

Rank Incentives:
Evidence from a Field Experiment¹

Iwan Barankay²
University of Pennsylvania
CEPR IZA

Preliminary Version
March 2010

Abstract

Even though purely monetary incentives can be effective it has been argued that people may also have various social preferences that affect their effort provision. This paper studies the extent to which employee's effort choice is affected by the knowledge of their rank in the wage distribution. Using data from a large randomized controlled study at a leading furniture company we show that when people don't know their rank in the wage distribution they are significantly more productive than when they do. It is shown that this result is driven by workers in the top 30% of the distribution who exhibit significantly lower effort when they are informed of their rank. Interestingly we find that the effect is gender specific in that only men change their behavior in light of rank information. This paper sheds new light on the role of status in the workplace and how it relates to other distributional social preferences like inequity aversion. It also has important policy implications as many countries use or consider rankings as an incentive device for public service provision.

Keywords: Rank, Status, Incentives, Field Experiment

JEL Classification: D23, J33, M52

¹ Financial support from the ESRC, the Wharton Center for Human Resources, and the William and Phyllis Mack Center for Technological Innovation at the Wharton School are gratefully acknowledged.

² The Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104, USA; Tel: +1-215-898-6372; E-mail: barankay@wharton.upenn.edu.

Extended Abstract

Introduction

Even though monetary incentives can often achieve desirable outcomes (Lazear, 2000) it can fail to achieve the desired outcome for a number of reasons. In most productive environments incentive contracts can generate gaming and other unintended consequences due to moral hazard or adverse selection (Prendergast 1999). In addition to that people might care about other rewards besides pecuniary compensation as they may have social preferences (Fehr and Schmidt, 1999) or care about their image generated by their actions (Ariely et. al, 2009). This paper explores how status generated by the knowledge about one's ranking in the wage distribution affects effort provision.

It has often been argued that people systematically compare their situation to that of others (Veblen, 1949, Frank 1985). In the context of wages, people may derive status from the knowledge about how they fair compared to others. Despite the large literature on this topic, there is scant evidence as to how status concerns translate to actual workplace behavior. Indeed people might have very heterogeneous responses when they learn about their rank. They may be driven to work harder to improve their status, they may not care about their status, or they may work less hard as they become complacent or lose motivation pressed about their status depending on their position in the distribution. To my knowledge this is the first paper to employ a randomized-controlled field experiment in a large company to shed light on the question of whether people change their effort depending on whether they are told their rank in terms of compensation in a company and to also explore the heterogeneity in the treatment response.

There are three key results that emerge when randomize whether people know their rank in terms of compensation in a company.

Result 1: People who are not told their rank work significantly harder on average compared to those who are told their rank.

Result 2: The result is driven by the people in the top 30% of the distribution working less hard when they are told their rank.

Result 3: The result is gender specific. The change in behavior only exists among men; women don't seem to act on or ignore the rank information

This paper speaks to three strands of the literature. First, the area of feedback and interim performance evaluations highlighted that feedback can have complex effects when the information is not verifiable. In our context the feedback is about hard and verifiable information so that the incentives of the principal to give a biased result to the agent is muted. There has also been work that suggests that feedback can help to establish or reinforce social norms and communicate to the employer what the company cares about (Kluger and Denisi, 1996). Depending on the context this can help or reinforce multi-tasking problems. Second, a large literature, especially in social psychology, discusses the role and the importance of status and rank order preferences. Status is often defined to be utility generated by how favorably a person's compensation (or other metric) compares to that of a reference group.

social psycholand is

in workplace (Bandiera et al, 2009, Blanes-i-Vidal and Nossol, 2009, Cueni et al, 2009) . This in many ways is also related to conspicuous consumption (Veblen 1949). It is also worth noting how concerns for status typically contrasts sharply from inequity aversion (Fehr and Schmidt, 1999; Fehr and Gächter, 2000, Rey-Biel, 2008) where people have a preference for compressed wage distributions.

This paper also complements the literature that explored wage inequality and rank with laboratory experiments (Charness and Kuhn, 2007, Eriksson et al. 2008)

Third, there has been a very interesting literature on heterogeneity in behavior in competitive environments. Information about rank (league tables, wall of honor, employer or team member of the month) can generate a competitive environment just like tournament do. There is now growing evidence that establishes gender-heterogeneity in competitive behavior. Women shy away more often from competitive environments and do, on average, less well in them compared to men. (Gneezy et al 2003, Niederle and Vesterlund, 2007, Gneezy and List, 2008).

Empirical Context and Challenges

It is rather hard to generate compelling evidence on the role of rank from purely observational data. Wages and rank are determined endogeneously (Kwon and Meyersson Milgrom, 2009) and as rank is often linked to pecuniary benefits either in static (period specific performance pay) or dynamic form (promotions or access to additional resources) it may just be a proxy for monetary compensation.

Furthermore there is now growing recognition that one has to have intimate knowledge about the production function to establish the mechanism at play in a specific workplace. In a team context (Bandiera et al., 2009, Dur et al, 2009) when people may be forced to leave teams in light of feedback, ranking information generates incentives to ensure retention.

In other contexts when rank is arguably not linked to monetary incentives, the absence of clean control group and randomized changes in the environment may limit the interpretation of results (Blanes-i-Vidal and Nossol, 2009, Cueni et al, 2009, Ehrenberg and Bognanno, 1990)

Context

To address these empirical challenges this paper uses data generated by a field experiment with the salespeople of a large office furniture company in North America. The company produces the furniture yet this furniture is then sold by salespeople employed and working in franchised dealerships. As is typical for this kind of task, the compensation of these salespeople is purely commission based. The salespeople derive all their compensation from commissions which are a function of the value of furniture sold.

There are 238 dealership with 1165 salespeople working at some stage during 2009. For the experimental evaluation we use salespeople active at the time of treatment 600 salespeople.³ All salespeople log their furniture sales onto a webpage.

³ Salespeople may leave the dealerships over time. The departure and arrival of salespeople was not affected by the treatment over the period of six months of data after the treatment started. In other contexts it has been shown that selection can be attributed to be just as important as the incentive effects (Lazear, 2000).

On that webpage they can see their total bonuses received and verify that all the items sold are paid the correct rewards. On that webpage they get a monthly update on their bonus payouts. This is very accurate according to both the company and based on a survey of 294 salespeople.

The Treatment

Before the start of the treatment, all salespeople were told their rank in term of salesbonuses in the North America and within the dealership on the webpage. Only the current rank information was displayed and they cannot see the performance or the ranking of any other specific salespeople. In August of 2009, the salespeople were randomly divided into three groups. The first group continued to see their rank, the second group did no longer see their rank, and the third group continued to see their rank but it was in relation to a different reference group.⁴ The treatment remained in place for the rest of 2009.

It is important to make four remarks here. First, note that Salespeople are paid for their performance alone so there is no reward at the dealership or team level. Second, there were no rank based compensations in the form of tournaments. Third, salespeople are typically the sole person in charge of closing and seeking out deals. Finally, no other elements of the environment were changed in particular none of the other terms of the incentive scheme changed.

As is apparent, the treatment is very subtle so that we can interpret this as a marginal change in the environment (compared to deep interventions with general equilibrium effects). The treatment is very subtle as it only changes information on the webpage. No other documentation of information on this was shared.

Salespeople were no aware that they were part of a field experiment (debrief at end of study)

Power calculations were conducted based on pre-treatment data that guided the stratification of the treatment group by gender and by rank.

To assess the timing of the treatment note that the salespeople first have to notice the change on the webpage; then they need to sell furniture and record their sales into the system. As the data gets updated monthly on the webpage it is reasonable to assess the treatment effect at least two months after the change.

Results

We have been given access to personnel records of all salespeople at the company. The key variables of interest are the value of sales per month which is the basis of the ranking information. We also have access to the complete data of each sale, the frequency and timing or login activity to the webpage and some background characteristics of the salespeople.

⁴ In that third group the ranking was given in terms of a “normalized rank”. Those salespeople were informed on the webpage that the normalized rank took into account differences in local market conditions. It was established by basing rank on residuals of state fixed effect regressions.

We find that compared to the group who continued to receive their rank information salespeople in the “normalized rank group” did not change their behavior so that the change in reference group did not matter.

Salespeople in the group that no longer received their rank information sold on average over 20% more furniture compare to those who continued to see their rank. Decomposing this effect it turns out that it were the salespeople in the top 30% of the distribution (either in terms of prior rank or in terms of quantile regressions) who sold more furniture once they no longer saw their rank information. Salespeople at or below the median did not change their behavior.⁵

Given prior work on competitive environments we stratified the treatment by gender. Looking at differences in means by gender we notice there is no change in behavior for women. Only men change their sales-performance in light of the presence or absence of rank information.

Discussion

Some prior work found positive effects of giving ranking feedback on performance (Dur et al, 2009, Bandiera et al, 2009, and). One important difference to those papers is that here we have individual and not team work. In team work the threat of being pushed out of the team is very strong due to the common pool problem. With individual pay we see that knowing the rank does not affect anyone at the bottom or in the middle of the distribution so there is no catching up effect nor does feedback on relatively poor performance lead to a loss of motivation. What is remarkable is that even the top people don't seem to get any additional boost in performance from positive feedback as is often hypothesized.

What we find instead is a complacency effect for people at the very top.

The results are also in contrast to predictions of inequity aversion. If all salespeople are inequity averse, we should see symmetric changes at the top and the bottom of the sales-distribution. Yet, there is only a significant change at the top of the distribution.

The ranking information so far has been very coarse. Indeed one may argue that giving feedback about one's rank only matters when one is also informed what performance improvement are needed to make a difference in terms of ranking. In the beginning of 2010 we therefore continued our experiment at the furniture company. Salespeople have now been randomly assigned to three groups. The first group is told their rank in North America and at their dealership; the second group also is told their ranking in North America and the dealership but there are also told the current performance in terms of the value of furniture sold of salespeople just within the top 50%, 25% and 10% of the distribution; the third group does not get ranking information.

⁵ We have the login data from the Webpage (preliminary). There is no difference in the frequency and of logins by gender. Nor does it change after the treatment.

Conclusions

We implemented a large scale RCT in the field to study how knowledge about one's rank affects effort.

In contrast to some other studies the complacency effect seems to be the dominant behavioral response. Not knowing one's rank leads to higher performance compared to being told the rank. The experiment is still ongoing and will shed light if the degree of information about the rank matters i.e. knowing how much one needs to do to make a significant change to one's rank.

The results have important implications on how we may quantify the effect of status as induced by ranking information onto different parts of the distribution and to different types or people. We find that knowing the rank only male people at the top of the distribution.

The results also informs the vivid debate in the public sector on how to employ ranking and league tables for schools and hospitals as discussed in the U.K. and the US.