Master thesis Economics & Business:
Economics of Markets, Organizations and Policy

Title: Employee Learning Dynamics under the Introduction of Relative Performance Information: a Theoretical Model

Student: Bastiaan Dessing (ID 314087)
Contact: bastiaandessing@hotmail.com

Supervisor: Prof. dr. Robert Dur (department of economics)
Contact: a.j.dur@ese.eur.nl

Employee Learning Dynamics under the Introduction of Relative Performance Information: a Theoretical Model

Bastiaan Dessing

Erasmus University Rotterdam
Erasmus School of Economics

December 29, 2011

Abstract

The empirical literature is generally supportive of a positive effect on employees' performance of the introduction of relative performance information (RPI). However, different studies give different explanations for this result. The thrust of this paper is to theoretically introduce a learning dynamic which may offer an alternative, non-psychological, explanation. We consider both the case where employees do and do not expect RPI to be provided at some intermediate stage. We also take into account different forms of RPI and show that the level of refinement of RPI is of consequence. Although placed in a piece-rate pay context, we argue that the learning dynamic is equally valid in other pay contexts. Finally, we look at the dynamic from an employer perspective and develop an optimal feedback policy.

Keywords: relative performance information, piece-rate pay, learning dynamic, feedback policy

I. Introduction

It is well known that it is easier to find a solution to a problem if one knows the solution exists or that it is attainable. Cognitive psychologists argue that you may already have the necessary skill-set to solve a problem, but you may not have the *metaskill* to know when to exercise these skills (Mayer 1998). Problem-solvers may also have a low self-efficacy which means that some may not even embark on a problem-solving process even though they would be able to solve the problem at hand. Information that comparable individuals have been able to reach a solution may then increase self-efficacy, show to individuals that they indeed have the necessary skills to successfully attack a problem, and thereby start the ball rolling towards effective problem-solving.

We can also look at this phenomenon from an economic perspective. If you receive information on a solution's existence or attainability, the expected benefits of searching for the solution go up. This learning dynamic may easily be translated to a professional setting. For instance, an employee in a firm is more motivated to search for an optimal work routine if he knows a better work routine exists or is attainable. He may get this inference from the revelation of relative performance information (RPI). If similar colleagues generate higher output, the employee in question may start to wonder why. True, a good employer would have extensive guidelines in place to steer employees to an efficient routine. But even from the employer's perspective clear cut guidelines should not reach too far. No employee is the same and learns in the same way.

This paper aims to formalize this learning dynamic and aims to establish that the introduction of relative performance information may increase employee performance under piece-rate pay schemes. Moreover, we show that the increase in output due to this learning dynamic should be most prevalent among low performing individuals. We put the learning dynamic in a piece-rate pay context as this is analytically tractable and abstract from other employee incentives as much as possible in the interpretation of the model. Nonetheless, the learning dynamic should be applicable to other pay contexts as well.

This article is organized as follows: we will embed our learning dynamic in the theoretical and empirical literature first (II). We then develop a theoretical model (III), where we will consider the possibility of an unexpected RPI introduction (IV.A) and an RPI introduction employees anticipate on (IV.B). Section IV.C differentiates between different kinds of RPI. Next, Section V analyzes the learning dynamic from the employer's perspective and develops whether RPI should be given and should be announced. In section VI we discuss the results and possible extensions and, finally, we conclude and summarize in section VII.

II. Related literature

The broader application of economic research on RPI is embedded in its capacity to be a non-monetary employee performance driver which may be steered by an employer. In certain circumstances, monetary incentives may fail to create the right motivation or may do so only in a cost-intensive way. Apart from RPI, there are other non-monetary ways too that an employer may try to influence employee performance. He may motivate employees by giving recognition (e.g. Bradler et al. 2011 and Neckermann et al. 2010), attention or by delegating certain responsibilities (e.g. Swank & Visser 2006). Task design in general may trigger intrinsic motivation (Banabou & Tirole 2003). He may also simply try to boost confidence by verbally differentiating between employees (Crutzen et al. 2010). This may be considered different from RPI, due to its informal, sporadic and unsubstantiated nature.

Returning to RPI, it is important to start out with a clear distinction between relative performance information (RPI) and relative performance evaluation (RPE), as these two terms are often nonchalantly used as substitutes in the scientific literature. Whereas the former states the presence of knowledge on an employee's performance ranking, the latter comprises the act of differentiating pay or career prospects according to different spots in the performance ranking. RPI may be present without RPE, and RPE may be conducted without sharing RPI. That is, employees may work in a *tournament setting* without having access to intermediate standings over the course of the tournament. It is surprising to find that only RPE has been covered quite extensively in the economic literature (e.g. Lazear & Rosen 1981 and Dye 1992). Often RPE is placed in the context of CEO compensation, where a lot of noise makes it difficult to make an accurate inference on the CEO's isolated performance (e.g. Gibbons & Murphy 1991). Tying pay to the performance of others who incur the same shocks in performance is efficient in the presence of risk-aversion.

The performance effects of sharing RPI in a tournament setting (i.e. in an RPE-context) are theoretically deemed ambiguous. On one hand, frontrunners may feel safe from competition and reduce effort, just as lagging employees may feel it is impossible to catch up and decide to slack off. Both reduce performance. On the other hand, when competition is close both frontrunner and underdog may decide to go for a final push and increase effort. Even when competition is not close, pride on behalf of the frontrunner or shame on behalf of the underdog may induce higher effort.

Void from attached monetary benefits (i.e. under a pay scheme other than based on relative performance) the introduction of RPI still has several effects, some of which fall beyond the scope of economic science. The first and foremost effects are of a psychological nature. Few people are left untouched by feedback of absolute performance (See Kluger & DeNisi 1996 for an overview), let alone by learning to be better or worse than expected compared to co-workers. Social comparison theory (Festinger 1954) formulates that people have a continuous desire to compare themselves with others. If this comparison turns out unfavorably, people's sense of self-identity may suffer (Tesser & Campbell 1980). In order to avoid or change this, behavior and hence performance may change. This is clearly showed by Tran & Zeckhauser (2009) who find that students are willing to pay substantial amounts of money in order to increase their rankings, even in a *private* RPI-context (i.e. learning about your place in, and the shape of, the performance distribution rather than learning where others stand in the rankings). If RPI is not private, employees may exhibit status concerns instead or in addition to effects on their self-identity. Empirical work shows that performance increases of others within your reference group may decrease your own utility (e.g. Luttmer 2005 and Charness & Kuhn 2007). However, it often proves difficult to completely disentangle pure status concerns from career concerns where an ultimate monetary reward may drive performance. Another issue is that the mere sharing of RPI may show to workers that competitive behavior is appropriate (Beck & Seta 1980). This intuition closely resembles ideas put forward by management theory (e.g. Milkovich & Newman 1996), stating that RPI may create animosity among employees thereby hurting morale and hence performance. The extent to which all these factors discussed play a role may be contingent on heterogeneities in general and gender differences in particular. Taking this into account makes it difficult to theoretically assess the value of sharing RPI in terms of performance effects. A lot seems to depend on the specific circumstances and people involved.

Empirically, RPI in a tournament setting has not been proven worthwhile yet. In their laboratory experiment Eriksson et al. (2009) find no evidence for performance improvements across the performance distribution. Hannan et al. (2008) find evidence for deteriorating average performance in a laboratory experiment, where especially the lower performing individuals suffer from a dramatic performance fall.

Sharing RPI under a fixed wage regime has been given empirical attention by Falk & Ichino (2006) in a field experiment and Kuhnen & Tymula (2011) in a laboratory experiment. Both find increases in average performance due to an RPI introduction, although neither of them allows explicitly or implicitly for learning dynamics. Kuhnen & Tymula's research indicates a performance increase under top performers especially, while Falk & Ichino find exactly the opposite.

Under a piece-rate pay scheme, Blanes I Vidal & Nossol (2009) find a significant increase in performance under both high and low- productivity workers in their field experiment. This is partly due to the announcement of a future RPI introduction, but mostly due to the RPI introduction itself. Their main interpretation of these results deals with status concerns, although they do not exclude the possibility of the presence of some learning effects. In the light of a private RPI-context this could at least partially explain the results. Hannan et al. (2008) find an increase in performance due to RPI, both for low and high performing individuals. They do not offer a preferred explanation of these results. Eriksson et al. (2009) find no increase in performance, yet do not allow for learning dynamics in their experimental setup. Barankay (2011) finds a negative performance response to private RPI in a natural field experiment among salesmen in an office furniture. There is no significant difference between top and bottom performers, but more so between men and women. Whereas men decrease their performance, the performance of women is virtually left unchanged. In a related field research by Barankay (2011), he again finds a performance decrease, this time among participants in a crowd-serving platform. There are no significant heterogeneity effects. Interestingly, he looks also at RPI from a participation perspective. He finds that 20% fewer people return to the platform of the group which is confronted with RPI. Finally, in a natural experiment Azmat & Iriberri (2010) find evidence for an increase in performance under high school students due to an RPI introduction. As grades are different for different performances, this result can be thought of as placed in some kind of piece-rate pay context. The performance increase is significant across the performance distribution, although especially prevalent among the best and worst performers compared to average performers. They do not give an explicit preferred explanation, but do dismiss the effect of status concerns as one among them stressing the private nature of RPI in their experiment. Although the articles studying RPI in a piece-rate pay context (and in general: see table 1) hint in the way of a performance increase across the performance distribution, they are even less consistent in their preferred explanation of this result.

Table 1: Summary performance effect RPI

Paper	Type of reward	Total performance effect
Eriksson et al. (2009)	Tournament	No performance effect
	pay-for-performance	No performance effect
Hannan et al. (2008)	Tournament	Negative
	pay-for-performance	Positive
Barankay (2011)	pay-for-performance	Negative
Barankay (2011)	pay-for-performance	Negative
Falk & Ichino (2006)	Fixed	Positive
Kuhnen & Tymula (2011)	Fixed	Positive
Blanes I Vidal & Nossol (2009)	pay-for-performance	Positive
Azmat & Irriberry (2010)	pay-for-performance	Positive

The learning dynamic I aim to model may help at this point. While it has an unambiguous positive effect on performance, little theoretical attention has so far been dedicated to this intuition. Closest are Hannan et al. (2008) who notice that "relative feedback helps an individual assess the range of potential performance levels, and thereby facilitates learning and appropriate effort allocation". However, they do not venture to formalize this intuition. Somewhat similar is Ederer (2010)'s intuition, as he notes that "providing [relative] feedback on performance helps workers to do their jobs or plan their futures better by giving them better information on which to base their decisions". However, he specifically links this to differences in ability which is not necessarily a heterogeneity that is a prerequisite for our learning dynamic. By referring to job planning (and labeling his intuition a "sorting effect") he invokes a participation perspective which is best placed in a tournament setting (as is indeed the case in his research). This lack of theoretical attention to our learning dynamic is understandable to an extent that a learning dynamic may not be the most important effect of an RPI introduction in the short run. Changes in performance due to psychological effects may surface instantly when RPI is given. A learning dynamic is more tedious in this respect as RPI only implants the seed. It merely starts and enables the process towards a future performance increase. After all, when an employee comes to the understanding that comparable workers perform better he still needs to uncover why this is the case before performance can increase. This may take considerable time.

III. The Model

We assume a mass of employees normalized to 1. An individual employee i is unsure about the usefulness of investing in an efficient work routine given by θ. The actual usefulness can either be high θ_H or low θ_L, where $\theta_H > \theta_L \geq 0$. Each employee gets a private signal s on θ, given by s_H if positive and s_L if negative respectively. The chance that this signal is correct is given by ρ, where $1 \geq \rho \geq \frac{1}{2}$. Thus, $\Pr(\theta = \theta_H | s_H) = \Pr(\theta = \theta_L | s_L) = \rho$. Accordingly, $\Pr(\theta = \theta_L | s_H) = \Pr(\theta = \theta_H | s_L) = 1 - \rho$. Moreover, ρ is assumed to be common knowledge, yielding $E(\theta | s_H) = \rho\theta_H + (1 - \rho)\theta_L$ and $E(\theta | s_L) = (1 - \rho)\theta_H + \rho\theta_L$. Each individual employee needs to make two decisions on how much learning activities $l_{1,i} \in [0, \infty]$ and $l_{2,i} \in [0, \infty]$ he will undertake (i.e. how much time he will devote searching for an optimal work routine). The two choices to be made reflect two respective periods, the *initial stage* (before RPI is shared) and the *RPI stage* (after RPI is shared). Plausibly, learning activities cannot be destructive or made undone (so both $l_{1,i} \geq 0$ and $l_{2,i} \geq 0$). Individual employee productivity after the initial stage is given by $y_{1,i} = l_{1,i}\theta + \varepsilon_i$. The noise term ε_i captures differences in employee productivity due to other factors than learning activities and is distributed normally with mean 0 and variance σ_ε^2. Likewise, productivity of the RPI stage is given by $y_{2,i} = (l_{1,i} + l_{2,i})\theta$. Hence, total productivity is given by $y_{T,i} = (2l_{1,i} + l_{2,i})\theta + \varepsilon_i$. The fact that $l_{1,i}$ has a larger effect on productivity than $l_{2,i}$ can be interpreted as the opportunity costs of postponing the investment decision. We assume that through pay-for-performance an employee completely internalizes his productivity. Learning activities are costly as it takes effort to find new work routines. They are denoted by the convex functions $C(l_{1,i}) = \frac{1}{2}{l_{1,i}}^2$ and $C(l_{1,i}, l_{2,i}) = \frac{1}{2}(l_{1,i} + l_{2,i})^2 - \frac{1}{2}{l_{1,i}}^2$ respectively for the intial and RPI stage. Convexity is assumed as it should be increasingly difficult to find a better work routine. The total effect of learning activities on an individual's expected utility is now given by $U_i(l_{1,i}, l_{2,i}) = (2l_{1,i} + l_{2,i})\theta + \varepsilon_i - \frac{1}{2}(l_{1,i} + l_{2,i})^2$. The timing of the model is as follows:

1) Nature draws θ.

2) Each employee observes his private signal.

3) Each employee chooses l_1.

4) Each employee observes the distribution of y_1 through RPI.

5) Each employee chooses l_2.

6) Payoffs are realized.

Phase 1 , 2 and 3 correspond to the *initial stage*, while phase 4 and 5 reflect the *RPI stage*.

IV. Analysis

This section analyzes employee behavior in reaction to and in anticipation of RPI when they do not expect RPI at some point (section A) and when they do (section B). Section C considers employee behavior when the RPI they will receive is of a less refined kind.

A. Unexpected RPI introduction

In this first case we discuss, an employee maximizes his expected utility without anticipating on RPI. His objective function simplifies to $U_i(l_{1,i}, l_{2,i}) = (2l_{1,i} + l_{2,i})E(\theta) - \frac{1}{2}(l_{1,i} + l_{2,i})^2$. The employee does not expect to learn anything due to RPI, and thus we can analyze his investment choice at the initial stage as if he is able to set both l_1 and l_2 simultaneously. While subject to the same impact on utility costs, the marginal product for $l_{1,i}$ is higher than for $l_{2,i}$. We can therefore restrict $l_{2,i} = 0$, take the derivative of the objective function to $l_{1,i}$ and solve for $l_{1,i}$. This yields the optimal initial investment choice in learning activities $l_{1,i} = 2E(\theta)$. We can distinguish two different scenarios manifesting itself through either a positive or negative signal. $l_{1,i}|s_H = 2[\rho\theta_H + (1-\rho)\theta_L]$ and $l_{1,i}|s_L = 2[(1-\rho)\theta_H + \rho\theta_L]$. Obviously, the higher one estimates the usefulness of learning activities, the more one will pursue it at the initial stage ($l_{1,i}|s_H > l_{1,i}|s_L$). Only in the case of an uninformative signal (i.e. $\rho = \frac{1}{2}$) are both scenarios the same. Comparative statics on the first period investment choices show that both $l_{1,i}|s_H$ and $l_{1,i}|s_L$ increase in θ_H and θ_L. ρ, however, has a positive effect on $l_{1,i}|s_H$, but a negative effect on $l_{1,i}|s_L$, increasing the probability of θ_H and θ_L respectively.

Let us now turn to the RPI stage. At this instant, employees observe a bimodal distribution of y_1 with mean either

$$\mu_H = 2\theta_H\left[\rho[\rho\theta_H + (1-\rho)\theta_L] + [(1-\rho)[(1-\rho)\theta_H + \rho\theta_L]]\right]$$
$$\mu_L = 2\theta_L\left[\rho[\rho\theta_L + (1-\rho)\theta_H] + [(1-\rho)[\rho\theta_H + (1-\rho)\theta_L]]\right]$$

in case of the realization of θ_H and θ_L respectively.[1]

Since ρ, θ_H and θ_L are all known, an employee is now able to identify the state of the world. Note that this identification can take place even in the case of an uninformative private signal. In the mean of the performance distribution the effect of the error term vanishes, which establishes a perfect inference on the state of θ in any scenario. In other terms, one need not to rely on the fact that knowledge on the state of θ may already be transferred by identifying the shape of the distribution (thereby reading which signal has been more prevalent and therefore correct). Note also that knowledge on the state of θ will not be gained on the basis of observing only one's own output, as we deal with two unknowns: θ and ε_i.[2]

Whether re-engaging in learning activities (i.e. positive second period investing) is desirable depends on an employee's first period choice and the state of the world. Obviously, if the state of the world turns out to be θ_L no second period investing takes place. The interesting case emerges when the state of the world is θ_H. The objective function an employee faces still takes the shape of $U_i(l_{2,i}) = (2l_{1,i} + l_{2,i})\theta_H - \frac{1}{2}(l_{1,i} + l_{2,i})^2$. Taking the derivative to $l_{2,i}$ and solving for $l_{2,i}$ yields $l_{2,i} = \theta_H - l_{1,i}$ as the optimal second period investment choice. Filling in $l_{1,i}|s_H = 2[\rho\theta_H + (1-\rho)\theta_L]$ in the optimal second period investment choice, while considering $\rho \geq \frac{1}{2}$, learns that second period investing is not feasible

[1] Technically, the distribution in question can only be formally called bimodal if the two centers are sufficiently far apart. In the case of θ_H this would mean that $\theta_H[\theta_L + 2\rho\theta_H - 2\rho\theta_L - \theta_H] > \sigma_\varepsilon$. In general, the centers are formed by the two different investment decisions by employees with a correct and incorrect signal.

[2] Depending on the properties of both unknowns, minor information may be deduced from observing one's own output. However, this pales in comparison to the effect of the observation of the total distribution of output. Hence, I comfortably discard this notion.

for an employee with a positive private signal. His optimal investment choice would be negative. This is not possible, hence, $l_{2,i} = 0$. For an employee with a negative private signal we can fill in $l_{1,i}|s_L = 2[(1-\rho)\theta_H + \rho\theta_L]$. This leaves a second period investment choice of $l_{2,i} = (2\rho - 1)\theta_H - 2\rho\theta_L$ under the condition that the total is positive, $\theta_H > \frac{\theta_L}{(1-\frac{1}{2\rho})}$. Comparative statics on this last restriction shows that if ρ approaches one it is less restrictive then when in it is closer to one half. This makes sense, as when ρ is large an employee with a negative signal almost completely disregards the possibility that $\theta = \theta_H$. In contrast, if ρ is relatively small an employee with a negative signal already takes into account the significant possibility that $\theta = \theta_H$. In order to engage in second period investing, the difference between θ_H and θ_L now needs to be more substantial.

In equilibrium expected payoffs are given by the following functions:

$$E(U_i|s_H) = (1-\rho)4[\rho\theta_H + (1-\rho)\theta_L]\theta_L + \rho 4[\rho\theta_H + (1-\rho)\theta_L]\theta_H + \varepsilon_i - 2[\rho\theta_H + (1-\rho)\theta_L]^2$$
$$E(U_i|s_L) = (1-\rho)[3\theta_H - 2\rho\theta_H + 2\rho\theta_L]\theta_H + \rho 4[(1-\rho)\theta_H + \rho\theta_L]\theta_L + \varepsilon_i - \rho 2[(1-\rho)\theta_H + \rho\theta_L]^2 - (1-\rho)\frac{1}{2}[\theta_H]^2$$

The important thing to see is that productivity may increase for employees with an incorrect negative private signal through an unexpected RPI introduction. That is, productivity with second period investing, $[3\theta_H - 2\rho\theta_H + 2\rho\theta_L]\theta_H$, is larger than productivity without second period investing, $4[(1-\rho)\theta_H + \rho\theta_L]\theta_H$, when the condition on θ_H is met. All other employees remain on the same productivity level. As employees with an incorrect negative signal are expected to stand relatively low in the first period performance distribution, we can expect the increase in performance to be most prevalent (although not exclusively so) among the initially low-performing individuals. To what extent this is the case depends on the properties of the error term. The number of employees who are able to increase productivity in the case of $\theta = \theta_H$ corresponds to a fraction $(1-\rho)$. Plausibly, the less likely it is a private signal is correct, the more there are employees with a need to update their investment in learning activities.

B. Expected RPI introduction

If employees expect a RPI introduction we have to take into account a forward looking component. Employees know that conditional on certain behavior of others they will receive information on the state of θ after the first period. We can solve this game by backward induction.

For the moment we assume that employees expect to fully learn the state of θ after the first period (beliefs). Let us start with the possibility of a positive signal. If the signal turns out to be false, obviously no further investments take place ($l_{2,i} = 0$). If the signal is correct, the optimal second period investment choice would be $l_{2,i} = \theta_H - l_{1,i}$ as long as the total is positive. That is, $l_{1,i}$ should be sufficiently small to allow for any positive second period investment. If we fill in both second period investment choices and weigh them by their corresponding probabilities we have the following objective function:

$$U_i(l_{1,i}) = \rho\left[[l_{1,i} + \theta_H]\theta_H - \frac{1}{2}(\theta_H)^2\right] + (1-\rho)\left[(2l_{1,i})\theta_L - \frac{1}{2}(l_{1,i})^2\right]$$

Optimizing yields $l_{1,i} = 2\theta_L + \frac{\rho}{(1-\rho)}\theta_H$. Clearly, the restriction on $l_{1,i}$ is never met as $2\theta_L + \frac{\rho}{(1-\rho)}\theta_H \geq \theta_H$ for all values of ρ. This leaves the option of completely disregarding an RPI introduction as the best option for an employee with a positive signal. The opportunity costs of under investing in the first period are simply too large as these employees attach such a high possibility to $\theta = \theta_H$. We can therefore conclude that the forward looking component brings no changes to the behavior of employees with a positive signal, regardless whether it turns out to be a correct signal or not.

However, in the case of a negative signal things are different. The employee in question now again takes into account two scenarios. First, the signal might be correct hence no second period investing would take place ($l_{2,i} = 0$). Second, the signal might be incorrect and the second period choice would be $l_{2,i} = \theta_H - l_{1,i}$ as long as $\theta_H > \frac{\theta_L}{(1-\frac{1}{2\rho})}$. If we fill in these second period investment choices in the objective function and weigh both scenarios by their corresponding probability we have the following objective function conditional on $l_{1,i}$.

$$U_i(l_{1,i}) = \rho\left[2l_{1,i}\theta_L - \frac{1}{2}(l_{1,i})^2\right] + (1-\rho)\left[(l_{1,i} + \theta_H)\theta_H - \frac{1}{2}(\theta_H)^2\right]$$

Optimizing the objective function yields $l_{1,i} = 2\theta_L + \frac{(1-\rho)}{\rho}\theta_H$. This shows that for an employee with a negative signal second period investing is indeed feasible. Comparative statics on $l_{1,i}$ show that conform intuitive expectations $l_{1,i}$ increases in θ_L and θ_H, and decreases in ρ.

The last thing we have to check is whether the beliefs we assigned are consistent with equilibrium. Indeed they are, since after the first period employees will observe a bimodal performance distribution with either

$$\mu_H = \theta_H \left[\rho 2[\rho\theta_H + (1-\rho)\theta_L] + (1-\rho)[2\theta_L + \tfrac{(1-\rho)}{\rho}\theta_H]\right]$$
$$\mu_L = \theta_L \left[\rho \left[2\theta_L + \tfrac{(1-\rho)}{\rho}\theta_H\right] + 2(1-\rho)[\rho\theta_H + (1-\rho)\theta_L]\right]$$

in case of θ_H and θ_L respectively. Again, since θ_H, θ_L and ρ are known employees have perfect knowledge on the state of θ.

In equilibrium expected payoffs are given by:

$$E(U_i|s_H) = (1-\rho)4[\rho\theta_H + (1-\rho)\theta_L]\theta_L + \rho 4[\rho\theta_H + (1-\rho)\theta_L]\theta_H + \varepsilon_i - 2[\rho\theta_H + (1-\rho)\theta_L]^2$$
$$E(U_i|s_L) = \rho 2\left[2\theta_L + \tfrac{(1-\rho)}{\rho}\theta_H\right]\theta_L + (1-\rho)\left[2\theta_L + \tfrac{(1-\rho)}{\rho}\theta_H + \theta_H\right]\theta_H + \varepsilon_i - \rho\tfrac{1}{2}\left[2\theta_L + \tfrac{(1-\rho)}{\rho}\theta_H\right]^2 - (1-\rho)\tfrac{1}{2}[\theta_H]^2$$

It is interesting to compare the results of this section with the previous one. The principal difference is in the first period investment choice for a negative signal, $l_{1,i}|s_L = 2[(1-\rho)\theta_H + \rho\theta_L]$ in the former and $l_{1,i}|s_L = 2\theta_L + \frac{(1-\rho)}{\rho}\theta_H$ in the latter case. As long as the restriction on θ_H is met first period investment is lower in the case of an expected RPI introduction. Employees anticipate on the idea that they are still able to reap a part of the benefit of a high 'productivity parameter' in the case they wrongly follow their negative signal. This diminishes the opportunity costs of under investing in the first period. Therefore, the increase in performance may be greater due to an RPI introduction when it is actually expected. Nonetheless, the total productivity over both periods combined is higher when RPI is in fact not expected.

In general we can say that the learning dynamic remains present when RPI is expected, and again mostly so among the lower performing individuals.

C. Clarity of RPI

Until now we have accepted RPI as it comes in our model: as specific and detailed as possible. Clearly, however, RPI can be given in many, less refined forms. One commonly seen possibility is to form cohorts of employees and to share with employees in which cohort he ranks. Intuitively, this may be detrimental to any learning dynamic. Less specific information will reduce, if not eliminate learning opportunities compared to more specific information as an inferior inference will be made on the underlying variables. This section aims to prove this.

Consider the extreme example of an employee unexpectedly learning not the whole first period performance distribution, but merely whether he performs under (s_U) or above (s_A) the median employee (i.e. whether he belongs to the top or bottom 50 percent). For an employee with a positive signal this will spell no difference in behavior as more specific information is already insufficient to establish positive second period investment. For an employee with a negative signal the following happens. Let us assume for the moment that performance is virtually only driven by learning activities (i.e. σ_ε approaches 0). If σ_ε approaches zero, the probability of learning to be under the median $P(s_U)$ approaches $\frac{\frac{1}{2}}{\rho}$ in case of θ_L and a negative signal. Formally,

$$\lim_{\sigma_\varepsilon \to 0} P(s_U) \mid \theta_L, s_L = \frac{\frac{1}{2}}{\rho}$$

$$\lim_{\sigma_\varepsilon \to 0} P(s_A) \mid \theta_L, s_L = 1 - \frac{\frac{1}{2}}{\rho}$$

$$\lim_{\sigma_\varepsilon \to 0} P(s_U) \mid \theta_H, s_L = 1$$

$$\lim_{\sigma_\varepsilon \to 0} P(s_A) \mid \theta_H, s_L = 0$$

where $P(s_A)$ is the probability of learning to be performing above the median employee. An employee i uses this to update his subjective probability ρ_i that his private signal is correct. Obviously, if he learns to be performing above the median there must be a significant amount of employees who received a negative signal as well. This makes the employee perfectly confident about the accuracy of his private signal ($\rho_i | s_A$, $s_L = 1$) and will certainly induce no further investments in learning activities. Alternatively, the employee may learn to be performing below the median. This yields $\rho_i | s_U \,, s_L = \frac{\frac{1}{2}}{1\frac{1}{2} - \rho}$. Clearly, by learning to be performing among the worst half of employees the employee in question has become a bit less secure about the accuracy of his private signal ($\rho_i | s_U \,, \; s_L < \rho$). His subjective expectation of θ is now given by $E(\theta) | s_U \,, s_L = \rho_i \theta_L + (1 - \rho_i)\theta_H$. As the employee is less secure about his negative signal, his expectations on θ have risen ($E(\theta) | s_U \,, s_L > E(\theta) | s_L$). Filling in this new expectation in the second period objective function gives the following:

$$U_i(l_{2,i}) = (2l_{1,i} + l_{2,i})(\rho_i \theta_L + (1 - \rho_i)\theta_H) - \frac{1}{2}(l_{1,i} + l_{2,i})^2$$

Filling in $l_{1,i} | s_L$ and optimizing yields $l_{2,i} = (2\rho - \rho_i - 1)\theta_H + (\rho_i - 2\rho)\theta_L$ given that the total cannot be negative. Clearly, however, $l_{2,i}$ does not fit that restriction as the first term is always non-positive and the second term strictly negative. That means that no further investments in learning activities will be made by an employee with a negative signal, regardless of any RPI.

Now consider briefly that we loosen our restriction on σ_ε. The larger σ_ε the less information RPI will bring on the state of θ. To illustrate this, think of σ_ε approaching infinity. Performing below or above the median will then be completely decided by the error term, which means that no further knowledge on θ will be transferred at all. If no useful information is to be obtained by RPI, no second period investing will emerge.

Finally, consider that RPI of this kind is anticipated on by employees. An employee with a negative signal is confronted with the choice to either act as if RPI is not expected or choose a first period investment choice that allows any feasible positive second period investing. If an employee learns to be performing below average, the optimal investment choice would be $l_{2,i} = \rho_i \theta_L + (1 - \rho_i)\theta_H - l_{1,i}$ as long as the total is positive. Filling in this second period choice in the objective function and weighing the objective function by the ex ante probabilities of the two possible scenarios of s_A and s_U yields the following objective function:

$$U_i(l_{1,i}) = (\rho - \frac{1}{2})\left[[2l_{1,i}]\theta_L - \frac{1}{2}(l_{1,i})^2\right] + \left(1\frac{1}{2} - \rho\right)\left[(l_{1,i} + \rho_i\theta_L + (1-\rho_i)\theta_H)\rho_i\theta_L + (l_{1,i} + \rho_i\theta_L + (1-\rho_i)\theta_H)(1-\rho_i)\theta_H - \frac{1}{2}(\rho_i\theta_L + (1-\rho_i)\theta_H)^2\right]$$

Optimizing yields $l_{1,i} = 2\theta_L + \frac{\left(1\frac{1}{2}-\rho\right)}{(\rho-\frac{1}{2})}[\rho_i\theta_L + (1-\rho_i)\theta_H]$ which clearly conflicts with the restriction put on $l_{2,i}$. We can therefore say that a scenario without any second period investment is optimal for an employee with a negative signal, even if he anticipates on RPI of this less refined kind.

In general, we have seen that less refined information reduces the extent to which a good inference on the state of θ can be made. This in turn may inhibit the surfacing of any learning dynamic regardless of expectations on RPI.

V. Employer

An important player in our model has so far been left un-modeled. The employer – the principal in a principal-agent context – has interests at stake in the investment decisions of his employees. Moreover, he has the tools to influence the decisions of the agents by various means as he is the one securing and sharing RPI. An important question is what timing is in the employer's interest and whether it is recommendable to raise expectations on a possible RPI introduction. That is, whether the employer prefers an unexpected RPI introduction or an expected RPI introduction. Let us extend the base model by allowing for an employer who may decide on the timing of RPI and whether RPI will be introduced at all. As RPI can only be given after some performance has been measured we slightly alter the dynamic of the model by introducing a third period. The choice of the employer now comprises whether RPI will be given after the first period, after the second period or not at all. We allow employees to anticipate on the employer's behavior. That is, we extend on section IV.B rather than on section IV.A. Determining the objective function of the employer is rather arbitrary, but nonetheless potentially of importance. Therefore we will elaborate on two different options. First we look into a productivity maximizing employer (A); second we will discuss an employer who internalizes wage considerations (B). The last subpart discusses an employer who tries to use his RPI announcement strategically (C). But let us first analyze behavior of employees in this new framework.

Employee payoffs are now given by $U_i(l_{1,i}, l_{2,i}, l_{3,i}) = (3l_{1,i} + 2l_{2,i} + l_{3,i})\theta + \varepsilon_i - \frac{1}{2}(l_{1,i} + l_{2,i} + l_{3,i})^2$. Accordingly, $y_{1,i} = l_{1,i}\theta + \varepsilon_i$, $y_{2,i} = (l_{1,i} + l_{2,i})\theta$ and $y_{3,i} = (l_{1,i} + l_{2,i} + l_{3,i})\theta$. If employees are not expecting any RPI at all, they undertake all their learning activities in the first period. Formally, $l_{1,i} = 3E(\theta)$ and $l_{2,i} = l_{3,i} = 0$. Expected utility comes at $E(U_i|s_H) = 9[\rho\theta_H + (1-\rho)\theta_L]^2 - \frac{1}{2}(3[\rho\theta_H + (1-\rho)\theta_L])^2$ and $E(U_i|s_L) = 9[\rho\theta_L + (1-\rho)\theta_H]^2 - \frac{1}{2}(3[\rho\theta_L + (1-\rho)\theta_H])^2$.

In expectation of RPI after the first period, employees with a negative signal take into account two scenarios after the first period. They either learn that their negative signal has been correct, resulting in $l_{2,i} = l_{3,i} = 0$, or incorrect, leaving $l_{2,i} = 2\theta_H - l_{1,i} - l_{3,i}$ and $l_{3,i} = 0$ under the restriction that $l_{2,i}$ is non-negative. Filling in the second and third period investment choices and weighing both scenarios by their respective probabilities gives the following objective function:

$$U_i(l_{1,i}) = \rho\left[3l_{1,i}\theta_L - \frac{1}{2}(l_{1,i})^2\right] + (1-\rho)\left[(l_{1,i} + 4\theta_H)\theta_H - \frac{1}{2}(2\theta_H)^2\right]$$

Optimizing yields $l_{1,i} = 3\theta_L + \frac{(1-\rho)}{\rho}\theta_H$. To fulfill the restriction on $l_{2,i}$, we assume that $\theta_H \geq \frac{3\rho}{(3\rho-1)}\theta_L$. Interestingly, under these new conditions second period investment is not unthinkable anymore for employees with a positive signal either. Consider the following objective function:

$$U_i(l_{1,i}) = (1-\rho)\left[3l_{1,i}\theta_L - \frac{1}{2}(l_{1,i})^2\right] + \rho\left[(l_{1,i} + 4\theta_H)\theta_H - \frac{1}{2}(2\theta_H)^2\right]$$

Optimizing yields $l_{1,i} = 3\theta_L + \frac{\rho}{(1-\rho)}\theta_H$. Fulfilling the restriction on $l_{2,i}$ gives rise to the following restrictions: $\rho \leq \frac{2}{3}$ and, if satisfied, $\theta_H \geq \frac{3(1-\rho)}{(2-3\rho)}\theta_L$. The intuition behind the first restriction is that second period investing can only be present when the employee only loosely takes into account his positive signal. The second restriction stipulates that even if the employee gets only a loosely informative signal, the difference between θ_H and θ_L needs to be substantial enough. Needless to say, if any of the two restrictions is not met the employee acts as if no RPI is given. Utility conditional on signal is now given by $E(U_i|s_L) = \rho\left[3(3\theta_L + \frac{(1-\rho)}{\rho}\theta_H)\theta_L - \frac{1}{2}(3\theta_L + \frac{(1-\rho)}{\rho}\theta_H)^2\right] + (1-\rho)\left[\left((3\theta_L + \frac{(1-\rho)}{\rho}\theta_H) + 4\theta_H\right)\theta_H - \frac{1}{2}(2\theta_H)^2\right]$ for an employee with a low signal. For an employee with a high signal, expected utility is as follows for $\rho \leq \frac{2}{3}$ and $\rho \geq \frac{2}{3}$ respectively:

$$E(U_i|s_H, \rho \leq \frac{2}{3}) = (1-\rho)\left[3(3\theta_L + \frac{\rho}{(1-\rho)}\theta_H)\theta_L - \frac{1}{2}(3\theta_L + \frac{\rho}{(1-\rho)}\theta_H)^2\right] + \rho\left[\left((3\theta_L + \frac{\rho}{(1-\rho)}\theta_H) + 4\theta_H\right)\theta_H - \frac{1}{2}(2\theta_H)^2\right]$$

$$E(U_i|s_H, \rho \geq \frac{2}{3}) = 9[\rho\theta_H + (1-\rho)\theta_L]^2 - \frac{1}{2}(3[\rho\theta_H + (1-\rho)\theta_L])^2$$

In expectation of RPI after the second period, employees with a negative signal take into account two scenarios. They either learn that their signal has been correct or not, resulting in $l_{2,i} = l_{3,i} = 0$ and $l_{2,i} = 0$, $l_{3,i} = \theta_H - l_{1,i} - l_{2,i}$ respectively. Again, under the restriction that the total has to be non-negative. The objective function is now as follows:

$$U_i(l_{1,i}) = \rho\left[3l_{1,i}\theta_L - \frac{1}{2}(l_{1,i})^2\right] + (1-\rho)\left[(2l_{1,i} + \theta_H)\theta_H - \frac{1}{2}(\theta_H)^2\right]$$

Optimizing yields $l_{1,i} = 3\theta_L + \frac{(1-\rho)}{\rho}2\theta_H$. To satisfy the restriction we again have to make two assumptions. First of all $\rho \geq \frac{2}{3}$, regardless of the difference between θ_H and θ_L, and also $\theta_H \geq \frac{3\rho}{(3\rho-2)}\theta_L$. The intuition behind the first restriction is that learning activities undertaken in the third period are only contributing a relatively small amount to utility, which means that it is not worthwhile if θ_H is already somewhat expected upon regardless of its height in respect to θ_L. The second restriction stipulates that even if the realization of θ_H is not really taken into account (i.e. ρ is relatively high), it still needs to be substantially higher than θ_L as, for instance, compared to the restriction imposed in order to establish second period investment. Again, because third period investment contributes less than first or second period investment. Expected utility is now given by $E(U_i|s_H) = 9[\rho\theta_H + (1-\rho)\theta_L]^2 - \frac{1}{2}(3[\rho\theta_H + (1-\rho)\theta_L])^2$ for employees with a high signal. Employees with a negative signal have the following expectations for $\rho \geq \frac{2}{3}$ and $\rho \leq \frac{2}{3}$ respectively: $E(U_i|s_L, \rho \geq \frac{2}{3}) = \rho\left[3(3\theta_L + \frac{(1-\rho)}{\rho}2\theta_H)\theta_L -\right.$

$\frac{1}{2}(3\theta_L + \frac{(1-\rho)}{\rho} 2\theta_H)^2\Big] + (1-\rho)\Big[\Big(2(3\theta_L + \frac{(1-\rho)}{\rho} 2\theta_H) + \theta_H\Big)\theta_H - \frac{1}{2}(\theta_H)^2\Big]$ and $E(U_i|s_L, \rho \leq \frac{2}{3}) = 9[\rho\theta_L + (1-\rho)\theta_H]^2 - \frac{1}{2}(3[\rho\theta_L + (1-\rho)\theta_H])^2$.

A. Production maximizing

Consider that the employer is production maximizing. Thus, $U_E = \sum_{i=0}^{n} y_{T,i}$ where $y_{T,i}$ is now defined as $y_{T,i} \equiv y_{1,i} + y_{2,i} + y_{3,i}$. The moment the employer decides on RPI, we assume he assigns an ex ante probability of ½ to either θ_H and θ_L. That is, he receives no signal like the employees do. This is a plausible assumption if it takes time to prepare the provision of RPI, so that the employer has to decide in an early stage when θ is not yet known. Alternatively, the employer may simply just not be aware about what drives employee performance to the same extent as the employees themselves are aware of that. The choice he makes is then between the following options, normalized to one employee, where we make a distinction between the case of $\rho \leq \frac{2}{3}$ and $\rho \geq \frac{2}{3}$.

$$E(U_E|no\ RPI) = \frac{1}{2}\theta_H\big[[\rho 9(\rho\theta_H + (1-\rho)\theta_L)] + [(1-\rho)9(\rho\theta_L + (1-\rho)\theta_H)]\big] + \frac{1}{2}\theta_L\big[[\rho 9(\rho\theta_L + (1-\rho)\theta_H)] + [(1-\rho)9(\rho\theta_H + (1-\rho)\theta_L)]\big]$$

$$E\left(U_E \,\middle|\, RPI\ after\ first\ period, \rho \leq \frac{2}{3}\right) = \frac{1}{2}\theta_H\left[\left[\rho((3\theta_L + \frac{\rho}{(1-\rho)}\theta_H) + 4\theta_H)\right] + \left[(1-\rho)\left((3\theta_L + \frac{(1-\rho)}{\rho}\theta_H) + 4\theta_H\right)\right]\right] + \frac{1}{2}\theta_L\left[\left[\rho 3(3\theta_L + \frac{(1-\rho)}{\rho}\theta_H)\right] + \left[(1-\rho)3(3\theta_L + \frac{\rho}{(1-\rho)}\theta_H)\right]\right]$$

$$E\left(U_E \,\middle|\, RPI\ after\ first\ period, \rho \geq \frac{2}{3}\right) = \frac{1}{2}\theta_H\left[[\rho 9(\rho\theta_H + (1-\rho)\theta_L)] + \left[(1-\rho)\left((3\theta_L + \frac{(1-\rho)}{\rho}\theta_H) + 4\theta_H\right)\right]\right] + \frac{1}{2}\theta_L\left[\left[\rho 3(3\theta_L + \frac{(1-\rho)}{\rho}\theta_H)\right] + [(1-\rho)9(\rho\theta_H + (1-\rho)\theta_L)]\right]$$

$$E\left(U_E \,\middle|\, RPI\ after\ second\ period, \rho \leq \frac{2}{3}\right) = \frac{1}{2}\theta_H\big[[\rho 9(\rho\theta_H + (1-\rho)\theta_L)] + [(1-\rho)9(\rho\theta_L + (1-\rho)\theta_H)]\big] + \frac{1}{2}\theta_L\big[[\rho 9(\rho\theta_L + (1-\rho)\theta_H)] + [(1-\rho)9(\rho\theta_H + (1-\rho)\theta_L)]\big]$$

$$E(U_E \left| RPI\ after\ second\ period, \rho \geq \frac{2}{3}\right)$$

$$= \frac{1}{2}\theta_H \left[[\rho 9(\rho\theta_H + (1-\rho)\theta_L)] + \left[(1-\rho)[2(3\theta_L + \frac{(1-\rho)}{\rho}2\theta_H) + \theta_H] \right] \right]$$

$$+ \frac{1}{2}\theta_L \left[\left[\rho 3(3\theta_L + \frac{(1-\rho)}{\rho}2\theta_H) \right] + [(1-\rho)9(\rho\theta_H + (1-\rho)\theta_L)] \right]$$

If $\rho \leq \frac{2}{3}$ and all other restrictions are met, the option of RPI after the first period is always optimal. The intuition behind this is that the difference between θ_H and θ_L that is necessary to satisfy the restrictions at the same time makes the under investing in the first period less consequential compared to additional investing when θ turns out to be high. To see this, imagine that $\theta_L = 0$. If this would be the case, any productivity would only emerge when $\theta = \theta_H$. Which means in turn that the employer is only interested in having 'appropriate' employee behavior when the high θ arises and couldn't care less about any sense of under investing when θ turns out to be low. To establish this 'appropriate' behavior an employer would be very willing to provide the employee with RPI. If $\rho \geq \frac{2}{3}$, again it is optimal to provide RPI after the first period. Providing it after the second period is then still better than providing no RPI at all.

B. Internalizing wage costs

Consider now the employer to be internalizing wage considerations. Wages are paid both via a fixed (w_i) and a flexible component (this case is developed more generally by Sappington 1991). As stated earlier, employees completely internalize their productivity. In other terms, the flexible wage component exactly equals productivity and therefore does not show directly in the employer's utility function. Formally, $U_E = \sum_{i=0}^{n} -w_i$. In order to maximize utility, the employer likes to set the fixed wage component as low as possible. We do not restrict w_i directly, but we do assume that employees are able to reject a contract if this will give them a negative expected utility. That is, employees have an outside option that will yield them a utility of zero. The timing is as follows:

1) An employer offers a contract consisting of w_i.

2) An employee either accepts or rejects the contract.

3) Nature draws θ.

4) Each employee observes his private signal.

5) Each employee chooses l_1.

6) Etc.

It is clear that the employee's expected utility is largest when he receives as much information as early as possible (as he always has the option to ignore it). It is therefore clearly also in the employer's interest to provide RPI in a quick fashion as this will allow him to extract the largest surplus from the employees. The expected utility of the employer, normalized to one employee, will then amount to

$U_E = -w_i = \frac{1}{2}E(U_i|s_L, RPI\ after\ first\ period) + \frac{1}{2}E(U_i|s_H,\ RPI\ after\ first\ period)$.

Correspondingly, the ex ante expected utility of the employees now drops to zero exactly.

C. Strategic RPI denunciation

In the previous parts, we assumed the employer had no problem agreeing or committing to the sharing of RPI. This section devotes attention to the case where the employer tries to use his announcement strategically.

The first and most straightforward case is the one of the employer who internalizes wage considerations. As the employee is residual claimant, the employer has no direct benefit of production. All the employer cares about is that the employee agrees to the lowest possible fixed wage. In order to do so, the employer will promise to deliver RPI after the first period and, as we assume no costs of RPI, has no reason to not do so indeed.

The interesting case is the one where the employer merely maximizes production. Intuitively, the employer would like to see an as high as possible first period investment choice and, if applicable, see a lot of further investments too. In order to equip the employee with the possibility to engage in further investments the employer would like to give RPI after the first period, but in order to strengthen first period investment he would rather not say so. Instead, he would like to nullify expectations on the sharing of any RPI. The extent to which this will succeed depends on the extent to which employees are aware of the strategic considerations of the employer.

We can place this issue of awareness also in a broader context that goes beyond our theoretical framework. Just how trustworthy is shared RPI? If every employee is given information from which he or she infers the realization of $\theta = \theta_H$, productivity is higher. This may be in the interest of the employer. Without a penalty to this manipulation of information, the employees in question may generally disregard any RPI shared. In that case, consequential information transmission between employer and employee cannot be established. Obviously, if RPI is shared publicly and not privately, the employer has no opportunity to manipulate the ranking as all employees are able to check their own performance. This could be one way to prevent any manipulation of information on the employer's behalf. Note that this does not mean that RPI *should* be shared publicly as this may invite all kinds of difficult to predict psychological effects on the employee, which may be directly or indirectly harmful to the interests of the employer. Other ways to overcome a cheap talk problem could therefore be more appropriate. For more on this subject see the cheap talk literature (e.g. the classical paper of Crawford & Sobel (1982) or specifically the paper of Chakraborty & Harbaugh (2007) on comparative cheap talk).

Concluding this section, we can say that sharing RPI in a quick fashion seems always optimal for the employer from the perspective of our learning dynamic. Whether raising expectations on RPI is optimal depends on the specific circumstances and the characterization of the employer. In practice, neither a (merely) production maximizing employer, nor one who leaves all the residual claims with the employees is very realistic. Nor is it realistic to assume that regardless of the employer's announcement, employees will always expect RPI. Again, our learning dynamic is hardly the only performance and utility effect driven by RPI and most employees would at least be vaguely aware of this. Note also that, in comparison to the presence of a learning dynamic in itself, the optimality of sharing RPI may be more dependent on specific assumptions. Aoyagi (2010) shows in a theoretical research that under a tournament pay scheme, sharing RPI may not be optimal indeed from the perspective of the employer. The reason for this is that in this context RPI also gives information on the likelihood of attaining a reward and not just on relative productivity. Interestingly, although due to very different reasons, in a tournament context we may still detect the characteristic tradeoff between improved ex post and lower ex ante performance instigated by announced RPI (e.g. Goltsman & Mukherjee 2011).

VI. Discussion

Regarding the nature of the employee, it is clear from the model that we look upon him as risk-neutral. The employee gets exactly what he produces. This variable wage rate cannot be optimal in the presence of risk-aversion on the employee's behalf, due to the noise term ε_i that is beyond his control. If the employee would indeed be risk-averse (or at least to another degree than the employer) a standard tradeoff between insurance and incentive has to be made by the choice of the bonus rate (e.g. see Gibbons 1998). The optimal bonus rate would not be the rate as assumed in our model. However, our theoretical results are not qualitatively affected by an adjustment of the variable pay rate.

Another issue concerning the nature of the employee is that in our model employees are heterogeneous in two ways (their private signal and the error term), if not regarding their actual usefulness of learning activities (i.e. θ is the same for everyone). It would not be unreasonable to assume that the realization of the parameter θ is slightly different for employees according to their cognitive abilities. As stated in the introduction, one reason for limited guidelines on introduction procedures might exactly be this heterogeneity. Nonetheless, I think the results are largely robust to this if we place the model in a context of blue-collar workers. In this environment, the sheer cognitive ability to handle certain lessons learned (i.e. θ) seems less important than having the conviction that learning activities are useful (i.e. your private signal on θ). Even in a white-collar context where the ability to learn might be relatively more important than in a blue-collar context, the learning dynamic would remain present albeit somewhat diluted as RPI will bring less clarity on the underlying, diverging parameter θ.

Besides the nature of an employee, the number of employees is of importance in the establishment of our learning dynamic. In our model we assume an endless amount of employees, leading to a perfect performance distribution and therefore to perfect learning. In general, however, the fewer colleagues one has to compare oneself to, the less certain one is of the exact value of the productivity parameter. The smaller the amount of comparable employees, the less learning would take place and the less likely it is for any second period investing to emerge.

As argued before, we believe this learning dynamic is equally applicable to other pay contexts as well. The piece-rate pay context is most convenient as it provides a direct incentive to act upon extra knowledge which may enhance performance. Expecting this incentive to be present in other pay contexts – RPE or a fixed wage – is reasonable, as some intrinsic motivation or career concerns will usually be present. Nonetheless, these motivations may not be as clear cut or working with the same force.

Finally, our model is in line with empirical findings in its hint towards a positive performance effect due to an RPI introduction. However, the empirical literature does not generally support a larger performance increase among low performing individuals compared to top performers. I would argue that this is not troublesome. Although tentative, it could very well be that performance increases across the performance distribution may be driven by very different factors. A learning dynamic could be part of a syndicate of factors at the lower part of the performance distribution, while psychological factors are the main drivers of performance increases at the top. Additional empirical research on this subject would be needed to make definite statements. Although disentangling these effects empirically might be difficult to accomplish in practice, there comes one reasonable starting point in mind. Most psychological effects would arguably drive performance effects through an increase or decrease in effort (as the perceived marginal benefit of effort changes through shame or pride). A learning dynamic could be detected by a change in the activities this effort is allocated to. In other terms, a performance effect through psychological mechanisms and through a learning dynamic would respectively be more quantitative and more qualitative in nature. In any sense, in order to be able to detect a learning dynamic an experiment would therefore need to have a certain minimum time frame. As argued before, psychological effects may change performance instantly exactly because it would mainly affect effort quantitatively. A learning dynamic could then be detected via the observation of long terms effects in general and the differences therein between top and bottom performers in particular. Finally (and obviously), the tasks that subjects within an experiment undertake should be of such a nature that it allows for any conceivable learning. A simple and familiar laboratory task such as adding numbers does not seem appropriate for this. To be more clear I will give one brief, but vivid example of an experiment that I would see fit (admittedly not in a professional context).

In your local grocery store, especially around the end of a working day, shoppers line up to pay at one of the cash registers. Clearly, most people have a comparable goal when they choose a line: minimize waiting time. As your waiting time will decrease (i.e. pay increases) on average when your line-picking tactics (i.e. productivity) improve, we can interpret this as a pay-for-performance context. Now let's say that yours and others' average waiting time is recorded and adjusted for the amount of people inside the store at the moment you arrive at any cash register. If this could be recorded for a substantial period, an interesting RPI experiment may present itself. I would assume that most people would be quite surprised if they (unexpectedly) learn to spend say 15 or 20 percent more seconds in line compared to an average shopper. Suddenly the shopper in question may be struck with an awareness that picking a line smartly is actually not as straightforward as it seems. The next time he or she shops, he or she may more intensely observe line picking behavior of others. For instance, the shopper could learn that it is not just a matter of picking the line that has progressed most towards the exit as not all registers are situated along one straight line. Or that one should not only observe the number of people in line, but could also form an estimate on how filled the shopping carts of these people are. Or that some lines actually serve more than one cash register (as is the case in the author's grocery store), which means that these lines are often more attractive than they at first appear to be. Clearly, if any, it will be the 'underperforming' shoppers who will search for a new line picking-tactic. I do not expect to identify psychological or social issues such as pride, shame or status as determinants for a better future line-picking performance, especially when RPI is private and not expected to be given again. Nor do I expect shoppers to actually spend consistently more time (i.e. effort) on the actual line-picking choice in the long run. At first they may do so indeed, but I would argue that this can more appropriately be interpreted as an investment in learning activities. To complete the translation to our model, I would say most people would not have consciously engaged in any line-picking learning activities before. That is, a first period investment choice would be more or less zero (at least for the 'underperformers').

VII. Conclusion

This article shows that the introduction of relative performance information may increase employee performance through a learning dynamic. Employees may infer from RPI that better work routines exist and may reengage in learning activities in order to discover these work routines. Although this intuition has been approached by others, this is to my best knowledge the first article that tries to capture this learning dynamic in a theoretical framework. We have also seen that there are different effects when employees do or do not expect to receive RPI at some intermediate stage. The performance effect of the RPI introduction is qualitatively the same, yet the announcement of RPI may trigger a reluctance to search for an optimal work routine ex ante. If we abstract from psychological issues, both an expected and unexpected RPI introduction are in the employees' interest, as employees always have the option to ignore it. Whether it is in the employer's interest to announce RPI depends on how we characterize the employer. The actual introduction of RPI, however, is in our framework firmly in the employer's interest and should be proceeded with rather sooner than later. We have also seen that the level refinement of RPI is crucial in the establishment of our learning dynamic. More detailed information brings a better inference on why different employees exhibit diverging performances. Finally, an important feature of our model lies in its prediction that the learning dynamic should be most prevalent among bottom performers. This can be a starting point for further empirical research on the matter.

VIII. References

Aoyagi, Masaki. 2010, *Information feedback in a dynamic tournament,* Games and Economic Behavior Vol. 70 (2): 242-260.

Azmat , Ghazala & Nagore Iriberri. 2010, *The importance of relative performance feedback information: Evidence from a natural experiment using high school students*, Journal of Public Economics Vol. 94: 435-452.

Barankay, Iwan. 2011, *Gender differences in productivity responses to performance rankings: Evidence from a randomized workplace experiment*, Working paper, The Wharton School, University of Pennsylvania.

Barankay, Iwan. 2011, *Rankings and Social Tournaments: Evidence from a Crowd-Sourcing Experiment*, Working paper, The Wharton School, University of Pennsylvania.

Beck, Hall & Seta, John. 1980, *The effects of frequency of feedback on a simple coaction task*, Journal of Personality and Social Psychology, Vol. 38 (1): 75-80.

Bénabou, Roland & Jean Tirole. 2003, *Intrinsic and Extrinsic Motivation,* Review of Economic Studies, Vol. 70: 489-520.

Blanes I Vidal, Jordi & Mareike Nossol. 2011, *Tournaments without prizes: evidence from personnel records*, Management Science Vol 57 (10): 1721-1736.

Bradler, Christiane., Robert Dur, Susanne Neckermann & Arjan Non. 2011, *Employee Recognition and Perfomance – A Field Experiment*, Working paper.

Chakraborty, Archisman & Rick Harbaugh. 2007, *Comparative cheap talk*, Journal of Economic Theory, Vol. 132: 70-94.

Charness, G. & P. Kuhn. 2007, *Does pay inequality affect worker effort? Experimental evidence*, Journal of Labour Economics, Vol. 25 (4): 693-725.

Crawford, Vincent & Joel Sobel. 1982, *Strategic Information Transmission*, Econometrica, Vol. 50 (6): 1431-1451.

Crutzen, Benoit., Otto Swank & Bauke Visser. 2007, *Confidence management: on interpersonal comparisons in teams,* Tinbergen Institute Discussion Paper Series.

Dye, Ronald. 1992, *Relative Performance Evaluation and Project Selection*, Journal of Accounting Research, Vol. 30 (1): 27-52.

Ederer, Florian. 2010, *Feedback and motivation in dynamic tournaments,* Journal of Economics & Management Strategy, Vol. 19 (3): 733-769.

Eriksson, Tor., Anders Poulsen & Marie Claire Villeval. 2009, *Feedback and Incentives: Experimental evidence*, Labour economics, Vol. 16: 679-688.

Falk, A. & A. Ichino. 2006, *Clean evidence peer pressure*, Journal of Labor Economics Vol. 24 (1): 39-57.

Festinger, Leon. 1954, *A Theory of Social Comparison Processes*, Human Relations Vol. 7: 117-140.

Gibbons, Robert & Kevin Murphy. 1991, *Relative Performance Evaluation for Chief Executive Officers*, NBER Working paper No. 2944.

Gibbons, Robert. 1998, *Incentives in Organizations*, NBER Working Paper 6695.

Goltsman, Maria & Arijit Mukherjee. 2011, *Interim Performance Feedback in Multistage Tournaments: The Optimality of Partial Disclosure*, Journal of Labor Economics, Vol.29 (2): 229-265.

Hannan, Rebecca., Ranjani Krishnan & Drew Newman. 2008, *The Effects of Disseminating Relative Performance Feedback in Tournament versus Individual Performance Compensation Plans*, The Accounting Review Vol. 83: 893-913.

Kluger, Avraham & Angelo DeNisi. 1996, *The Effects of Feedback Interventions on Performance: A Historical Review, a Meta-Analysis, and a Preliminary Feedback Intervention Theory*, Psychological Bulletin Vol. 119 (2): 254-284.

Kuhnen, Camelia & Agnieszka Tymula. 2011, *Feedback, Self-Esteem, and Performance in Organizations*, Articles in Advance, Management Science: 1-20.

Lazear, E. & Sherwin Rosen. 1981, *Rank-Order Tournaments as Optimum Labor Contracts*, The Journal of Political Economy Vol. 89 (5): 841-864.

Luttmer, E. 2005, *Neighbors as negatives: Relative earnings and well-being*, Quarterly Journal of Economics, Vol. 120 (3): 963-1002.

Mayer, Richard. 1998, *Cognitive, metacognitive and motivational aspects of problem-solving,* Instructional Science, Vol. 26: 49-63.

Milkovich, G. & J. Newman. 1996, *Compensation*, McGraw-Hill.

Sappington, David. 1991, *Incentives in Principal – Agent Relationships*, Journal of Economic Perspectives Vol. 5 (2): 45-66.

Neckermann, Susanne., Cueni Reto & Bruno Frey. 2010, *Awards at work*, IEW – Working Papers iewwp411, Institute for Empirical Research in Economics – University of Zurich.

Swank, Otto & Bauke Visser. 2006, *Motivating through Delegating Tasks or Giving Attention*, Journal of Law, Economics and Organization Vol. 23 (3): 731-742.

Tesser, A. & J. Campbell. 1980, *Self-definition: The impact of the relative performance and similarity of others*, Social Psychology Quarterly Vol. 43: 341-347.

Tran, Ahn & Richard Zeckhauser. 2009, *Rank as an incentive*, HKS Faculty Research Working Paper Series RWP09-019, JFK School of Government, Harvard University.