MANAGEMENT SCIENCE
Vol. 00, No. 0, Xxxxx 0000, pp. 000–000
ISSN 0025-1909 | EISSN 1526-5501 | 00 | 0000 | 0001

INFORMS
DOI 10.1287/xxxx.0000.0000

Tournaments without Prizes: Evidence from Personnel Records*

Jordi Blanes i Vidal
Managerial Economics and Strategy Group, London School of Economics
Houghton Street, London WC2A 2AE, United Kingdom
Tel: +44-207-955-6041, j.blanes-i-vidal@lse.ac.uk.

Mareike Nossol

We use a quasi-experimental research design to study the effect of giving workers feedback on their relative performance. The setting is a firm in which workers are paid piece rates and where, for exogenous reasons, management begins to reveal to workers their relative position in the distribution of pay and productivity. We find that merely providing this information leads to a large and long-lasting increase in productivity that is *costless* to the firm. Our findings are consistent with the interpretation that workers' incipient concerns about their relative standing are activated by information about how they are performing relative to others.

Key words: Tournaments, Relative Concerns, Status Concerns, Relative Performance Feedback, Relative Performance Evaluation

History: This paper was first submitted on XXXX and has been with the authors for XXXX years for XXXX revisions.

1. Introduction

Does receiving feedback on relative performance make workers strive for success or does it lead instead to dissatisfaction and loss of motivation? Does this response depend on whether the comparison with other workers is favorable or unfavorable? In this paper we use a unique dataset from a firm level quasi-experiment to investigate these issues.

An important body of work has studied the use of relative performance from the perspective of explicit incentive theory (Lazear and Rosen, 1981). According to tournament theory, an agent's performance can provide a useful benchmark upon which other agents' performance can be evaluated and rewarded. The notion that competing with others for monetary rewards can have motivation

* We thank David de Meza, Daniel Ferreira, Thomas Kittsteiner, Ignacio Palacios-Huerta and Yona Rubinstein for insightful comments. We are grateful to all those involved in providing the data. Joy Blundell and Vera Radyugina provided excellent research assistance. This paper has been screened to ensure that no confidential information is revealed. All errors remain our own.

effects, has received empirical support in various economic settings (e.g. Knoeber and Thurman, 1994; Eriksson, 1999; Casas-Arce and Martinez-Jerez, 2009).

Alternatively, it may be that receiving feedback on relative performance has *in itself* substantial consequences for employee satisfaction, motivation and productivity (Kluger and DeNisi, 1996). Consider, for instance, a setting in which workers are compensated using piece rates, so that better performers also receive higher compensation. If workers display relative concerns (Festinger, 1954; Frank, 1985; Buunk and Gibbons, 2007), information on relative performance will affect well-being and effort decisions even when pay is independent of relative performance per se.

Understanding this direct motivation effect is important for several reasons. At a general level it can contribute to our knowledge of the interlinks between preferences, behaviour and the information available to individuals. Secondly, it can shed light on the additional, probably unintended, effects of explicit incentive schemes, like tournaments, that rely on comparisons across workers. Thirdly, it has direct implications for optimal communication policies. In organizations where information on relative standing can be computed easily, there is clearly potential for a judicious communication policy to increase productivity at little or no cost to the firm.

In this paper we use a field quasi-experiment to study the direct effect of informing workers about their own position in the productivity and pay distribution. There is a scarcity of evidence using real workers in real economic conditions to study the effects of relative performance feedback. In addition to the usual difficulties in obtaining data from organizations, there is the additional difficulty that, when relative performance is communicated in firms, it typically has explicit or implicit monetary consequences. For instance, a worker being informed that he has performed better than his colleagues may reasonably conclude that he has a chance of receiving a promotion to a better paid position, which could translate into higher effort. This makes it difficult to disentangle the direct effect of relative performance feedback from the effect of the monetary reward to which such feedback is often linked.

In this paper we take advantange of a unique micro level data set, constructed from a German wholesale and retail organization. In the main warehouse of this organization, workers were being paid piece rates. From some point onwards, however, they began to receive *private* information on their own position in the distribution of pay and productivity. For reasons that we describe in detail below, this relative performance information exercise (henceforth RPIE) did not have (and was known to not have) any type of compensation or career consequences. Its introduction was triggered by exogenous reasons and was not part of any wider managerial policy.

Although not a randomized trial, a key feature of our dataset, namely the availability of information on individual *daily* performance, provides the basis for a quasi-experimental research design.

We make use of this rich information to estimate the productivity increase following the introduction of the RPIE. We find the immediate average increase to be 6.8% and find no evidence that this effect vanishes over time. The increase is consistent across the distribution of productivity. We also show that this long term increase in productivity (i.e. in 'quantity') did not occur at the expense of either a long term decrease in the quality of the work produced or an increase in the quit rate among the workforce.

To understand the economic mechanism behind this increase in productivity, we exploit the differential timing in the two parts which constitute the RPIE. The information was communicated on the first day of the month, and referred to the relative performance during the month *before* the previous one. For instance, on the 1st of November workers learnt about their performance during the month of September. As a result of this timing, we can distinguish the 'kickoff moment' -the date in which workers performance first started to count towards the information that they would receive- from the 'revelation moment' -the date in which workers first learnt their relative performance. We find that productivity increases not only at the revelation moment, but *also* at the kickoff moment. The fact that productivity starts increasing even before workers actually learn any new information, casts doubt on most potential explanations of our findings, including those based on pure worker learning about absolute ability or the distribution of wages.

Our preferred interpretation suggests that the observed increase in productivity was caused by an increase in concerns about relative standing. As we discuss in detail below, our results are consistent with information about relative performance, and the anticipation of such information, leading to an increase in relative concerns that shifted the distribution of productivity to the right, as all types of workers tried to improve their relative position in terms of pay and/or productivity. Our findings therefore add to previous studies arguing that workers are affected by relative concerns, whereby increases in others' performance always has a negative effect on own utility (Luttmer, 2005; Charness and Kuhn, 2007; Clark, Frijters and Shields, 2008). It has been argued that the existence of such concerns leads to excessive levels of effort and consumption, creating a role for welfare increasing public policy (Oswald, 1997). We interpret the findings of our paper as providing evidence on the existence of 'contests for relative position' using field data.

1.1. Related Literature

Psychologists and management scholars have long recognized that workers are motivated by more than monetary rewards (Etzioni, 1971; Deci, 1975). Individuals exert effort at work partly because of interest or enjoyment in the task itself or because completing the task provides a feeling of accomplishment or self-worth. Typically, the presence of extrinsic rewards is believed to have the effect of 'crowding out' intrinsic motivation (Kohn 1993; Deci, Koestner and Ryan, 1999). At the

same time, Deci (1972) and Anderson, Manoogian and Reznick (1976) find that (positive) feedback increases intrinsic motivation. In our setting, we argue that relative performance feedback does indeed increase relative concerns and therefore intrinsic motivation. The existence of performance pay most likely reinforced the strengthening of relative concerns following the RPIE and therefore served to enhance the increase in intrinsic motivation.

Feedback interventions can sometimes lead to a decrease in performance (Kluger and DeNisi 1996). A substantial body of work in human resource management and applied psychology has therefore analyzed the circumstances that make job feedback enhance satisfaction and performance (Nadler 1979, Latham and Wexley 1981, Moore and Klein 2008). Since in most organizations evaluations are mainly subjective, research efforts have often focused on which organizational actors should contribute to the evaluation, what format should the rating take, and how best to present it to be evaluated worker (for an overview see, for instance, London 2003). In our firm, performance was easily quantifiable and therefore it was possible to be both objective and detailed in the information presented to workers[1]. Furthermore, the objective nature of performance and the clarity of the goal implied that workers dissatisfied with their relative performance could easily understand that a reduction in the gap between expectations and feedback was possible through the exertion of higher effort (Kluger and DeNisi 1996, Bandura and Cervone 1983). This fact may contribute to explain the positive effect on productivity that we document.

In economics, Eriksson, Poulsen and Villeval (2008) find no effect of providing relative performance feedback to laboratory subjects compensated via piece rates. Kuhnen and Tymula (2010) develop a theoretical model of self-esteem, study subjects in a laboratory flat wage setting and find that relative performance feedback has both ex ante (i.e. for agents anticipating it) effects and ex post (for agents receiving it) effort effects. Their result echoes our finding that the RPIE affects performance both at the 'kickoff stage' and at the 'revelation stage'.

Two recent experimental papers try to separately measure the importance of cardinal versus ordinal considerations. Clark, Masclet and Villeval (2010) find that agents' rank in the income distribution affects effort while, controlling for rank, relative income does not. Dohmen et al. (2010) observe laboratory subjects using functional magnetic resonance imaging. They find that both relative income and rank affect regions of the brain associated with the processing of basic rewards. In our firm, the RPIE consisted of both cardinal and ordinal information and unfortunately their relative importance cannot be separately identified.

[1] 'Constructive' feedback is regarded as more likely to lead to increased performance (Baron 1988). London (2003) argues that feedback is constructive when 'it offers concrete information that can be used', 'it is clear and easily understood', and 'it is interpreted similarly by the source and recipient'. All these characteristics clearly apply to the RPIE that we study.

Several papers study relative performance feedback using field data. In a context where workers are compensated according to team performance, Delfgaauw et al. (2009) find that information on relative performance leads to higher productivity, while Bandiera et al. (2009) find the opposite. In a schooling context, Azmat and Iriberri (2009) and Tran and Zeckhauser (2009) argue that learning about relative performance leads to higher student effort.

Lastly, this paper contributes to the literature on tournament schemes by documenting a new margin through which competition operates, namely the 'pure contest for relative position' which arises when individuals derive utility directly from outperforming others in terms of performance or pay. Importantly, incentives along this margin are free to the principal: the 6.8% average increase in productivity that we observe in our setting was achieved at *no cost* to the firm, other than the negligible cost of communicating with workers. Thus, even tournaments *without* prizes can have large incentive effects.

1.2. Outline of the Paper

Section 2 discusses the institutional setting in which the RPIE occured. Section 3 presents and interprets our main results. In Section 4 we examine other effects of the RPIE. Section 5 concludes.

2. Institutional Setting

2.1. Workplace Conditions and Compensation

We study an intervention occuring in the second half of 2001 in the main warehouse of a German wholesale and retail organization. The company's wholesale division handles a large number of food and non-food products, which are sold to independent supermarkets as well as the company's own retail division.

During our sample period, the firm employed around 65 workers to perform the core task of the warehouse. This task involves picking up customer orders, assembling the requested products, packing them onto a trolley, and moving the trolley to the goods-out area of the warehouse[2]. Work is strictly individual, and opportunities to observe each other's work are limited. Worker's salary comes in three parts: a fixed base salary, which is the same for every worker[3], a 'productivity' or 'quantity' performance component, and a 'quality' performance component. The structure of the compensation scheme was constant between August 1999 and August 2002. The levels of the piece rates have however changed over time.

[2] The allocation of orders to workers took place following a 'cab rank system', so that a worker finishing an order would be matched with the order at the top of the list, and so on. In regressions available upon request we found that a worker's productivity in a month could not predict the characteristics of the orders that, on average, he would receive the following month.

[3] In this firm there is a one-to-one mapping between performance and pay. For this reason, we will not be able to disentangle whether concerns about relative standing refer to performance or to pay.

The 'quantity' performance component has two parts, with their corresponding linear piece rates: the first part depends on the number of orders completed and the second on the number of goods dispatched. The number of goods dispatched is obviously the product of the number of orders and the average number of goods per order. In 2001 the 'quantity' performance component represented around 25% of the average worker total compensation.

The 'quality' performance component depends on the number of wrongly dispatched goods and, in practice, is much less important. When a customer is sent the wrong good and he complains to the firm, this fact is recorded and tracked back to the individual worker in charge of that particular good. A small linear discount is applied to a worker's total compensation when his monthly mistake rate exceeds a certain threshold. In 2001 the average discount was less than 1% of total compensation.

Both performance components are paid with a one month delay. Hence in the pay slip received, say, on the 1st of November, workers learn about performance and wage per hour corresponding to the month of September, rather than October.

Our measure of productivity is the average number of goods dispatched per hour by a given worker on a given day. Since customer orders contain a varying number of goods, we always control for the average number of goods per order. Otherwise we would be overestimating the productivity of those workers that handled a few large orders rather than many small ones. Once the number of goods per order is held constant, the number of goods per hour, and the number of orders per hour are equivalent measures of productivity.

Upon joining the firm, workers face a six months probationary period, and they may not be offered a regular contract at the end of this period. However, once they survive the probationary period they are very unlikely to be fired. After two years, the restrictions imposed by German law, the fact that less productive workers are cheaper to the firm and the screening performed during the probationary period, make termination of contracts all but unheard of.

If the job that we study is unusually secure, it is also a cul-de-sac, since the skills acquired in it do not transfer easily to other positions in the firm. Among the 207 workers holding this job during the last 10 years, only 2 workers have ever been promoted. The reason for this is that there is no natural higher-level position in which the skills gained in this job provide either a good training or a good signal of future productivity. Thus, any form of career concerns is a truly unlikely source of motivation for the workers in our study.

Lastly, there is a certain amount of learning and skill in this occupation. This is because being able to identify the optimal route to gather the goods ordered by a customer leads to a significant gain in time and productivity. This ability/skill component translates into a strong serial correlation in workers' relative performance across time. During our sample period workers' rank-order position

in the productivity distribution during a certain month had a correlation of .88 with their position in the previous month.

2.2. The Relative Performance Information Exercise

In the summer of 2001 a few members of the workforce -employees are not unionised- requested from management access to information about the wage per hour being earned by the average worker. Given the low computational cost of creating and distributing this information, the firm agreed to this demand. All workers were then notified that their September performance -that is, the one included in the 1st November 2001 pay slip- would be the first one on which they would receive this information.

In addition, management decided to also communicate to workers their own rank-order position in the productivity/wage per hour distribution. The reason for this was very specific: two workers had been constantly complaining about the conditions of the job and spreading discontent among the workforce. These two workers happened to be amongst the worst performers, so management thought that privately revealing to them this fact would help to alter their behaviour[4].

The introduction of the RPIE provides the basis of a quasi-experimental research design. Note that the RPIE was first triggered by the demand of a small number of workers and then by a very specific managerial motivation which was orthogonal to the compensation scheme in place. Further, company insiders stated to us that the introduction of the RPIE did not coincide with any type of shock to the dispatching technology. The identification strategy, which we discuss further below, uses both individual daily productivity and a short time window around the introduction of the RPIE to ensure that potential pre-existing productivity trends are not confounding our empirical findings.

The information finally communicated on the 1st of November and subsequent months consisted of the minimum, maximum and average wage per hour and a worker's rank-order position in the wage per hour distribution. Workers were notified prior to the 1st of September that they would receive this information.

3. The Productivity Effects of the Relative Performance Information Exercise

Our discussion above suggests two moments at which workers' behaviour may have changed as a result of the RPIE. Firstly, on the 1st of September -the 'kickoff moment'-, workers' daily productivity first started to count towards the RPIE and they may have altered their effort in response

[4] According to company sources, these two workers were indeed embarrased to learn that they ranked at the bottom of the distribution and stopped complaining. Their contracts were not terminated nor were they disciplined or even threatened in any way, a fact that should alleviate any lingering concerns that workers' reaction may have been due to the fear of job terminations.

to this fact. If workers anticipate that they will derive utility tomorrow from learning that they outperformed their colleagues today, then increasing effort today to obtain that future utility seems like a reasonable response. Secondly, on the 1st of November -the 'revelation moment'- when workers received information about their relative performance for the first time. This may have led to a further behavioural response if receiving such information further increased the salience and importance of relative concerns.

3.1. Descriptive Analysis

We first provide graphical evidence that the RPIE was associated with a large and sudden shift in workers' productivity, and not with similar discontinuities in observable determinants of productivity.

Figure 1A shows the evolution of daily productivity, averaged across workers. It is clear from the figure that the introduction of the RPIE coincided with a sharp increase in productivity. Before September 2001 an average worker dispatched less than 145 goods per hour, whereas by November 2001 productivity had increased to around 157 goods per hour. Such increase seems out of line both with the flat or mildly positive pre-existing trend and with pre-existing fluctuations around the trend.

In Figures 1B and 1C we display the evolution of characteristics of the goods dispatched by workers. Goods can be either fresh or dry, and it is claimed by company insiders that fresh goods take longer to dispatch since the aisles where they are placed are narrower and this creates jams in the warehouse. In early 2001 a weekend shift was added, leading to a more extreme concentration of fresh and dry goods in different days of the week. However, it is clear from Figure 1B that the RPIE did not coincide with any big discontinuity in the ratio of fresh goods. Similarly, we plot the evolution of the average value of the goods dispatched by workers over time in Figure 1C. Again, no large and permanent discontinuity seems apparent from the figure.

In Figure 1D we display the evolution of average worker experience over time. Average experience trends upwards if there are no changes to personnel and instead decreases in periods, like the second part of 2000, in which the firm hires new workers. Such fluctuations are obvious in the figure, but, again, no large changes are apparent during our period of interest.

It seems difficult from Figure 1A to attribute the sharp increase during the September 2001-November 2001 period to any pre-existing trend in productivity. To ensure, however, that potential mispecification of pre-existing time trends does not confound our estimates, we focus our main analysis on a short time window comprising 6 months in total. Our main sample consists of three periods: (1) The 'baseline period' includes the two months (July and August) just before the introduction of the RPIE; (2) The 'kickoff period' includes the two months (September and October)

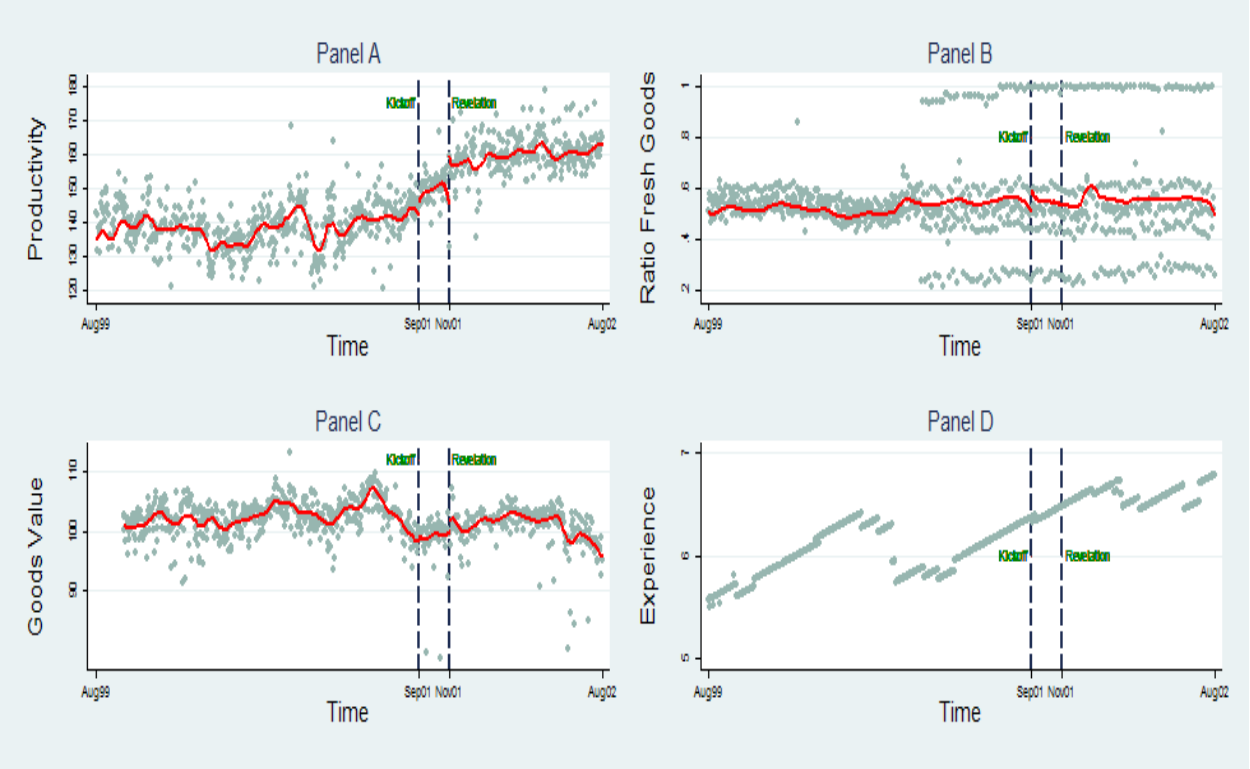

Figure 1 Evolution of Productivity and Other Variables over Time.

Note: Dependent variables are: the number of goods dispatched per hour worked (Panel A), the ratio of fresh goods over fresh+dry goods (Panel B), the average value of a good (Panel C) and worker experience (Panel D). Each dot is the (averaged across workers) dependent variable of a different working day. Vertical dashed lines display the 1st of September 2001 (kickoff date) and 1st of November 2001 (revelation date). In Panels A-C solid lines display local linear regressions (with 8 days bandwidth) computed separately before the kickoff date, after the revelation date, and in between.

during which workers did not have any new information but knew that their daily productivity would form the basis on which relative performance would be calculated; and (3) The 'revelation period' includes the two months (November and December) right after workers first learnt their relative performance.

Using this sample, Panel A of Table 1 displays the main input and output variables at firm level. In Panel B of Table 1 we display a number of worker characteristics. The average worker in our sample has been at this job for four years and is 39 years of age. More than 60% of workers in our sample were still at the job six and a half years later, which implies that the yearly separation rate was just 6%. We also display information on the rank-order performance during August 2001,

which has been constructed by us but was not computed by the firm[5].

Table 1: Descriptive Statistics

Panel A: Firm-Level Characteristics				
	(1) Mean	(2) S.D.	(3) Minimum	(4) Maximum
Goods per day	1,107.24	152.73	474.78	1561.79
Goods per hour	149.73	7.85	126.84	172.14
Average Value of Good	100	28.57	3.6	122.22
Ratio of Fresh Goods	.56	.22	.22	1
Number of Observations	152			
Panel B: Individual-Level Characteristics				
	(1) Mean	(2) S.D.	(3) Minimum	(4) Maximum
Male	.63	.48	0	1
Experience (in years)	3.99	4.06	.50	23.88
Age (in years)	38.75	8.35	21.27	58.48
Left 31/12/2007?	.38	.49	0	1
Ranking August 2001	34.01	18.35	3	66
Number of Observations	63			

Note: Descriptive statistics based on sample comprising 1Jul2001-28Dec2001. The number of worker-day observations included in the sample is 5,275. Panel A displays variables at firm-day level by average worker. Average Value of a Good is normalised at 100. Panel B displays variables at worker level. Sample is restricted to those workers with at least 6 months of experience in the firm.

3.2. Econometric Framework

Our basic estimating equation takes the following linear form:

$$y_{it} = \beta_k K_t + \beta_r R_t + \mathbf{X}'_{it} \cdot \delta + \mu_i + g(t) + \epsilon_{it} \tag{1}$$

where y_{it} is (the log of) worker i's productivity, measured in number of goods dispatched per hour, on day t, $K_t = 1$ in the post 1st September period and $R_t = 1$ in the post 1st November period. Our model also contains a vector of covariates $\mathbf{X}'_{it}$, including as discussed above the (log of the) average number of goods per customer order by worker i on day t. We allow for individual time-invariant unobserved effects, μ_i, to account for the possibility that individual workers may differ across time periods.

[5] We restrict our study to workers with more than 6 months of experience. We do this for several reasons. Firstly, the returns to experience are both very steep and very concave among workers in their first six months. As a result, controlling for a linear trend may not adequately capture the evolution of productivity for these workers, even in such a short time horizon as the one used here. Secondly, these workers are in their probationary period and may have different incentives from those faced by the rest of the workforce. In particular, we cannot rule out the possibility that they may perceive the information on their relative position as a signal on the likelihood that their contract will be renewed. By focusing on workers unlikely to be making those considerations, we will be able to argue that pure concerns for relative position lie behind their reaction to the RPIE. This restriction reduces our sample by just 3 workers.

Our identification strategy relies on the comparison of the same workers over time. Relying on such variation implies that we cannot be completely non-restrictive in accounting for time effects. We instead allow productivity to evolve smoothly over time through the parametric function $g(t)$. We show however that our results are robust to the specification of $g(t)$ as well as to the period over which $g(t)$ is calculated. ϵ_{it} captures iid person-specific idiosyncratic shocks.

To account for correlation of the observations from the same date, we adjust the standard errors by allowing for arbitrary variance-covariance matrix within each date across individuals. In practice, this does not have much effect on the standard errors.

Our parameters of interest are β_k and β_r. The identifying assumption in estimating these parameters from equation (1) is that, without the introduction of the RPIE, productivity would not have deviated from the trend $g(t)$.

3.3. Baseline Results

Table 2 displays our baseline results. Column (1) shows that productivity was 2.7% higher in the kickoff period and 5.9% higher in the revelation period, relative to the baseline period.

Table 2: Baseline Results

	(1) Unconditional	(2) Time Trend	(3) Other Covariates
Kickoff	.027***	.033***	.028***
	(.005)	(.010)	(.006)
Revelation	.032***	.038***	.040***
	(.006)	(.010)	(.007)
(Log) Goods per Order	.308***	.308***	.336***
	(.011)	(.011)	(.006)
Time Trend		-.0001	-.0001*
		(.0001)	(.00007)
(Log) Goods Piece Rate			-.045
			(.104)
Ratio Fresh Goods			.013
			(.010)
Average Value of Goods			-.174***
(In Euros)			(.022)
Day of Week Dummies	No	No	Yes
Worker Experience Dummies	No	No	Yes
Worker Fixed Effects	No	No	Yes
Adjusted R^2	.22	.22	.77
Number of Observations	5,275	5,275	5,275

Note: Dependent variable is (log of) number of goods dispatched per hour worked. Data comprises 1Jul2001-28Dec2001. Kickoff is a dummy variable taking value 1 in the post-1Sep2001 period. Revelation is a dummy variable taking value 1 in the post-1Nov2001 period. Worker experience dummies capture the quarter of experience of the worker. Sample is restricted to those workers with at least 6 months of experience in the firm. Robust standard errors allow for cluster at day level.

The fit of the model is very large even in this first specification with just three explanatory variables. The reason is that the number of goods per order is a very strong determinant of the number of goods dispatched per hour, as workers can dispatch more goods when they do not have to make constant trips to the goods-out area of the warehouse and to pick up new orders. This fact underlines the importance of controlling for this variable in every specification.

Column (2) adds a linear time trend to the specification. The estimated coefficients of interest are very similar in magnitude to those of Column (1) and the estimated linear trend is not statistically different from zero. Thus, unobserved time-varying factors correlated with productivity are unlikely to be biasing our estimates.

In Column (3) we add extra covariates, day of week dummies, experience dummies and worker fixed effects. We find the average effect of kickstarting the RPIE to be 2.8%, while the average effect of actually revealing the relative performance information is 4.0%. The cumulative effect is 6.8%.

The coefficients on the covariates are interesting in their own right. Consider the level of the goods piece rate, which the firm altered in every month of our sample. These changes were by very small percentages (usually 1 or 2 percent), so it is plausible that the workforce barely (if at all) reacted to them. The negative and statistically insignificant coefficient on the piece rate indicates indeed that workers did not seem to work harder in months in which the piece rate was marginally higher.

We find however that workers were more productive when the goods dispatched were cheaper. Company insiders suggested that this may be either because cheaper goods are smaller in size or because they are more common and workers are better aware of their location.

3.4. Placebo RPIE and Robustness Checks

As the RPIE was introduced at the same time for all workers, identification of its effect on productivity arises from a comparison over time of the same worker. The estimated effects are therefore biased upward to the extent that they capture factors that cause productivity to rise independently of the RPIE and that are not captured by the linear trend. Table 3 presents a number of robustness checks to address this and other concerns.

To judge whether our findings are due to seasonality, we repeat the estimation of (**??**) substituting our baseline sample from 2001 with an equivalent sample from the previous year. This new 'placebo' sample spans between 1Jul2000 and 31Dec2000 and consists of 3,836 observations. We define a 'placebo kickoff' dummy to be equal to one after 1Sep2000 and a 'placebo revelation' dummy to be equal to one after 1Nov2000. We include worker fixed effects, a time trend, experience dummies and the other controls. Column (2) of Table 3 suggests that our main results were not due to seasonality.

The estimated coefficients are much smaller in magnitude, and are either not statistically significant or weakly significant but with a negative sign.

Table 3: Placebo RPIE and Robustness Checks

	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline Results	Placebo	12 Months Window	Quadratic Trend	Pre-Treatment Trend	Pre-Treatment Effects
Kickoff	.028***		.020***	.029***	.027***	.032***
	(.006)		(.005)	(.006)	(.005)	(.005)
Placebo Kickoff		.008				
		(.006)				
Revelation	.040***		.035***	.049***	.027***	.040***
	(.006)		(.003)	(.004)	(.005)	(.006)
Placebo Revelation		-.011*				
		(.006)				
Table 2 Covariates	Yes	Yes	Yes	Yes	Yes	Yes
Standard Errors	Day Cluster	Day Cluster	Day Cluster	Day Cluster	Boot-strapped	Boot-strapped
Adjusted R^2	.74	.71	.74	.74	.74	.74
Window Months	6	6	12	12	6	6
Observations	5,275	3,836	10,176	10,176	5,275	5,275

Note: Dependent variable is (log of) number of goods dispatched per hour worked. Column (1) is baseline regression (i.e. Column (3) from Table 2). Column (2) data comprises 1Jul2000-31Dec2000 and the placebo kickoff and placebo revelation dummies take value 1 in the post-1Sep2000 and post-1Nov2000 periods respectively. In columns (3) and (4) data comprises 1Apr2001-31Mar2002. Column (4) includes a quadratic time trend. In Column (5) the linear time trend is estimated exclusively on the baseline period (July and August 2001). In Column (6) the linear time trend, the worker fixed effects and the effects of the covariates are estimated exclusively on the baseline period (July and August 2001). Robust standard errors allow for cluster at day level unless specified.

We next examine the robustness of our estimates to the sample width by estimating (**??**) on a sample comprising 12 months in total. As Column (3) of Table 3 shows, the estimated coefficients are still statistically significant and only slightly different in magnitude.

Secondly, we study whether our results are robust to allowing the parametric time trend $g(t)$ to take different formats. We first allow $g(t)$ to take a quadratic form. Since a problem of severe collinearity between t and t^2 arises in our baseline time window of 6 months, we use a window of 12 months. Column (4) of Table 3 shows that our main findings are indeed robust to allowing $g(t)$ to take a more flexible form.

In our benchmark specification from Column (3) in Table 2 the linear trend is estimated using both pre-treatment and treatment observations. Alternatively, we may want to estimate the trend *exclusively* using pre-treatment information. This allows us to evaluate the effect of the RPIE above

or below what we would expect productivity to be given its pre-treatment trend. To do this, we estimate the empirical model

$$y_{it} = \mathbf{X}'_{it} \cdot \delta + \mu_i + \theta(1 - K_t) \cdot t + \epsilon_{it} \tag{2}$$

and use the estimated coefficients to project the expected value of productivity $\hat{y}_{it}$ given the estimated pre-treatment trend $\hat{\theta}$. We then compute the difference between actual and predicted productivity $\hat{v}_{it} = y_{it} - \hat{y}_{it}$ and estimate

$$\hat{v}_{it} = \beta_k K_t + \beta_r R_t + \xi_{it} \tag{3}$$

The estimated coefficients $\hat{\beta}_k$ and $\hat{\beta}_r$ are displayed in Column (5) of Table 3. The coefficients are very similar to those of our benchmark specification from Table 2, suggesting that our results are robust to the form of the parametric time trend.

In Column (6) we use only pre-treatment information to estimate the individual fixed effects μ_i and the effects of other covariates δ as well as the linear trend θ. Again, we find that the estimated coefficients are statistically significant and similar in magnitude[6].

3.5. Heterogeneous Effects

Table 2 shows that the *average* effect of the RPIE is 6.8%. The effect may however be different for different types of workers, perhaps even negative for some.

In Table 4 we split the workforce on the basis of observable characteristics and run our benchmark specification (**??**) separately for different groups of workers. Panel A of Table 4 shows that the reactions of males and females to the RPIE are very similar. While the estimated coefficients are slightly different in magnitude, these differences are not statistically significant. In Panel B we find that only workers in the middle and top third in terms of experience react to the kickoff of the RPIE. On the other hand, we find no differences in the revelation coefficient for workers with different levels of experience. In Panel C we use the (undisclosed) ranking from August 2001 to split the sample by pre-existing level of performance. We find weak evidence that workers in the bottom third increase their productivity more at the kickoff stage, while workers in the top third increase in productivity more at the revelation stage. The overall effect of the RPIE is very similar for workers with different pre-existing levels of performance.

Instead of aggregating workers on the basis of their observable characteristics, we can study separately the reaction of every individual worker to the RPIE. To do this, we ran 58 separate

[6] One important limitation of our study is the absence of a control group. To partially remedy this drawback, we examined the performance of a comparable German food trader, Metro Group. Productivity measures at the warehouse worker level were unfortunately unavailable. Instead, we studied the evolution over time of a quarterly measure of operating profits, EBITDA. We found that this measure of performance was not characterized by any significant discontinuity in the second part of 2001 (details available upon request).

worker-level regressions equivalent to (**??**) and computed the sums of the estimated kickoff and revelation coefficients. While there was substantial heterogeneity across workers, the reaction to the RPIE was be positive for the overwhelming majority of workers (i.e. 51 out of 58). Furthermore, of the 19 workers for whom the computed overall effect was statistically different from zero at the 10% level, 18 had a positive sign, while only 1 had a negative sign.

Table 4: Decomposition by Gender Experience and Performance (August Ranking)

Panel A: Decomposition by Gender

	(1) Males	(2) Females	(1)-(2) P-value
Kickoff	.023***	.038***	.28
	(.008)	(.011)	
Revelation	.047***	.026**	.17
	(.008)	(.012)	
Worker F.E.	Yes	Yes	
Covariates	Yes	Yes	
Obs.	3,307	1,968	

Panel B: Decomposition by Experience

	(1) Least Experience	(2) Middle Experience	(3) Most Experience	(1)-(2) P-value	(1)-(3) P-value	(2)-(3) P-value
Kickoff	-.010	.029**	.037***	.02	.00	.41
	(.012)	(.010)	(.010)			
Revelation	.045***	.046***	.036***	.98	.55	.56
	(.013)	(.012)	(.011)			
Worker F.E.	Yes	Yes	Yes			
Covariates	Yes	Yes	Yes			
Obs.	1,484	1,833	1,907			

Panel C: Decomposition by Performance (August Ranking)

	(1) Bottom Ranked	(2) Medium Ranked	(3)Top Ranked	(1)-(2) P-value	(1)-(3) P-value	(2)-(3) P-value
Kickoff	.052***	.015	.019*	.02	.04	.78
	(.012)	(.009)	(.011)			
Revelation	.035**	.028**	.053***	.72	.25	.11
	(.014)	(.011)	(.011)			
Worker F.E.	Yes	Yes	Yes			
Covariates	Yes	Yes	Yes			
Obs.	1,553	1,860	1,862			

Note: Dependent variable is (log of) number of goods produced per hour worked. Data comprises 1Jul2001-28Dec2001. Sample is restricted to those workers with at least 6 months of experience in the firm. Covariates as in Column (3) of Table 2, except in Panel B, which excludes experience quarter dummies. Robust standard errors allow for cluster at day level.

We argue in Section 3.7 that two findings from this Section are unlikely to be compatible with basic models of learning about ability or the distribution of performance. These are: (a) very experienced workers do not react less at the revelation stage than less experienced workers, and (b) the effect of the RPIE can be regarded as positive for almost all workers.

3.6. Effects on Inequality

We now study whether the RPIE affected the amount of inequality across workers. Figure 2 plots the kernel distribution of productivity for August, September and November 2001. It seems that the RPIE shifted the distribution to the right, without noticeable effects on the dispersion of productivity. To confirm this finding formally, we computed the coefficient of variation for the productivity variable for each day of our baseline sample. We found that this measure of dispersion was not different on average following the introduction of the RPIE (details available upon request from the authors). Thus, static inequality did not change.

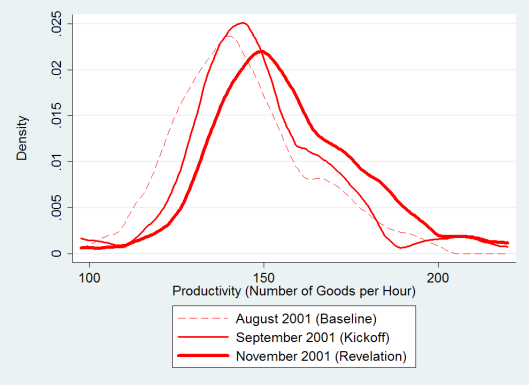

Figure 2 Distribution of Productivity over Time.
Note: Figure displays the kernel densities of the productivity variable for three selected months: August 2001 (the last baseline month), September 2001 (the first kickoff month) and November 2001 (the first revelation month). Productivity is defined in terms of goods dispatched per hour.

We can also study inequality from a dynamic perspective by examining the likelihood that a worker at a position in the performance distribution in a particular month repeats position the following month. In Figure 3 we plot workers' rank position in terms of wage per hour across consecutive months, for each month in our baseline sample. There is significant persistence[7]. Overall, it does not seem as if persistence in rank position varied following the introduction of the RPIE. We confirmed this by regressing a worker's rank position in a month on his position in the previous month and interacting this variable with the kickoff and revelation variables. We found that

[7] We do not find the existence of strong serial correlation in rank incompatible with the notion that agents exert more effort when their relative concerns have stronger. This is because (a) relative concerns could be cardinal as well as ordinal, (b) rank mobility over time can be high in the presence of strong serial correlation provided that the number of workers is large, and (c) increased relative concerns could be affecting the *marginal* return to effort even while having no effect on *equilibrium* rank positions.

the coefficients associated with these interactions were not statistically different from zero (details available upon request from the authors).

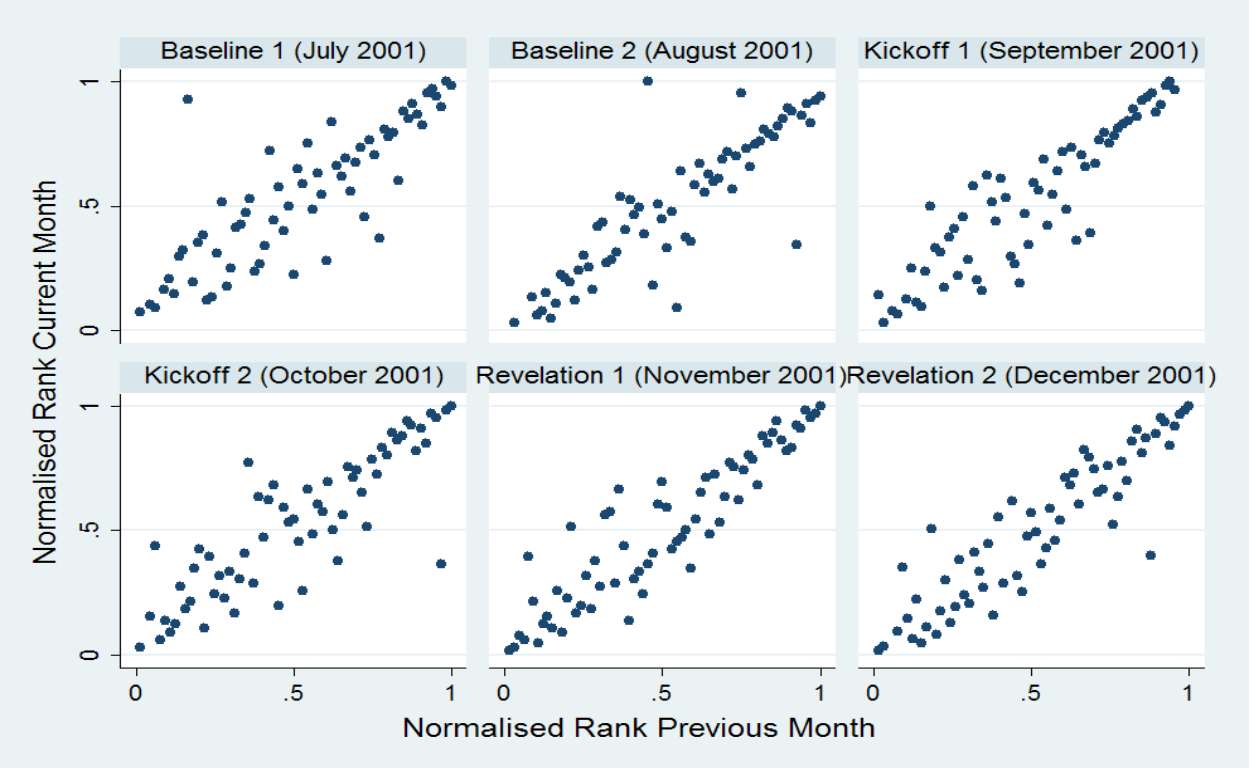

Figure 3 Relation between Ranks in Two Consecutive Months

Note: This Figure presents a scatter plot capturing the relation between the rank (in terms of the wage per hour) obtained by an individual in a month and the rank obtained by the same individual in the previous month. The ranks are all normalized to one by dividing by the total number of workers working in the firm on that month. Six different scatter plots are presented for each of the months comprising our baseline sample.

To summarize, the RPIE led to an increase in average productivity and wages while leaving both static and dynamic disparities between workers essentially unchanged.

3.7. Interpretation and Alternative Explanations

To summarise the main findings: (a) the effect of the RPIE is on average positive both at the kickoff stage and at the revelation stage (Figure 1A and Table 2), (b) individual regressions show that the effect is positive for a large proportion of workers and it is negative for just one worker (Section 3.5), (c) the increase in productivity at the kickoff stage is larger for workers with more experience (Panel B of Table 4) and for workers at the bottom of the performance distribution (Panel C of Table 4) (d) the overall increase in productivity is however statistically equal for workers with different observable characteristics (Table 4).

We now discuss the potential explanations of our findings.

3.7.1. Preferred Interpretation: Increase in Relative Concerns We regard the evidence presented here as consistent with the notion that workers' incipient concerns about relative standing were activated by the RPIE. This interpretation is consistent with the documented productivity increase at the kickoff stage. By increasing effort at that stage a worker would have increased his chances of outperforming his colleagues; a fact that, on being learnt later, would generate a utility increase. This explanation is also consistent with the revelation of relative performance further increasing the salience of these relative concerns, and leading to the additional increase in productivity documented in Table 2. Lastly, the existence of concerns for relative pay is perfectly compatible with more experienced workers reacting more strongly at the kickoff stage, as it is natural to think that these concerns should be stronger for individuals that have been members of the workforce for a longer period of time (Frank, 1985).

What *type* of relative concerns can explain our findings? Consider first relative concerns of the 'inequity aversion' type (Fehr and Schmidt, 1999; Fehr and Gächter, 2000). These preferences would be consistent with workers at the bottom of the distribution increasing their effort by a large amount, with workers at the top not working harder and even decreasing their effort to allow the others to catch up. Yet, this is the opposite of what we find in Section 3.5.

We instead regard relative concerns whereby increases in others' performance always has a negative effect on own utility (Frank, 1984; Fershtman et al., 2005; Moldovanu et al., 2007), as consistent with our findings. As these authors discuss, an increase in these concerns translates into everybody increasing their level of effort and nobody decreasing it, as we find in Section 3.5.

We now discuss alternative explanations.

3.7.2. Alternative Explanation: Career Concerns Could concerns about either termination of employment or a potential promotion present an explanation of our findings? As discussed in Section 2, the likelihood of both terminations and promotions in our firm is close to zero. Working conditions and base pay are fixed and identical for every worker, regardless of performance. Furthermore, workers were explicitly told that the relative performance information would *not* be used in any way by management. Thus, we do not regard career concerns as a valid explanation.

Career concerns interpretations would also find it difficult to explain why productivity did not decrease for almost any worker in the firm. Consider for instance the incentives of a weak performer who, previously to the introduction of the RPIE, is unaware of this fact and exerts some extra effort with the hope of ranking among the top and meriting a promotion. We would expect that, upon learning his mediocrity, he would regard a promotion as virtually impossible and decrease his efforts to attain it. Yet we do not find this to be the case for almost any worker. Similarly, very able workers fearing a potential termination should have put *less* effort, not more, after learning that they ranked at the top of the distribution. Again, this is the opposite of what we find.

3.7.3. Alternative Explanation: Learning about Own *Absolute* Ability One potential explanation of our findings is that the RPIE may have helped workers to learn about their own ability for the job, and this may have led to a readjustment in the chosen level of effort. One potential benchmark for this argument is Ertac (2006), who presents a theoretical model in which effort and ability are strategic complements in a worker's production function and workers are uncertain about their own ability. Past production can help workers learn about their own ability but is subject to an unobserved firm wide productivity shock that makes learning imperfect. In this setting, Ertac (2006) shows that information on how much other workers produced can help a worker to learn about their own ability. For instance, a worker with low productivity in a month will revise expectations about his own ability upwards after observing that others also did badly, as his low productivity is likely to be the result of an unfavorable firm wide productivity shock rather than the result of low ability.

We see two problems with this explanation. The main problem is that any interpretation based exclusively on learning cannot explain the existence of a positive productivity effect at the kickoff stage, in which workers do not learn anything.

An additional problem is that the predictions of the model outlined above do not seem consistent with our empirical findings. In the model outlined above the arrival of new information makes some workers revise their expected ability upwards and exert more effort, while other workers revise it downwards and exert less effort[8]. This is not what we find in the data. We found in Section 3.5 that there is almost no heterogeneity in the sign of the reaction to the RPIE: almost all workers react by increasing their effort[9].

Nevertheless, it may be that a different model of learning about absolute ability can explain the revelation stage effect.

3.7.4. Alternative Explanation: Learning about the Distribution of Performance One potential explanation of our findings is that the RPIE may have helped workers to learn about a productivity parameter that is common to all workers, and, again, this may have led to a readjustment in the chosen level of effort. For instance, a worker with bad luck in a particular month may, in the absence of RPIE, wrongly believe that the returns to effort in the company are not very high. However, by observing the higher production of his co-workers, he will update his belief about the returns to effort upwards, which may lead to more effort.

[8] To gain intuition, consider a situation in which all workers have the same prior expected ability and learn the average level of production of their co-workers. Roughly speaking, workers producing more (respectively, less) than the average will update their ability upwards (downwards) and exert more (less) effort upon learning this fact.

[9] A second prediction of Ertac (2006) is that workers less uncertain about their own ability should react less to the information conveyed by the RPIE. Again, this is inconsistent with our findings. We find in Panel B of Table 4 that highly experienced workers (presumably well-informed about their own ability) do not react less at the revelation stage that highly inexperienced workers.

We believe that this explanation is unlikely. Any learning explanation is inconsistent with the documented positive reaction at the kickoff stage, in which there can be no learning. Secondly, we would expect that some workers should react to the arrival of new information by increasing their effort *while others should decrease it*. We find instead in Section 3.5 that the sign of the reaction to the RPIE is positive for almost all workers.

Again, it may be that a different model of learning can explain the revelation stage effect.

3.7.5. Alternative Explanation: Certification to Outside Employers It could be argued that the relative performance information included in the pay slip could have been used by workers as a verifiable signal of high productivity towards prospective outside employers. Under this hypothesis, the interest in improving this signal could have generated the increase in effort and productivity that we have documented.

While we cannot completely rule out this notion, we find it unlikely for two reasons. Firstly, we would not expect workers at the bottom third of the distribution to plan on disclosing such a signal to outside employers, since little can be gained by admitting one's relative incompetence[10]. As a result, we would expect their effort to remain unchanged, contrary to our findings. Secondly, the availability of such certification should have led to an increase in the quit rate among the workforce (and especially among the best performers), as information about workers' ability becomes less asymmetric. We show in Section 4 that this is not the case.

4. Other Effects of the Relative Performance Information Exercise: Mistakes Rate and Separation Rate

In this Section we briefly investigate whether the productivity increase documented in the previous Section was accompanied by a permanent change in: (a) the quality of the work done; and (b) the separation rate.

4.1. Effect on the Mistakes Rate

The 'quality' component represents a tiny determinant of total compensation (on average the discount in compensation due to mistakes is less than 1%). Yet, it is worth studying whether the RPIE was accompanied by a permanent increase (or occured at the expense of a permanent decrease) in the quality of work done. Unfortunately, in our firm information on mistakes is only recorded at monthly intervals, rather than at daily frequencies. Since this makes it more difficult to study the instantaneous behavioural response to the introduction of the RPIE, our findings in this Section should be regarded as suggestive, rather than conclusive.

[10] This is of course provided that (a) the most that such workers can aspire to is relative mediocrity and that (b) outside employers are unaware that our firm provides this information, two notions that we find highly plausible.

We display in Figure 4 the evolution of the average mistakes rate over time. The introduction of the RPIE did not coincide with a dramatic discontinuity in the evolution of the mistakes rate. The mistakes rate is on average higher during the 12 months previous to the RPIE than during the months after it, but this declining trend is apparent well before the kickoff date. The month of September 2001 (the first month after the kickoff date) is associated with a mistakes rate that is much higher than that of the two previous months (July and August). However, in October, November and December the mistakes rate was back to normal pre-treatment levels.

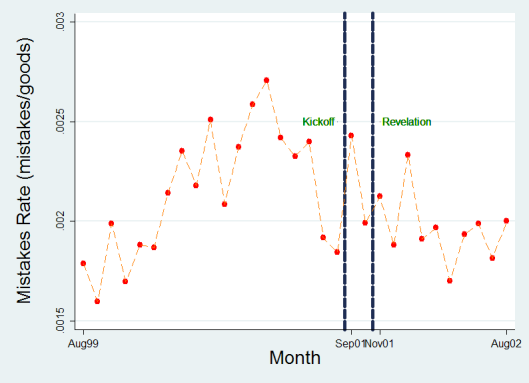

Figure 4 Evolution of Mistakes Rate over Time

Note: Dependent variable is (log of) number of mistakes per good dispatched. Vertical dashed lines displaying the 1st of September 2001 (kickoff date) and 1st of November 2001 (revelation date) displayed. Each dot is the (averaged across workers) mistake rate of a different month.

We now use our baseline 6 months time window to estimate the following panel data specification:

$$M_{it} = d_t + \mathbf{X}'_{it} \cdot \delta + \mu_i + \upsilon_{it} \tag{4}$$

where M_{it} is the (log of) mistakes rate of worker i on month t, d_t is a dummy variable capturing month t, μ_i is an individual time-invariant unobserved effect, $\mathbf{X}'_{it}$ is a vector of covariates similar to that of Column (3) of Table 2, and υ_{it} are iid person-specific idiosyncratic shocks.

The main parameters of interest are $d_{September}$, $d_{October}$, $d_{November}$, and $d_{December}$, where the first two parameters capture the kickoff period and the second two parameters capture the revelation period. The parameter d_{July} is omitted, and the parameter d_{August} captures the pre-treatment time trend in the mistakes rate.

Table 5 report OLS estimates of (**??**). In Column (1) we do not include worker fixed effects or extra covariates. We find that in August the mistakes rate is smaller than in July, whereas in September it is larger. From October onwards we do not find any statistical difference in the mistakes rate relative to the omitted month of July. Column (2) shows that this finding is robust to conditioning on individual fixed effects and a vector of covariates similar to that of Column (3)

Table 2.

Table 5: Evidence on the Mistakes Rate

	(1) Unconditional	(2) Controls
July (Baseline Month 1)	Omitted	Omitted
August (Baseline Month 2)	-.167***	-.224*
	(.076)	(.112)
September (Kickoff Month 1)	.192***	.194**
	(.069)	(.081)
October (Kickoff Month 2)	.017	.114
	(.072)	(.083)
November (Revelation Month 1)	.036	.143
	(.063)	(.118)
December (Revelation Month 2)	-.056	-.016
	(.080)	(.112)
Worker Fixed Effects	No	Yes
Table 2 Covariates	No	Yes
Adjusted R^2	.03	.70
Number of Observations	332	332

Note: Dependent variable is (log of) number of mistakes per good dispatched. Data comprises 1Jul2001-28Dec2001. July and August are the two months previous to the introduction of the RPIE. September and October are the two months comprising the kickoff period. November and December are the first two months of the revelation period. Sample is restricted to those workers with at least 6 months of experience in the firm. Robust standard errors allow for cluster at worker level. Covariates as in Column (3) of Table 2 (except for the linear time trend and the (log of) the goods piece rate). The covariates are monthly averages.

Thus, the results in this Section are consistent with a transitory (i.e. in the month of September) increase in the mistakes rate following the kickoff date. There is however no evidence to suggest that the RPIE led to a permanent change in the quality of work produced.

4.2. Effect on Workers' Separation Rate

In this Section we investigate whether the separation rate can be regarded as being different before and after the introduction of the RPIE. We also study whether workers of different ability reacted differently to this introduction. To do this, we estimate a Cox semiparametric proportional hazard model, which allows for a fully flexible, nonparametric baseline hazard. In our model, the covariates shift the baseline hazard proportionally, through the function

$$\theta_{it} = \phi R_t + \eta A_i + \rho(R_t \cdot A_i) + \mathbf{X}_i' \cdot \delta \tag{5}$$

where $R_t = 1$ if the first date of the quarter of tenure t occurs after the revelation moment, A_i is a measure of individual ability[11] and $\mathbf{X}_i'$ is a vector of worker characteristics.

[11] To compute this measure, we run a regression where productivity (i.e. the number of goods per hour) is the dependent variable, and day fixed effects, worker fixed effects and quarter of experience dummies are the explanatory

We estimate this model exclusively on workers who joined the firm after the end of the probationary period (i.e. for whom the quarter of tenure t is at least 3). By focusing on workers past their probationary period we can interpret our findings as referring to the workers' *quit* rate.

The small size of the workforce, together with the fact that separations are relatively rare events, makes it unfeasible to circumscribe the analysis to a narrow time window centered around the introduction of the RPIE. We instead use 8 years of data, between February 1999 and December 2007. Due to the large width of the sample time window, the evidence from this Section should again be regarded as suggestive, rather than conclusive.

We include in our analysis only workers who joined the firm after August 1999 (the first month for which we have observations), to avoid problems of left censoring. In total, our sample contains 127 workers, out of which 41 had left the firm by December 2007.

Table 6: Likelihood of a Separation
Before and After the Revelation of Relative Performance

Dependent Variable: =1 if Worker i Left the Firm in Quarter t			
	Cox Regression Unconditional	Cox Regression Ability	Cox Regression Other Covariates
	(1)	(2)	(3)
Revelation	-1.70***	-1.72***	-1.64***
	(.359)	(.367)	(.537)
	[.181]	[.178]	[.193]
Worker Ability		-6.24***	-7.88***
		(1.19)	(1.42)
		[.001]	[.0003]
Revelation*		-.008	.216
Worker Ability		(.491)	(.489)
		[.486]	[1.24]
Worker Covariates	No	No	Yes
Number of Observations	719	719	719
Number of Subjects	127	127	127
Number of Separations	41	41	41

Note: Revelation takes value 1 if the first date of quarter t occurs after 1Nov2001. Worker ability is the estimated worker fixed effect of a regression that has productivity as a dependent variable and includes experience dummies and day fixed effects, as well as worker fixed effects, on the right hand side. Worker covariates include gender, marital status, age entry in the firm and a linear cohort effect (the calendar year in which the worker joined the firm). Data comprises 1Aug1999-31Dec2007. The sample is restricted to workers joining the firm after August 1999. The sample is restricted to workers surviving the probationary period (i.e. the first 6 months of employment). Robust standard errors in parentheses in Column (1). Robust bootstrapped standard errors in parentheses in Columns (2) and (3). Hazard ratios in square brackets.

Table 6 reports the results from the proportional Cox hazard model (5). In Column (1) we observe that the likelihood of a separation is lower in the period after the introduction of the RPIE than

variables. Our measure of individual ability, A_i is the estimated worker fixed effect from this regression. The estimated standard errors are bootstrapped in equation (5) to account for the fact that A_i is generated, rather than observed.

in the period before. The hazard ratio is .181, indicating that the likelihood of a separation in the period after the RPIE is just 18% of the likelihood in the period before. In Column (2) we introduce our measure of worker's ability both in levels and interacted with R_t. We find that workers with a higher value of our ability measure are less likely to leave the firm, suggesting that our measure does indeed capture an important dimension of worker heterogeneity. However, the difference in separation rates across workers with different abilities does not change after the introduction of the RPIE. Lastly, we find in Column (3) that our findings are robust to the introduction of a number of worker characteristics, including gender, marital status, age at entry in the firm and a linear cohort effect (captured by the calendar year in which the worker joined the firm).

To summarise, we do not find any evidence indicating that workers are more likely to quit the firm after the introduction of the RPIE. Subject to the caveats mentioned above, our evidence in fact suggests the opposite.

4.3. Relation with Monetary Incentives and Effect on Profits

We now relate the effects of the RPIE to the effects of monetary incentives. To estimate how workers respond to an increase in the goods piece rate, we use data between January 2002 and December 2007[12]. We regress (log of) individual productivity on the (log of) the goods piece rate, the (log of) the orders piece rate (which the firm changed during this time period), a cubic time trend, individual fixed effects and the time varying controls from Table 2. Our estimate of the effect of increasing the goods piece rate implies an elasticity of 43%. This elasticity implies that the 6.8% achieved by the introduction of the RPIE would be equivalent to an increase in the goods piece rate of around $15.8\% = \frac{6.8\%}{.43}$. Thus, the non-pecuniary incentives provided by the RPIE are reasonably large not only in absolute terms but also relative to the effects achieved by monetary incentives.

We can also calculate the potential reduction in labor costs induced by the RPIE. Since the existing 65 workers started producing on average 6.8% more goods per hour, the firm was in principle able to produce the same amount of goods in the same amount of time with 4.20 (6.5%) less workers. How would this have translated into labor cost savings? To calculate this, we need to take into account the fact that the performance component of workers compensation is independent of the number of workers in the firm. The performance component is given by the product of the piece rate(s) and total production, which is exogenously given by the demand for the firm products. The firm could, however, have saved on the fixed component of workers compensation, which represented around 75% of the average worker pay during this period. Producing the same amount with 6.5%

[12] We use this longer time period because the fact that the firm changed the level of the piece rate in repeated occasions and by significant amounts allows us to obtain good estimates of the effects of monetary incentives.

less workers therefore implies potentially saving $4.9\% = 6.8\% \times 75\%$ in terms of labor costs. In our opinion, this represents a economically significant, although not huge, amount[13].

5. Concluding Remarks

We have provided evidence that social comparisons affect human behaviour independently of the monetary incentives to which these comparisons are usually linked.

In our firm, simply receiving information about relative standing -and anticipating such future information- leads to an increase in individual productivity that is statistically and economically significant. We have found no evidence of this increase being purely transitory. Apart from a transitory (i.e. one month) slight decrease in the quality of work done, we have found no evidence of workers reacting along other dimensions in ways that are detrimental to the firm. Taken together, our results suggest that making information about relative performance available is unambiguously beneficial to the firm's profitability.

We believe that the notion that workers display relative concerns provides the best possible explanation of our findings. Other interpretations fall short in one way or another, notably in explaining why workers already increase their effort at the kickoff stage.

We have been able to identify the causal effects of social comparisons thanks to unique access to the detailed personnel records of a firm that, for largely exogenous reasons, started to provide information on relative performance to workers. Our findings are based on a single firm study, and may not necessarily generalize to other settings. To understand their external validity, there are two features of our firm and work environment that we regard as particularly significant:

(1) The effects of providing information on relative productivity and pay may depend on the technology of production, on the incentive scheme in place and the nature of workers affected. In our firm, the existence of absolute performance pay implies a one-to-one mapping between performance and compensation[14]. Further, work is individual, with little scope for cooperation among workers. As a result, workers dissatisfied with their relative position have one single natural channel through which they can remedy their situation: the exertion of more effort. The effects of information on social comparisons may obviously be different in fixed-wage settings or in settings where industrial politics are potentially important (Lazear, 1989). Similarly, workers' attitudes may be critical in

[13] We refer to these savings as potential savings rather than actual savings, because no worker was actually fired as a result of the increase in productivity. The firm does not typically fire employees past the probationary period and during this time period the demand for the firm products and the size of the workforce was growing over time. The firm therefore adjusted by failing to expand the workforce at the same rate as it would have in the absence of the RPIE. Note lastly that the decrease in the separation rate found in Section 4.2 will also reduce search costs and therefore increase profits.

[14] Obviously this one-to-one mapping also implies that our results may reflect workers caring about relative productivity *per se*, as well as about relative pay. Disentangling these two possibilities falls outside the scope of this paper.

determining whether social comparisons are morale-boosting or have the opposite effect. In short, while we may find similar effects among, say, door-to-door salesmen, our results do not necessarily generalize to, for instance, corporate lawyers, since for these team-work is important and there is large scope for mutual sabotage.

(2) Perhaps more importantly, in our firm the information was disclosed in a purely *private* manner to the workers (although, of course, some may have shared that information among themselves). We are unfortunately unable to judge the extent of information sharing among workers. However, to the extent that such information sharing was not complete, we may observe different, possibly stronger, results under a public disclosure of relative standing. Falk and Ichino (2006), for instance, find that the ability to observe other workers' productivity leads to an increase in effort in a flat wage setting. While the distinction between the impact of private and public information is clearly beyond the scope of this paper, we regard it as an interesting direction for future work.

REFERENCES

Anderson, R., Manoogian, S. and Reznick, S. (1976), "The undermining and enhancing of intrinsic motivation in preschool children", *Journal of Personality and Social Psychology* 34: 915-92.

Azmat, G., and Iriberri, N. (2009), "The Importance of Relative Performance Feedback Information: Evidence from a Natural Experiment using High School Students", unpublished manuscript.

Bandiera, O., Barankay, I. and Rasul, I. (2009), "Team Incentives: Evidence from a Field Experiment", unpublished manuscript.

Bandura, A. and Cervone, D. (1983), "Self-evaluative and self-efficacy mechanisms governing the motivational effects of goal systems", *Journal of Personality and Social Psychology*, 45: 1017-1028.

Baron, R. (1988), "Negative effects of destructive criticism: Impact on conflict, self-efficacy and task performance", *Journal of Applied Psychology*, 73: 199-207.

Buunk, A. and Gibbons, F. (2007), "Social Comparison: The end of a theory and the emergence of a field", *Organizational Behavior and Human Decision Processes*, 102: 3-21.

Casas-Arce, P., and Martinez-Jerez, J. (2009), "Relative Performance Compensation, Contests and Dynamic Incentives", *Management Science*, forthcoming.

Charness, G. and Kuhn, P. (2007), "Does Pay Inequality Affect Worker Effort? Experimental Evidence", *Journal of Labor Economics*, 25(4): 693-725.

Clark, A., Frijters, P. and Shields, M. (2008), "Relative Income, Happiness and Utility: An Explanation for the Easterlin Paradox and Other Puzzles", *Journal of Economics Literature*, 46(1): 95-145.

Clark, A., Masclet, D. and Villeval, M. (2010), "Effort and Comparison Income: Experimental and Survey Evidence", *Industrial and Labor Relations Review*, 63(3): 407-426.

Deci, E. (1975), "Intrinsic Motivation", New York: Plenum Press.

Deci, E. (1972), "Effects of contingent and non-contingent rewards and controls on intrinsic motivation", *Organizational Behavior and Human Performance*, 8: 217-229..

Deci, E., Koestner, R., and Ryan, R. (1999), "A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation.", *Psychological Bulletin*, 125(6): 627-668.

Delfgaauw, J., Dur, R., Sol, J. and W. Verbeke (2009), "Tournament Incentives in the Field: Gender Differences in the Workplace", unpublished manuscript.

Dohmen, T., Falk, A., Fliessbach, K., Sunde, U. and B. Weber (2009), "Relative versus absolute income, joy of winning, and gender: Brain imaging evidence", unpublished manuscript.

Etzioni, A. (1971), "Modern Organizations", Englewood Cliffs, NJ: Prentice-Hall.

Eriksson, T. (1999), "Executive Compensation and Tournament Theory: Empirical Test on Danish Data", *Journal of Labor Economics*, 17(2): 262-280.

Eriksson, T., Poulsen, A. and Villeval, M.C. (2008), "Feedback and Incentives: Experimental Evidence", IZA Discussion Paper No. 3440.

Ertac, S. (2006), "Social Comparisons and Optimal Revelation: Theory and Experiments", unpublished manuscript.

Fehr, E. and Gächter, S. (2000), "Fairness and Retaliation: The Economics of Reciprocity", *Journal of Economic Perspectives*, 14(2): 159-181.

Fehr, E. and Schmidt, K. (1999), "A Theory of Fairness, Competition and Cooperation", *Quarterly Journal of Economics*, 114(3): 769-817.

Fershtman, C., Hvide, H. and Weiss, Y. (2005), "Cultural Diversity, Status Concerns and the Organization of Work", *Research in Labour Economics*.

Festinger, L. (1954), "A Theory of Social Comparison Processes", *Human Relations*, 1, 117-140.

Falk, A. and Ichino, A. (2006), "Clean Evidence on Peer Effects", *Journal of Labor Economics*, 24, 39-57.

Frank, R. (1984), "Are Workers Paid their Marginal Products?", *American Economic Review*, 74(4): 549-571.

Frank, R. (1985), "Choosing the Right Pond: Human Behavior and the Quest for Status", New York: Oxford University Press.

Gneezy, U., Niederle, M. and Rustichini, A. (2003), "Performance in Competitive Environments: Gender Differences", *Quarterly Journal of Economics*, 118: 1049-1074.

Kluger, A. and Denisi, A. (1996), "The Effects of Information Interventions on Performance: A Historical Review, a Meta-Analysis and a Preliminary Information Intervention Theory", *Psychological Bulletin*, 119(2): 254-284.

Knoeber, C. and Thurman, W. (1994), "Testing the Theory of Tournaments: An Empirical Analysis of Broiler Production", *Journal of Labor Economics*, 12(2): 155-179.

Kohn, A. (1993), "Punished by Rewards", Boston: Houghton-Mifflin.

Kuhnen, C. and Tymula, A. (2010), "Feedback, Self-Esteem and Performance in Organizations", unpublished manuscript.

Latham, G. and Wexley, K. (1981), "Increasing productivity through performance appraisal", Reading, MA: Addison-Wesley.

Lavy, V. (2008), "Gender Differences in Market Competitiveness in a Real Workplace: Evidence from Performance-Based Pay Tournaments among Teachers", NBER Working Paper, 14338.

Lazear, E. (1989), "Pay Equality and Industrial Politics", *Journal of Political Economy*, 97(3): 561-580.

Lazear, E. and Rosen, S. (1981), "Rank-Order Tournaments as Optimum Labor Contracts", *Journal of Political Economy*, 89(5): 841-864.

London, M. (2003), "Job Feedback: Giving, Seeking and Using Feedback for Performance Improvement", Mahway, NJ: Lawrence Erlbaum Associates.

Luttmer, E. (2005), "Neighbors as Negatives: Relative Earnings and Well-Being", *Quarterly Journal of Economics*, 120(3): 963-1002.

Main, B., O'Reilly, C. and Wade, J. (1993), "Top Executive Pay: Tournament or Teamwork?", *Journal of Labor Economics*, 11(4): 606-628.

Moore, D. and Klein, W. (2008), "Use of Absolute and Comparative Performance Feedback in Absolute and Comparative Judgments and Decisions", *Organizational Behavior and Human Decision Processes*, 107: 60-74.

Moldovanu, B., Sela, A. and Shi, X. (2007), "Contests for Status", *Journal of Political Economy*, 115(2): 338-363.

Nadler, D. (1979), "The effects of feedback on task group behavior: A review of the experimental research", *Organizational Behavior and Human Performance*, 23: 309-338..

Niederle, M. and Vesterlund, L. (2007), "Do Women Shy Away from Competition? Do Men Compete Too Much?", *Quartely Journal of Economics*, 122(3): 1067-1102.

Oswald, A. (1997), "Happiness and Economic Performance", *Economics Journal*, 107: 1815-1831.

Tran, A., and Zeckhauser, R. (2009), "Rank as an Incentive", unpublished manuscript.