence will automatically appear for you to repeat this task on a different sentence. This will continue for 7 minutes.

You will be told on your screen whether you were correct or not after each sentence. If your answer is not correct, you cannot change it; you automatically proceed to the next sentence. You screen indicates the current number of correct letter-orderings and the remaining time.

Your result in units in part 2 is given by the number of correct letter-orderings in this part. For each letter-order that you get correct, you are paid 9 points. In addition, you receive a fixed payoff of 9 points.

> Your earnings in points in part 2 = (your result in units in part 2 * 9 points)
> + 9 points

You are kindly requested to practice this task by attempting to solve the 3 sentences that will appear on your screen once you have cliked the « practice » button. During this practice, the

next sentence will appear only after the current one has been accurately solved. These answers are not taken into account in your result or your earnings.

You can click the « practice » button as soon as you like. The 7-minute task-period will start after each participant has completed the practice. Your earnings during this part will be added to your previous earnings.

The instructions that have just been distributed describe parts 3 and 4. Part 4 will start immediately after part 3 has been completed and a break of 3 minutes.

Part 3

The task involves reading a short sentence that describes the order of 3 letters, as you already did in part 1 but with new sentences. As before, you must type in the order of the letters that the sentence implies. The sentences appear successively during 7 minutes and you will be told on your screen after each answer if this one is correct or not.
Your result in units in this part depends on two elements:
-   The number of correct answers
-   A personal random number

Your personal random number is determined as follows. Once you have performed your task, a « random draw » button appears on your screen. Once you click this button, the computer program generates your personal random number. Your number can take any value between -12 and +12, included. Each number has an equally likely chance to be drawn. There is an independent random draw for each participant. Your personal random number is independent of your performance at the task.

Your result in units in part 3 is given by the addition of the number of correct answers and your personal random number in this part.

Your result in units in part 3 = your number of correct answers in part 3 + your personal random number in part 3

For example, if you give 38 correct answers and you draw a random number of -3, then your result is 35 units (i.e. 38-3). If, in contrast, you give 27 correct answers and you draw a random number of +3, then your result is 30 units (i.e. 27+3).

Once you have drawn your number, you are informed on your result in units and on your earnings in points in part 3. For each unit, you get paid 3 points. In addition, you receive a fixed payoff of 108 points.

Your earnings in points in part 3 = (your result in units in part 3 * 3 points) + 108 points

It is important to note that this is **your result** in this part (your number of correct answers + your random number) that will determine the type of task you will perform in part 4.

Part 4

At the beginning of part 4, you are informed on your assignment to 3-letter sentences or to 5-letter sentences during this new part. Your assignment depends on your result in units during part 3.

- If your result in units in part 3 (i.e. your number of correct answers + your random number) amounts to **68 or more**, then you are assigned to solving **5-letter** sentences under the conditions detailed below.

- If your result in units in part 3 is **lower than 68 units**, then you are assigned to solving **3-letter** sentences under the conditions detailed below.

Whatever your assignment is, you will solve sentences during 7 minutes. Your result in units in part 4 is given by the addition of the number of correct answers and your personal random number in this part.

> Your result in units in part 4 = your number of correct answers in part 4
> + your personal random number in part 4

Here are described the rules for each type of assignment.

◊ **If your result in part 3 assigns you to 3-letter sentences**

Once you have performed your task, you click the « random draw » button that generates your personal random number. Your number can take any value between -12 and +12, included. Each number has an equally likely chance to be drawn. There is an independent random draw for each participant.

You are then informed on your result in units and on your earnings in points in this part. For each unit, you get paid 3 points. In addition, you receive a fixed payoff of 108 points.

> Your earnings in points in part 4 if assigned to 3-letter sentences = (your result in units in part 4
> * 3 points)
> + 108 points

◊ **If your result in part 3 assigns you to 5-letter sentences**

Once you have performed your task, you click the « random draw » button that generates your personal random number. Your number can take any value between -4 and +4, included. Each number has an equally likely chance to be drawn. There is an independent random draw for each participant.

You are then informed on your result in units and on your earnings in points in this part. For each unit, you get paid 9 points. In addition, you receive a fixed payoff of 9 points.

> Your earnings in points in part 4 if assigned to 5-letter sentences = (your result in units in part 4
> * 9 points)
> + 9 points

Although your assignment to one task or the other depends exclusively on your result in part 3, you will be asked to indicate which task you would prefer being assigned to in part 4 si you would have the choice :

Either the 3-letter sentences, with a random number between -12 et +12, a payoff of 3 points per unit and a fixed payoff of 108 points

Or the 5-letter sentences, with a random number between -4 et +4, a payoff of 9 points per unit and a fixed payoff of 9 points.

You will be requested to submit your answer at the beginning of part 4.
Your earnings during these two parts will be added up to your previous payoffs.

Please raise your hand if you have any questions regarding these instructions. We remind you that communication between participants is still forbidden.

Footnotes

[1] There are other sources of distortion in job assignment where the principals themselves manipulate information to favor specific employees (Meyer, 1991; Prendergast and Topel, 1996) or to protect against potential competitors (see Bernhard and Scoones, 1993, and Bernhardt, 1995, after Waldman, 1984). See also Gibbons and Waldman (1999). We do not consider here these other sources of distortion.

[2] This implies that non-naïve employers may adjust promotion rules such that the Peter Principal is of no concern. Fairburn and Malcomson (2001) have shown in a different theoretical environment that promotion rules can be efficient under certain assumptions, but such rules may also give rise to distortion in job assignments if employees are risk-averse.

[3] In Barmby et al. (2006), the importance of luck is proxied by a logit model of promotion depending on the distance between the evaluation rating of the individual and the average rating of his current grade.

[4] In Lazear (2004)'s general model, the agent is not better informed than the principal. He also considers the case of asymmetric information, which is what we consider given our interest in the self-selection question.

[5] When transitory components exist, Lazear (2004) shows that promotion cutoffs will be adjusted to take into account the Peter Principal. Because we focus on agent behavior, we exogenously impose a promotion standard rather than submit it to subject choice.

[6] Lazear (2004) includes predictions of principal (employer) behavior as well, which we do not examine in our experiments.

[7] We mix passive (e.g., "*is followed by*") and active (e.g., "*precedes*") statements, positive (e.g., "*follows*") and negative statements ("*does not precede*") across sentences. Subjects are

given a precise rule to use for lettering ordering, and with the rule there is a unique letter

ordering for each sentence description (see instructions in Appendix).

[8] A minor change in fixed payments was made for the *Promotion* and *Self-Selection*

treatments. The easy task fixed payment was lowered from 135 to 108 points (9 points is

equivalent to one output unit, which is one hard task solved), to increase anticipated promotions

to around 50%--this lowered the Fig 1. intercept to output=12 for the easy task (from 15 in the

Calibration treatment).

[9] This expectation is based on sample estimates from the *Calibration* treatment, but it is not

critical to our experiment whether more or less than 50% of subjects have ability beyond the

cross-over point. That is, the Peter Principle is fundamentally about whether random

components create matching inefficiencies. Our experiment does not test the ability of principals

to anticipate this Principle and adjust the promotion cutoff.

[10] Recall that this easy-hard task correspondence was from the pilot sessions (n=12), where

average easy task output was about 2 times that of the hard task output. In the *Calibration*

treatment we again find subjects completing approximately twice as many easy tasks as hard

tasks post-practice stage. It is important to realize that we *cannot* use the practice stages prior to

Stages 1 and 2 to generate a hard-to-easy task conversion rate for each individual subject because

of the potential learning confound in those data.

[11] The instructions describe the experiment as having 4 parts. Parts 1 and 2 correspond to

the practice stages. This is a semantics issue, as we referred to the practice rounds for both the

easy and hard task described in the text as "parts".

[12] Because the practice stages were 7 minutes long, as were the main stages, and because they were compensated similarly, they are useful for examining learning trends by comparing correct sentences in the practice easy task to the number correctly sorted in Stage 1.

[13] More precisely, we multiply hard task outcomes by 2.05 to convert to easy task equivalent outcomes. This conversion ratio was decided upon after pilot experiments using the easy task and hard task average outcomes in Stage 2 of the pilot sessions, which were 71.50 and 34.83, respectively.

[14] It is also not the case that less able subjects are unaware of their incompetence because they lack the metacognitive ability to realize it (Kruger and Dunning, 1999). Indeed, the subjects are informed about their score before they draw their transitory ability component.

[15] These findings are consistent with van den Steen (2004), whose theoretical model shows that, if rational agents differ in their prior beliefs, they are more likely to choose actions for which they overestimate the probability of success. Here, it means that we are more likely to observe optimistic subjects among those who have chosen the hard task. In other words, according to van den Steen (2004), there is a choice-driven endogenous overconfidence in the data. Therefore, allowing self-selection of one's task may itself generate optimism that reduces job choice efficiency

[16] A final approach is to examine the cost of task mis-assignment, under the assumption that effort distortion should increase with the cost of task mis-assignment. Unfortunately, the farther (closer) one's ability is from the task-indifference ability point, the lower (higher) probability of task mis-assignment—that is, the lower the chance a random term causes you to be mis-assigned—but the higher (lower) the cost of a mis-assignment (see Fig. 1). Thus, it is unclear

how expected mis-assignment costs behave as a function of one's ability without knowledge of

one transitory component expectation.