Tournaments without Prizes: Evidence from Personnel Records*

Jordi Blanes i Vidal[†] Mareike Nossol

January 2009

Abstract

We use a quasi-experimental research design to study the introduction of a relative performance evaluation *without* introducing relative performance pay. The setting is a firm in which workers are paid piece rates and where, at some point, 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 permanent 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.

Keywords: Tournaments, Relative Concerns, Status Concerns, Relative Performance Evaluation.

JEL Classification: J33, M12, M52, M54.

*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. This paper has been screened to ensure that no confidential information is revealed. All errors remain our own.

[†]Corresponding author: Managerial Economics and Strategy Group, London School of Economics, Houghton Street, London WC2A 2AE, United Kingdom; Tel: +44-207-955-6041; Email: j.blanes-i-vidal@lse.ac.uk.

1 Introduction

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

Following Lazear and Rosen (1981) an important body of work has considered these questions from the perspective of explicit incentive theory. 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 explicit monetary rewards can have motivation effects, has received strong empirical support in various economic settings (e.g. Ehrenberg and Bognanno, 1990; Knoeber and Thurman, 1994; Eriksson, 1999).

Alternatively, it may be that the act of being compared to others has *in itself* substantial consequences for employee satisfaction, motivation and productivity. Psychologists (Kluger and Denisi, 1996) and management scholars (Milkovich and Newman, 1996) support this view. Economists have argued that if workers display relative concerns, information on relative performance can affect effort decisions *even* when pay is independent of relative performance (Frank, 1985b; Fershtman et al., 2005; Moldovanu et al., 2007). Furthermore, the weight that workers attach to relative concerns may itself depend on whether or not they are able to observe how others are doing¹.

Unfortunately we lack field evidence regarding the *direct* motivation effects of generating comparisons across workers. In addition to the usual difficulties in obtaining data from organizations, there is the additional difficulty that relative performance is often evaluated and communicated in firms only as part of an explicit monetary compensation scheme. This makes it difficult to disentangle the two mechanisms mentioned above.

Understanding these direct motivation effects 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.

The objective of this paper is to evaluate the causal effect on workers' performance of informing workers about their own position in the productivity and pay distribution without introducing relative

¹Frank (1985a), for instance, argues that the category of 'nonpositional goods' (i.e. goods for which relative comparisons are not important), includes goods that are not readily observed by outsiders. See also the distinction between local and global hierarchies in Frank (1985b).

performance pay. We take advantage of a unique micro level data set, constructed from a German wholesale and retail organization, to disentangle the direct effects of information on relative performance from the effects of the monetary rewards to which this information is usually linked. 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 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 also 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 6.6% and find no evidence that this effect vanishes over time. 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 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 and to an effort ‘rat race’, as employees try to improve their situation within the distribution of relative positions².

There is a dearth of field evidence showing the effects on human behaviour of relative concerns of the

²In this respect, our findings echo Landers, Rebitzer and Taylor (1996), who show that the promotion system in law firms, together with the unobservability of workers’ disutility of effort, leads to effort ‘rat races’ among associates, as they strive to signal their willingness to work long hours.

‘status’ type, whereby increases in others’ performance always has a negative effect on own utility^{3,4}. This type of relative concerns has been used to explain why, conditional on absolute income, individuals at the bottom of an income distribution report lower subjective well-being than those at the top, either within societies (Clark, Frijters and Shields, 2008) or within organizations (Brown et al., 2008). It has also 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; Hopkins and Kornienko, 2004; Frank, 2005; and Layard, 2005). To the best of our knowledge this is the first paper to provide direct evidence on the existence of ‘rat races’ or ‘contests for relative position’ using field data.

Lastly, this paper also 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 6.6% 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.

Section 2 discusses the institutional setting in which the RPIE occurred. 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 occurring 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 being 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. Work is strictly individual, and opportunities to observe each other’s work are limited. Worker’s salary comes in three

³One suggestive exception using survey data is however the finding by Neumark and Postlewaite (1998), that women’s labour supply decisions are correlated with the income of their husbands relative to the income of their sister’s husbands. For recent laboratory experiments addressing the effects of ‘status’ relative concerns on effort see Charness and Kuhn (2007) and Eriksson, Poulsen and Villeval (2008).

⁴Note that this is in sharp contrast with the extensive literature studying the effects of ‘fairness’ or ‘inequity aversion’ concerns (Fehr and Schmidt, 1999; Fehr and Gächter, 2000; Bandiera et al., 2005).

parts: a fixed base salary, which is the same for every worker⁵, a ‘productivity’ or ‘quantity’ performance component, and a ‘quality’ performance component. The structure of the compensation scheme has remained unaltered through time. The levels of the piece rates have however changed over time.

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 (the number of goods dispatched to the wrong customer, divided by the total number of goods dispatched) exceeds a certain threshold. In 2001 the average discount was less than 1% of total compensation⁶.

For technical reasons, 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, all our regressions below will control for the average number of goods per order on that day. 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 (see Table 3 for evidence that using these alternative measures yields equivalent results).

Upon joining the firm, workers face a six months probationary period, and a significant proportion of workers are not offered a regular contract at the end of this period. However, once they survive the probationary period they are very unlikely to be fired. In practice, this likelihood becomes zero after the worker has been with the firm for two years. 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

⁵In 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.

⁶While our main focus is on the quantity of work produced, we also study in Section 4 whether the mistakes rate is different before and after the introduction of the RPIE. Unfortunately, the power of this test is limited by the availability of data: while productivity is recorded at daily intervals, information on mistakes is only available at monthly intervals.

do not transfer easily to other positions in the firm. To illustrate as to how slim promotion prospects are, 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 productivity⁷. Thus, any form of career concerns is a truly unlikely source of motivation for the workers in our study⁸.

Lastly, there is a certain amount of learning and skill in this occupation. This is apparently due to the fact that being able to identify the optimal route to gather the goods ordered by a customer leads to a significant gain in time and productivity. Evidence of this is the observation that workers are still learning how to do things better even after holding the job for several years (see Figure A1 in the Appendix for a plot of how productivity increases with experience). This ability/skill component translates into a strong serial correlation in workers' relative performance across time. For example, 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 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 behaviour⁹.

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

⁷The 2 promotions that did occur were to a clerical position that requires computer literacy, a knowledge irrelevant to the job description of the workers studied here.

⁸We also discuss in Section 3.6 why our findings are anyway inconsistent with a career concerns interpretation.

⁹According to company sources, these two workers were indeed embarrassed to learn that they ranked at the bottom of the distribution and stopped complaining. Their contracts were not terminated nor were they disciplined 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 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 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. Importantly for our purposes, and as indicated earlier, 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 specific moments at which workers' behaviour may have changed as a result of the RPIE. Firstly, on the 1st of September -what we call the 'kickoff moment'-, when workers' daily productivity first started to count towards the exercise and they may have altered their effort in response to this fact. Intuitively, 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 is a reasonable response. Secondly, on the 1st of November -what we call the 'revelation moment'- when workers received information about their relative performance for the first time. Again, this may have led to a further behavioural response if receiving such information, for instance, further increased the importance of relative concerns.

3.1 Descriptive Evidence

Figure 1 shows the evolution of daily productivity, averaged across workers, over time. Two vertical lines displaying the kickoff moment and the revelation moment are also plotted. A cursory look at this figure immediately reveals 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 160 goods. This increase seems to be quantitatively quite out of line with the normal time-pattern in productivity. To be more precise and gain some intuition, note that this 2-month increase accounts for over 30% of all the increase in productivity in the 5 year period between February 1999 and February 2004.

In Figures A2-A3 of the Appendix we use 'placebo kickoffs' and 'placebo revelations' in equivalent calendar dates of 2000 and 2003 to show that Figure 1 is not simply capturing seasonal effects *We do not use 2002 as a placebo year as the firm introduced a change to the information revealed in the pay slip in the Autumn of 2002..*

It seems difficult from Figure 1 to attribute the sharp increase during the September 2001-November 2001 period to any pre-existing trend in productivity. To ensure, however, that potential misspecification of pre-existing time trends does not confound our estimates, we limit our analysis below to a short time window comprising 6 months in total (we show in Table 3 that very similar results are obtained with larger time windows). Our 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) 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.

Panel A of Table 1 displays the main input and output variables at firm level. The average number of workers in the payroll is 66, but only an average of 43 work on any given day -only a few workers are required on Sundays, for instance-. On an average day, work takes around seven and a half hours for the average worker, during which he dispatches around 150 goods per hour.

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 (as discussed in Section 2 and in the next paragraph, in practice this is the yearly *quit* rate) was just 6%. We also display information on two rank-order variables. The first one, based on performance during August 2001, has been constructed by us but was not computed by the firm. The second one, from September 2001, was the first rank-order information constructed by the firm and communicated to workers in the pay slip of the 1st of November.

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 (see Figure A1 in the Appendix for evidence on this). 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 (see Table 3 for evidence that our results are essentially unaffected by this restriction).

3.2 Econometric Framework

Our basic estimating equation takes the following linear form:

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

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.

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 $g(t)$. 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 . To reiterate, 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. Consistently with Figure 1, 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.

Note that the fit of the model is surprisingly 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 orders¹⁰. This fact underlines the importance of controlling for this variable in every specification.

¹⁰Obviously this does not mean that their wages are higher, as the wages productivity component is a linear combination of goods dispatched and orders processed.

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). This fact strengthens our claim that unobserved time-varying factors correlated with productivity are unlikely to be biasing our estimates. Moreover, the estimated linear trend is not statistically different from zero.

In our benchmark specification of Column (3), which includes extra covariates 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 3.8%. The cumulative effect is 6.6%.

The coefficients on the covariates are interesting in their own right. Consider first 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. Lastly, note that working on certain days of the week is associated with higher productivity.

To conclude this Section, we find that the productivity effects of the RPIE are on average positive both at the kickoff stage and at the revelation stage.

3.4 Robustness Checks

We now explore the robustness of our previous findings to a number of changes in our benchmark specification.

To investigate whether or not the increase in productivity may be transitory, we repeat our benchmark regression using longer time windows. The evidence is presented in Table 3, Columns (1)-(3). We find that our results are always statistically significant and only slightly different in magnitude if we use 10 months, 16 months and 26 months windows. Thus, we find no evidence to suggest that the increases in productivity found in Table 2 are not permanent.

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. To do this, we need to use a time window of 10 months¹¹. Column (4) of Table 3 shows that our main findings are indeed robust to allowing $g(t)$ to take a more flexible form. The estimated coefficients are in fact remarkably similar to those in Column (1), the linear trend 10 months window specification.

In our benchmark specification from Column (3) in Table 2 the linear trend is estimated using both

¹¹The reason is that a problem of severe collinearity between t and t^2 arises in our baseline time window of 6 months.

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

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

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

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

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.

We next include in the sample workers who are in their 6-month probationary period. These workers are obviously less experienced than those past their probation, so it is unsurprising to see in Column (7) that their inclusion yields slightly different coefficients¹². The estimated coefficients are still significant both at the kickoff stage and at the revelation stage.

We next cluster the standard errors at the worker-period (i.e. baseline period, kickoff period, revelation period) level, to allow for idiosyncratic worker characteristics that lead to worker productivity over different days being correlated. Column (8) reveals that the standard errors are only slightly larger than in our benchmark specification of Table 2.

Lastly, we provide evidence in support of our claim in Section 2 that, controlling for the average number of goods per order, the number of goods per hour and the number of orders per hour are equivalent measures of productivity. We estimate our benchmark specification in Column (9), but now our dependent variable is the number of orders per hour. Unsurprisingly, we find that our estimated coefficients of interest, $\hat{\beta}_k$ and $\hat{\beta}_r$, are identical whether we measure productivity in goods per hour or in orders per hour.

¹²The decrease in the kickoff coefficient when we include the probationary workers is consistent with the evidence to be shown in Table 4 that less experienced workers react by a lower amount at the kickoff stage

3.5 Heterogeneous Effects

Table 2 shows that the *average* effect of the RPIE is 6.6%. We would however expect this effect to be different for different workers, perhaps even negative for some. To understand the amount of heterogeneity in the reaction to the RPIE, we ran 58 regressions (one for each worker in the sample¹³) equivalent to (1) and computed the sums of the estimated kickoff and revelation coefficients. We display these individual-specific overall effects in Figure 2.

While Figure 2 confirms the existence of substantial heterogeneity across workers, the reaction to the RPIE can be regarded as positive for the overwhelming majority of workers. Firstly, most (i.e. 51 out of 58) of the effects have a positive sign. Secondly, of the 19 workers for whom the computed overall effect is statistically different from zero at the 10% level, 18 have a positive sign, while only 1 has a negative sign. Thirdly, if we enlarge the time window to 10 months (see Figure A4 in the Appendix) we find an even more unbalanced set of estimated effects, with 30 workers displaying a statistically significant and positive reaction and 0 workers displaying a statistically significant and negative reaction.

Thus, we find almost no evidence that any worker exerts less effort following the introduction of the scheme, while a large proportion of workers exert more effort. Very similar conclusions are reached by looking separately at the revelation effect, rather than the sum of both effects (see Figure A5 in the Appendix).

In Table 4 we take a more systematic look at this heterogeneity by splitting the workforce into different groups, according to gender, experience and pre-existing level of performance, and running our benchmark specification (1) separately for different groups. Panel A of Table 4 shows that reactions by males and females are very similar. While the estimated coefficients are slightly different in magnitude, these differences are not statistically significant¹⁴.

In Panel B we split the sample by experience level. We find that only workers in the middle and top third in terms of experience react to the kickoff of the RPIE. The differences between the top and the bottom third are significant at the 5% level. 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 that the bottom third increase their productivity by a larger amount at the kickoff stage. At the revelation stage, in which workers learn their position in the performance distribution, it is however the *best performers* that increase their productivity by more (i.e. by 5.4%,

¹³For five workers the data was unavailable for the whole of the sample period, as a result of them leaving or joining the firm.

¹⁴A recent literature focusing on laboratory settings (Gneezy, Niederle and Rustichini, 2003; and Niederle and Vesterlund, 2007) has argued the existence of gender performance differences in competitive environments. Lavy (2008) however shows that such gender differences are absent in a real tournament occurring in a workplace setting.

which is more than twice the increase by workers in the middle third; the p-value of the difference is .08).

3.6 Interpretation and Alternative Explanations

To summarise the findings in this Section: (a) the effect of the RPIE is on average positive both at the kickoff stage and at the revelation stage (Figure 1 and Table 2), (b) individual regressions show that the effect is positive for a large proportion of workers and it is not negative for any worker (Figure 2, Figure A4, Figure A5), (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 increase in productivity at the revelation stage is larger for workers at the top of the performance distribution (Panel C of Table 4).

We now discuss the potential explanations of our findings.

3.6.1 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 its communication to them. 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, 1985b).

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; Bandiera et al., 2005). These preferences would be consistent with workers at the bottom of the distribution increasing their effort by a large amount, while 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 Figure 2 and Table 4.

We instead regard relative concerns of the 'status' type, where increases in others' performance always has a negative effect on own utility (Frank, 1985b; Fershtman et al., 2005; Moldovanu et al., 2007), which is consistent with our findings. As these authors discuss, an increase in these concerns translates into a 'rat race' in which everybody increases their level of effort and nobody decreases it, as we find in Figure 2, Figure A4 and Figure A5. Secondly, if, as seems conceivable, the psychological

reward for ranking one position higher is stronger at the top of the distribution than at the bottom, this type of relative concerns is also consistent with our findings from Panel C of Table 4 that the best performers increase effort by a larger proportion than the worst performing workers.

We now discuss alternative explanations

3.6.2 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 *essentially* zero. As also discussed, both 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.

It is also important to note that career concerns interpretations would find it difficult to explain why productivity did not decrease for *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 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.

3.6.3 Alternative Explanation: Learning about Own *Absolute* Ability

Consider the following potential explanation, based on Ertac (2006). Firm-wide shocks to productivity are important, and therefore information about how much others produced in a month can help a worker learn about his own *absolute* ability for the job. For instance, a worker with low productivity in a month may want to revise expectations about his own ability upwards when others also did badly, as his low productivity may simply have been the result of an unfavorable firm-wide productivity shock. In a setting where effort and ability are complements, this will result in a higher level of future effort by this individual when he receives information about other workers' performance.

Can learning about one's own absolute ability represent a potential explanation? There are two problems with this interpretation. Firstly, 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. Secondly, the mechanism outlined above should have led at least some workers to revise their expected ability downwards, leading to lower effort. In general, in any Bayesian model of learning

workers should react differently depending on whether they are positively or negatively surprised relative to their prior expectations. Yet, we find that no worker decreases his effort following the introduction of the RPIE.

3.6.4 Alternative Explanation: Learning about the Distribution of Performance

A similar set of arguments reduces the plausability of an explanation based on ‘learning about the distribution of performance’. This explanation would posit that when workers learn at the revelation stage that higher productivity is achievable, they suddenly become motivated to achieve it or to learn from those at the top of the distribution. This alternative explanation is inconsistent with the documented positive reaction at the kickoff stage, in which there can be no learning. It is also inconsistent with the finding that no worker reacts by decreasing his effort at the revelation stage, as we would expect from workers at the top of the performance distribution, who suddenly find out that there is little or nil potential for improvement.

Summing up, we believe that the existence of concerns about relative pay and productivity is the most likely mechanism at play.

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

Recall, as indicated in Section 2, that 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 occurred 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.

We display in Figure 3 the evolution of the average mistakes rate over time. Figure 3 reveals that 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.

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

$$M_{it} = d_t + \mathbf{X}'_{it} \cdot \delta + \mu_i + v_{it} \quad (4)$$

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 v_{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 (4). 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.

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

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

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 ability¹⁵ and $\mathbf{X}'_i$ is a vector of worker characteristics.

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). As discussed in Sections 2 and 3.1, 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 February 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 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 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.

¹⁵To 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 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.

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 already at the kickoff stage. To the best of our knowledge, these results are the first to show that social comparisons can lead workers into engaging in effort 'rat races', whereby information about relative standing translates into efforts to outperform each other.

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. To understand the external validity of our findings, there are two features of our firm and work environment that we regard as particularly significant:

(1) We note first that the effects of providing information on relative productivity and pay may depend on the technology of production and on the incentive scheme in place. In our firm, the existence of absolute performance pay implies a one-to-one mapping between performance and compensation¹⁶. 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). Nevertheless, our findings provide an initial benchmark to which future work based on other workplace settings can contribute.

(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 results under a

¹⁶Obviously 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.

public disclosure of relative standing. While this 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

- Bandiera, O., Barankay, I. and Rasul, I.** (2005), “Social Preferences and the Response to Incentives: Evidence from Personnel Data”, *Quarterly Journal of Economics*, 120(3): 917-962.
- Brown, G., Gardner, J., Oswald, A., and Qian, J.** (2008), “Does Wage Rank Affect Employees’ Wellbeing?”, *Industrial Relations*, 47(3): 355-389.
- 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. and Oswald, A.** (1996), “Satisfaction and Comparison Income”, *Journal of Public Economics*, 61: 359-381.
- Charness, G. and Kuhn, P.** (2007), “Does Pay Inequality Affect Worker Effort? Experimental Evidence”, *Journal of Labor Economics*, 25(4): 693-725.
- Ehrenberg, R. and Bognanno, M.** (1990), “Do Tournaments Have Incentive Effects?”, *Journal of Political Economy*, 98(6): 1307-1324.
- 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 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*.
- Frank, R.** (1984), “Are Workers Paid their Marginal Products?”, *American Economic Review*, 74(4): 549-571.
- Frank, R.** (1985a), “The Demand for Unobservable and Other Positional Goods”, *American Economic Review*, 71(1): 101-116.
- Frank, R.** (1985b), “Choosing the Right Pond: Human Behavior and the Quest for Status”, New York: Oxford University Press.
- Frank, R.** (2005), “Positional Externalities Cause Large and Preventable Welfare Losses”, *American Economic Review, Papers and Proceedings*, 95(2): 137-141.

- Gneezy, U., Niederle, M. and Rustichini, A.** (2003), "Performance in Competitive Environments: Gender Differences", *Quarterly Journal of Economics*, 118: 1049-1074.
- Hopkins, E. and Kornienko, T.** (2004), "Running to Keep in the Same Place: Consumer Choice as a Game of Status", *American Economic Review*, 94(4): 1085-1107.
- 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.
- Landers, R., Rebitzer, J. and Taylor, L.** (1996), "Rat Race Redux: Adverse Selection in the Determination of Work Hours in Law Firms", *American Economic Review*, 86(3): 329-348.
- Lavy, V.** (2008), "Gender Differences in Market Competitiveness in a Real Workplace: Evidence from Performance-Based Pay Tournaments among Teachers", NBER Working Paper, 14338.
- Layard, R.** (2005), "Happiness: Lessons from a New Science", London: Allen Lane.
- 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.
- Main, B., O'Reilly, C. and Wade, J.** (1993), "Top Executive Pay: Tournament or Teamwork?", *Journal of Labor Economics*, 11(4): 606-628.
- Milkovich, G. and Newman, J.** (1996), "Compensation", Chicago: McGraw-Hill/Irwin.
- Moldovanu, B., Sela, A. and Shi, X.** (2007), "Contests for Status", *Journal of Political Economy*, 115(2): 338-363.
- Neumark, D. and Postlewaite, A.** (1998), "Relative Income Concerns and the Rise in Married Women's Employment", *Journal of Public Economics*, 70: 157-183.
- Niederle, M. and Vesterlund, L.** (2007), "Do Women Shy Away from Competition? Do Men Compete Too Much?", *Quarterly Journal of Economics*, 122(3): 1067-1102.
- Oswald, A.** (1997), "Happiness and Economic Performance", *Economics Journal*, 107: 1815-1831.

Figure 1: Evolution of Productivity over Time

Note: Dependent variable is (log of) number of goods dispatched per hour worked. Horizontal 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) productivity of a different working day.

Figure 2: Kernel Density of Overall Effects Coefficients from Individual Regressions

Note: Figure displays the kernel density of the sums of the Kickoff coefficient and Revelation coefficient arising from 58 individual regressions equivalent to the one in Column (3), Table 2. In 18 of these regressions the sum of the Kickoff coefficient and Revelation coefficient was **positive** and statistically different from zero. In only 1 regression the sum of the Kickoff coefficient and Revelation coefficient was **negative** and statistically different from zero.

Figure 3: Evolution of Mistakes Rate over Time

Note: Dependent variable is (log of) number of mistakes per good dispatched. Horizontal 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.

Table 1: Descriptive Statistics

Panel A: Firm-Level Characteristics				
	(1) Mean	(2) S.D.	(3) Minimum	(4) Maximum
Goods dispatched (total)	47,487.88	14,172.94	18,339	90,578
Customer orders (total)	927.44	272.35	338	1,541
Goods per order (by average worker)	51.36	3.72	42.99	62.59
Workers (total in payroll that day)	66.48	1.59	63	69
Workers (total working that day)	43.22	12.02	19	61
Hours worked (total)	316.65	90.36	137.75	558.5
Hours worked (by average worker)	7.39	.94	3.49	9.73
Goods per hour (by average worker)	149.73	7.85	126.84	172.14
Number of Observations	152			
Panel B: Individual-Level Characteristics				
	(1) Mean	(2) S.D.	(3) Minimum	(4) Maximum
Male	.63	.48	0	1
Married	.62	.49	0	1
Experience (in years)	3.99	4.06	.50	23.88
Age (in years)	38.75	8.35	21.27	58.48
Whether left job before 31/12/2007	.38	.49	0	1
Years to departure (if left job)	2.34	1.82	.16	5.41
(Undisclosed) Ranking in August 2001	34.01	18.35	3	66
(Disclosed) Ranking in September 2001	35.71	18.68	4	67
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. Panel B displays variables at worker level. Sample is restricted to those workers with at least 6 months of experience in the firm.

Table 2: Baseline Results

	(1) Unconditional	(2) Time Trend	(3) Other Covariates
Kickoff RPIE	.027*** (.005)	.033*** (.010)	.028*** (.006)
Revelation RPIE	.032*** (.006)	.038*** (.010)	.038*** (.007)
(Log of Average) Goods per Order	.312*** (.011)	.313*** (.011)	.336*** (.006)
Time Trend		-.0001 (.0001)	-.0001 (.0001)
(Log of) Goods Piece Rate			-.014 (.103)
(Log of Average) Value of Goods (In Euros)			-.176*** (.022)
Tuesday			.003 (.003)
Wednesday			.005 (.003)
Thursday			-.017*** (.003)
Friday			.017*** (.004)
Sunday			-.011** (.005)
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. 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.

Table 3: Robustness Checks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	10	16	26		Pre-Treatment	Pre-Treatment	With Probationary	Clustering	Dep.Var.:
	Months	Months	Months	Quadratic	Trend	Effects	Workers	S.E. by	Orders
	Window	Window	Window	Trend	Goods	Goods	Goods	Worker	per hour
Dependent	Goods	Goods	Goods	Goods	Goods	Goods	Goods	Goods	Orders
Variable:	per hour	per hour	per hour	per hour	per hour	per hour	per hour	per hour	per hour
Kickoff	.032*** (.005)	.016*** (.003)	.025*** (.003)	.034*** (.008)	.024*** (.005)	.029*** (.005)	.018*** (.006)	.028*** (.007)	.028*** (.006)
RPIE									
Revelation	.045*** (.003)	.033*** (.003)	.024*** (.003)	.042*** (.009)	.032*** (.005)	.048*** (.004)	.035*** (.006)	.038*** (.008)	.038*** (.007)
Quadratic Trend	No	No	No	Yes	No	No	No	No	No
Pre-Treat Trend	No	No	No	No	Yes	Yes	No	No	No
Pre-Treat Effects	No	No	No	No	No	Yes	No	No	No
Worker F.E.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Table 2 Covar.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Standard	Day	Day	Day	Day	Boot-	Boot-	Day	Worker	Day
Errors	Cluster	Cluster	Cluster	Cluster	strapped	strapped	Cluster	Period	Cluster
								Cluster	
Adjusted R ²	.74	.72	.73	.74	.74	.74	.77	.77	.85
Window Months	10	16	26	10	6	6	6	6	6
Observations	8,531	13,556	23,576	8,531	5,275	5,275	6,163	5,275	5,275

Note: Dependent variable in Columns (1)-(4) and (6)-(7) is (log of) number of goods dispatched per hour worked. Column (1) data comprises 1May2001-28Feb2002. Column (2) data comprises 1Feb2001-31May2002. Column (3) data comprises 1Aug2000-31Nov2002. In Column (5) the linear time trend is estimated exclusively on the baseline period (July and August) according to the description in Section 3.4. 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) according to the description in Section 3.4. 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.

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

Panel A: Decomposition by Gender				(1)-(2)		
	(1) Males	(2) Females		P-value		
Kickoff RPIE	.022*** (.007)	.037*** (.011)		.26		
Revelation RPIE	.045*** (.008)	.027** (.012)		.23		
Worker Fixed Effects	Yes	Yes				
Table 2 Covariates	Yes	Yes				
Observations	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 RPIE	-.011 (.012)	.025** (.010)	.037*** (.010)	.35	.02	.39
Revelation RPIE	.046*** (.013)	.046*** (.012)	.036*** (.011)	.95	.54	.57
Worker Fixed Effects	Yes	Yes	Yes			
Table 2 Covariates	Yes	Yes	Yes			
Observations	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 RPIE	.056*** (.011)	.015* (.009)	.014 (.011)	.00	.01	.90
Revelation RPIE	.030** (.014)	.026** (.011)	.054*** (.011)	.82	.19	.08
Worker Fixed Effects	Yes	Yes	Yes			
Table 2 Covariates	Yes	Yes	Yes			
Observations	1,595	1,825	1,811			

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. Robust standard errors allow for cluster at day level.

Table 5: Evidence on the Mistakes Rate

	(1) Unconditional	(2) Controls
July (Baseline Month 1)	Omitted	Omitted
August (Baseline Month 2)	-.170*** (.076)	-.220* (.110)
September (Kickoff Month 1)	.167*** (.072)	.194** (.081)
October (Kickoff Month 2)	.013 (.072)	.113 (.084)
November (Revelation Month 1)	.033 (.063)	.145 (.117)
December (Revelation Month 2)	-.059 (.080)	-.014 (.112)
Worker Fixed Effects	No	Yes
Table 2 Covariates	No	Yes
Adjusted R ²	.03	.77
Number of Observations	333	333

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.

**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 (1)	Cox Regression Ability (2)	Cox Regression Other Covariates (3)
Revelation RPIE	-1.70*** (.359) [.181]	-1.72*** (.367) [.178]	-1.64*** (.537) [.193]
Worker Ability		-6.24*** (1.19) [.001]	-7.88*** (1.42) [.0003]
Revelation RPIE*		-.008 (.491) [.486]	.216 (.489) [1.24]
Worker Ability			
Male			.315 (.409) [1.37]
Married			-.410 (.389) [.662]
Year Joined Firm			-.141 (.111) [.868]
Age at Entry			-.056** (.022) [.945]
Number of Observations	719	719	719
Number of Subjects	127	127	127
Number of Separations	41	41	41

Note: Revelation RPIE 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. Data comprises 1Feb1999-31Dec2007. The sample is restricted to workers joining the firm after February 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.

Tournaments without Prizes:
Evidence from Personnel Records

APPENDIX

NOT INTENDED FOR PUBLICATION

Figure A1: The Returns to Experience

Note: Figure displays Quarter of Experience Dummies calculated from a model that has productivity on the left hand side and includes worker fixed effects and day fixed effects, as well as experience dummies, on the right hand side. The interpretation of the dummies, noting that the dependent variable is in logarithmic form, is as follows: workers produce 16% more in their second quarter in the firm than in their first quarter; 20% more in the third quarter than in the first quarter, and so on. Data comprises 207 workers over the period 1Feb1999-31Dec2007. In total the data contains 127,531 worker-day observations.

Figure A2: Placebo Kickoff and Revelation - 2000

Note: Dependent variable is (log of) number of goods produced per hour worked. Horizontal dashed lines displaying the 1st of September 2000 (Placebo Kickoff date) and 1st of November 2000 (Placebo Revelation date) displayed. Each dot is the (averaged across workers) productivity of a different working day.

Figure A3: Placebo Kickoff and Revelation - 2003

Note:Dependent variable is (log of) number of goods produced per hour worked. Horizontal dashed lines displaying the 1st of September 2003 (Placebo Kickoff date) and 1st of November 2003 (Placebo Revelation date) displayed. Each dot is the (averaged across workers) productivity of a different working day.

Figure A4: Kernel Density of Overall Effects
from Individual Regressions - 10 Months Window

Note: Figure displays the kernel density of the sums of the Kickoff coefficient and Revelation coefficient arising from 62 individual regressions equivalent to the one in Column (3), Table 2. In 30 of these regressions the sum of the Kickoff coefficient and Revelation coefficient was **positive** and statistically different from zero. In **no** regression the sum of the Kickoff coefficient and Revelation coefficient was **negative** and statistically different from zero. Data comprises 1May2001-28Feb2002.

Figure A5: Kernel Density of Revelation Effect from Individual Regressions - 6 Months Window

Note: Figure displays the kernel density of the sums of the Revelation coefficient arising from 58 individual regressions equivalent to the one in Column (3), Table 2. In 12 of these regressions the Revelation coefficient was **positive** and statistically different from zero. In 1 regression the Revelation coefficient was **negative** and statistically different from zero. Data comprises 1Jul2001-28Dec2001.