Individual incentives and workers' contracts: evidence from a field experiment

M. Ali Choudhary^a

Vasco J. Gabriel^b

Neil Rickman^c

^a State Bank of Pakistan, Karachi

^b University of Surrey and NIPE-UM, Guildford, GU2 7XH, United Kingdom,
v.gabriel@surrey.ac.uk

^c University of Surrey, Guildford, United Kingdom

ABSTRACT

We present evidence on the operation of incentive pay from a field experiment in Pakistan, looking at piece rates and pay based on rank achieved in a tournament. Importantly, some workers are in contracts ‘tying’ them to the employer for several picking seasons; others are ‘untied’ in the sense of being employed for only the current season. We find that incentive pay (of either type) improves productivity by 30% on average, but that there are important differences across the types of workers: in particular, tournament incentives are less effective amongst the tied workers. We suggest that our main results have implications for tournament theory and the design of incentive pay schemes, particularly to the extent that they may discourage some workers and, thus, reduce incentives.

JEL CLASSIFICATION: D23, J23, J33, M52

1 Introduction

The question of how to encourage effort from workers who have private information about their abilities is a central focus of principal-agent theory. In a seminal paper that stimulated both theoretical and empirical work, Lazear and Rosen (1981) compare the effects of rank-order tournaments and piece rates on worker productivity in such settings. On the theoretical side, their finding (that the two schemes were equivalent in a static, deterministic setting) has been developed in a number of ways (surveyed by Prendergast, 1999; Konrad, 2009; Decheneaux et al., 2015). Tournaments have been shown capable of

inducing first best effort in the presence of risk neutral agents and common productivity shocks, but such results can be undermined by workers' concerns for fairness (Grund and Sliwka, 2005), collusive behaviour (Gürtler, 2010) and by performance feedback to workers within a repeated tournament (see Kräbmer, 2007; Münster, 2009; Konrad, 2009; Gürtler and Harbring, 2010). As such, evaluating the relative performance of paying by piece rates or tournaments has also required empirical attention.

Such empirical work has taken several directions. A number of papers have assessed the merits of performance-related pay using 'real effort' data from actual situations but, as the surveys by Bandiera et al. (2011) and Dechenaux et al. (2015) make clear, a large amount of empirical evidence has also been taken from field and laboratory experiments. This reflects the difficulties of measuring performance in many 'real' contexts and the potential for factors like career motivation or terms and conditions to blunt the incentive effects the pay schemes are designed to produce.

Prominent amongst field experiments (such as the one we present in the current paper) is the set of papers by Bandiera et al. (2005, 2010 and 2013). These authors look at the effects of financial incentive schemes on fruit pickers on a UK farm. They find that between-worker and manager-worker social relations play an important role in determining employees' behaviour and how they respond to financial incentives. One particularly relevant finding for the current paper is that, in the context of team production, tournament incentives improve productivity but less so than monetary piece rate ones.

Our paper adopts a similar approach by examining the effects of financial incentives (arising from piece rates and tournaments) on fruit pickers working on farms in Pakistan. Our study involves the seasonal picking of kinnow, an orange variety of which Pakistan was the tenth largest exporter at the time of our study, see Naseer (2010). Producers typically smooth their revenue stream by engaging in a variety of agricultural operations and therefore find it convenient to hire part of their workforce on a long-term basis in order to ensure stable labour supply. As the year, the type of produce in season and, hence, the location moves on, so the producers will top-up their workforce from (typically local) pickers who are retained only for a specific picking season and have low probability of graduating to permanent status because there are limited numbers of vacancies.¹ Hence, there are

¹Thus, the 'stepping-stones' effect in Engellandt and Riphahn (2005) may not serve as a motivation to work harder.

essentially two worker types: ‘tied’ (i.e. longer-term workers) and ‘untied’ (i.e. shorter-term workers).² The tied-untied distinction turned out to have an important bearing on workers’ reactions to the new pay structures we deployed.

The customary mode of pay for these types of workers is a fixed monthly salary (often paid in weekly portions). Thus, while the arrangements we have described may address labour supply issues, they are not likely to minimize the possibility of moral hazard amongst workers. This is clearly a problem for producers in this important export sector because, while export quality standards may regulate the quality of work, there is a danger of sub-optimal productivity. As a result, we approached a firm of fruit pickers/exporters in order to examine the effects of two incentive pay schemes: one involving a bonus payment for picking beyond a specified weight threshold (we call this a ‘piece-rate’ scheme) and the other involving prizes for the top three pickers on each day (a rank-order tournament). Both types of incentive scheme were applied in addition to fixed base pay, paid out on a ‘next-day’ basis; each lasted for two weeks.

Our setting had a number of interesting characteristics that help to distinguish it from the existing literature. First, our incentives were applied at the individual basis and across two different types of workers: the effects of incentives on workers operating under different contractual arrangements in the same setting has not received attention to date. Second, because we paid the workers on a daily basis, our tournament experiment was akin to a dynamic tournament amongst individual workers with feedback on (some elements of) performance. As is becoming increasingly apparent from theoretical work, dynamic tournaments may not share the positive incentive effects of their static counterparts because of the opportunities for learning that feedback creates. Berger and Pope (2011) and Bandiera et al. (2013) have examined such tournaments on group performance in field experiments, Cason et al., (2010) look at performance disincentives in winner-takes-all tournaments, while Brown (2011) and Gill and Prowse (2012) consider individual prizes/pay in a setting where agents apply efforts sequentially. The natural setting of our field experiment was more akin to individual agents competing in repeated, simultaneous-move tournaments and, as such, offers a new set of results. Third, placing it in a developing economy meant

²This type of work arrangement is common in Pakistani agriculture (see CABI, 2008), and elsewhere (e.g. Bardhan, 1979, 1980, identifies similar set-ups in India, while Macchiavello and Morjaria, 2015, study relational contracts amongst Kenyan flower exporters).

that we could provide a rare study of the effects (if any) of financial incentives in such a setting. It is important to do this at any time, but perhaps especially when microeconomic incentives are playing a role in aid policy, though the extent to which these are effective in settings with different cultural norms, educational characteristics and informational flows is still being learned (good examples of relevant studies include Banerjee and Duflo, 2011; Duflo et al. 2011; Muralidharan and Sundararaman, 2011).

We find that the introduction of incentives (either piece or rank) improved average hourly productivity compared to a simple fixed pay regime that we started the experiment with. As far as comparing the two incentive schemes is concerned, we find that piece rates generally outperformed tournaments. Most interesting, however, was the different ways in which the tied and untied workers responded to the financial incentives on offer. In particular, while untied workers performed at a similar level under both schemes, most tied workers generally underperformed in the tournaments compared to their outputs under piece rates. The exception was the tied workers who won tournaments; a highly concentrated group who (like the untied workers) performed comparably across both pay schemes.

We suggest that these results comparing fee schemes and types of workers can be explained by the fact that the tied workers were more familiar with each others' strengths and, as such, realized sooner than their untied counterparts whether they had realistic prospects of winning a prize. We empirically document the existence of a 'discouragement' effect amongst tied workers which resembles the 'quitters' effect found by Brown (2011) and the 'disappointment aversion' effect explored in Gill and Prowse (2012) in settings where there are patent differences in skill levels amongst agents.³ We perform several robustness checks and our results are regularly consistent with this interpretation. In addition to this, we also find that competition amongst the groups of workers appeared to spur the untied ones, indicating some non-financial motivations also at play.

The paper is organized as follows. Section 2 describes our experimental setting. Section 3 estimates the effects of rank-order pay and piece rates on labour productivity. Section 4 discusses our results before Section 5 concludes the paper.

³Similarly, Singh and Masters (2018) contrast winners-take-all incentives with proportional reward schemes, finding that the latter led to larger and more sustained gains, as lower-ability workers lower their efforts under the winner-takes-all contest.

2 Details of the experiment

With the aim of examining the operation of two forms of incentive pay (piece and rank), and having approached a number of producers with the aid of co-researchers from the University of Agriculture, Faisalabad, we found one willing (in principle) to participate.⁴ We were ultimately allowed to implement these schemes with the day-to-day cooperation of the fruit pickers' immediate line managers – on the understanding that we would build upon (rather than dismantle) current working arrangements. We spent a month advising the firm's CEO and local managers on the schemes we had in mind and the ways in which they could be implemented and monitored without drastically changing existing work methods. Concerns within the firm included its workers' pay levels (which it did not want to see damaged) and the potentially obtrusive nature of recording output per worker (as we required).

The success of such research (from both academic and business perspectives) depended on the design and implementation of the experiment used to generate the data so, in this section, we explain these and how we were able to accommodate the firm's conditions for supporting the research. Throughout the process, we received 'on the ground' assistance from our co-researchers and, in addition, one of us attended and oversaw the full data collection process.

Location The experiment was carried out on two groups of citrus fields located in the Punjab province of Pakistan. The fields were in an agricultural area highly regarded in Pakistan for its citrus production; indeed, it produces 95% of Pakistan's kinnow production (Naseer, 2010). The fields were in close proximity (5 kms apart) and similar in terms of soil type and preparation, size, plants, topography and conditions, thus ruling out (time-varying) effects on productivity due to differences in location or weather conditions.⁵ The fields' dimensions ensured no capacity constraints would bind over the period of the experiment: when one field was finished an adjacent, virtually identical, one was available

⁴The approach was in the form of a business case that sought to demonstrate the potential benefits to the firm from addressing the moral hazard problems implicit in its fixed pay set-up - further details of the business case, as well as on the picking conditions and data collection can be found in the supplementary appendix.

⁵For several days, the geography of the fields meant that the two groups of pickers operated in much closer proximity (to within a distance where they could shout at each other but not see the weights of fruit being picked). As shown below, we attempted to take advantage of this serendipity in order to proxy the effects of inter-group competition on outputs.

and the cost of moving from one field to the other was negligible. Thus, our experiment can be seen as an eight-week snapshot of a several months-long fruit picking process under fundamentally identical picking conditions.

Management arrangements and pickers The pickers (72 in total) were all deployed by contractors (which we shall refer to as Widget Enterprises) specialized in buying fruit from farmers and exporting it. As part of this operation, they provide agricultural inputs to growers. Widget operates across a number of fields and is overseen by a CEO, who is responsible for some aspects of hiring and taking other strategic decisions in relation to the selection of fields and their overall operation. Beneath the CEO, the fields' manager oversees the day-to-day running of the fields and recruitment of workers, while a team leader is used to manage pickers, who are directly beneath this tier.

Two groups of pickers were hired through two team leaders, to whom we refer as T (responsible for the 'tied' workers) and U (responsible for the 'untied' workers). These were supervised by the same fields' manager. Most of the untied workers lived in villages close to our chosen location, while the tied workers came from another region. Untied workers all had at least one year of picking experience and listed farming as their main occupation. All were male and average ages were similar across the two groups.

Nonetheless, the groups differed in one important respect: the basis of their employment status. Those employed with T were 'tied' labourers in the sense that they were contractually employed before the start of the picking season, and were paid a fraction of the base salary in advance. Once the citrus season is over, they move to pick other fruits of the season such as melon and mangoes under similar contractual arrangements. In contrast, U's workers were hired on an 'informal', or 'untied', basis, receiving wages on a weekly basis, with an understanding that the group members will work for the entire kinnow harvest season.⁶ There was little chance of crossing over from untied to tied status, as the tied contingent was already in place and while the untied workers were fairly local, tied pickers worked in different areas throughout the season. Furthermore, the two worker types were accommodated separately and did not mix during the experiment.

Our untied employees were akin to the ones used in Bandiera et al's (2005, 2010,

⁶This requirement also minimizes the possibility of self-selection into pay schemes, given that one of U's responsibilities was to ensure that roughly the same group of pickers was in place for the duration of the harvest; more on this below.

2013) experiments (they call this ‘casual’ labour) with the difference that our workers had some previous fruit picking experience. This minimized the need for learning-on-the-job in our case, and its potential impact on productivity.⁷ Note that our untied workers had not worked together as a group before: the group was assembled for the kinnow collection season. On the contrary, tied workers had been working together since the start of the fruit picking season, at least. In both cases, however, opportunities for observing colleagues’ outputs within the group were limited given the intense nature of the picking activity, the disposition of the fruit trees⁸ and the fact the output was weighed some distance from where it was picked.

Pay schemes Prior to launching the incentive schemes, production under the original method of pay (i.e. fixed base pay) was observed for four weeks. This allowed us to finalize the arrangements for the experiment, and to fine-tune data collection.⁹ This period also allowed workers to learn about these (and the fields to be used). We then required the workers to pick for fourteen days under a rank tournament structure, followed by fourteen days under piece rates that paid a bonus beyond a particular threshold level of output. Note that the nature and the durations of each pay scheme were not revealed to workers so as to minimize potential deadline and order effects.

We selected base pay, prizes and bonuses in consultation with the fields’ manager. Base pay under both our incentive schemes was a fixed sum set equal to its pre-experiment level, thereby meeting the producer’s concerns about workers’ pay levels during the research. The prizes in rank-order tournaments and the above-threshold piece-rates were then selected to achieve some comparability across the two regimes. We now explain this in more detail.

The pickers’ average monthly pay before the introduction of the schemes was Rs.4500 (Rs.150/day)¹⁰ and, so, this continued to be paid as the fixed level of base-pay that all the

⁷We distinguish here between learning about general fruit-picking techniques (which our workers had some experience of) and learning about the narrower specifics of the fields in question and the data collection process. Learning about the latter could have prejudiced our results, so we started the experiment with a period of fixed pay, as already remarked.

⁸These are leafy, with an average height of 22ft, an average distance from trunk to edge of 7.5ft and an average distance between two trunks of 20ft, see Qureshi et al. (1993).

⁹Ideally, we would have liked to have comparable individual data under fixed pay, but due to our time constraints, our need to gain workers’ confidence in the field and ‘bed in’ our data collection processes, and the need to implement rank-pay and piece-rates within a reasonable timeframe, we were not able to record individual daily production for all workers for a significant number of days – nevertheless, data on total daily production are still quite informative as an overall benchmark.

¹⁰Rs.150 is approximately equivalent to £1.10p, or \$1.45c.

workers received. The incentive schemes built on top and additional wages were calculated and paid on a next-day basis by the fields' manager. Thus, under the initial fixed pay regime (F), worker i 's pay on day t was $Y_{it}^F = B$, where $B = \text{Rs.150}$ (i.e. the base pay). In the case of the rank-order tournament, the top three outputs received prizes the following day, and the names of the workers were publicly distributed and displayed on a board for the rest of that (following) day; no other production/performance information was disclosed and the rankings below third place were not announced. The fields' manager was consulted about the appropriate levels at which to set the prizes and he suggested that reasonable incentives would be provided by allowing a worker the chance to double his pay. Hence, the first, second and third prizes were 100%, 66% and 50% of the daily base pay (Rs.150). Thus, if we assume that the rank-order tournament (R) has three monetary prizes, $M_1 > M_2 > M_3$, that a group has L workers, that worker i produces output on day t of q_{it}^R , and worker j wins the first prize and worker k wins the second prize, then

$$Y_{it}^R = B + \begin{cases} M_1 & \text{if } q_{it}^R = \max\{q_{it}^R\}_{i=1}^L \\ M_2 & \text{if } q_{it}^R = \max\{q_{it}^R\}_{i=1}^{L/j} \\ M_3 & \text{if } q_{it}^R = \max\{q_{it}^R\}_{i=1}^{L/j,k} \\ 0 & \text{otherwise} \end{cases}$$

Here $B = \text{Rs.150}$, $M_1 = \text{Rs.150}$, $M_2 = \text{Rs.100}$ and $M_3 = \text{Rs.75}$.

In the case of piece rates, incentive pay required calculation of a threshold at which the piece-rate kicked in, and calculation of the rate itself. Again, we relied on the fields' manager for guidance. We based the threshold on the upper levels of productivity that the workers achieved during our initial period of fixed pay. Data collection during the initial fixed pay phase supplied an upper output level for an individual worker of approximately 600kg per day; we therefore set this as our threshold. In order to achieve the fields' manager's target of the top workers just being able to double their daily pay but, also, the additional constraint that they should not be able to earn more than under the tournament, we set the piece-rate at $\text{Rs.150}/600\text{kg} = \text{Rs.0.25}/\text{kg}$ – i.e. the average daily pay per kilo achieved by the top picker(s) in the absence of incentives. Therefore, a worker picking 1200kg of fruit in a day would earn $\text{Rs.150} + \text{Rs.0.25} \times (1200 - 600) = \text{Rs.300}$ – the same as the rank tournament winner.

Thus, a worker's daily pay under the piece-rate scheme (P) with threshold $\bar{q}$ and piece-

rate s is

$$Y_{it}^P = B + \begin{cases} s(q_{it}^P - \bar{q}) & \text{if } q_{it}^P > \bar{q} \\ 0 & \text{if } q_{it}^P < \bar{q} \end{cases} \quad (1)$$

In this case, $B = \text{Rs.}150$, $s = \text{Rs.}0.25$ and $\bar{q} = 600\text{kg}$.¹¹

To summarize, the experiment provided data on the output of two groups of fruit pickers who were comparable and who worked on very similar fields in identical conditions throughout the experiment. They were hired by a common fields' manager and were all required to pick fruit of exportable standard. The output of each group was carefully recorded using simple processes and independent auditing (by academic researchers). We ran the tournament before the piece-rate stage and used actual data to set the various parameters of the pay schemes. In addition, we did not reveal these parameters to the workers until a new incentive scheme was introduced. All of these design features sought to minimize the potential for variation in outputs from sources beyond those central to the experiments: piece rates versus rank-order tournaments and the role of tied vs untied employment terms.

3 Empirical Analysis

3.1 Descriptive Analysis

Table 1 shows descriptive statistics for daily workers' productivity. First, notice that average hourly productivity under fixed pay is lower by more than 20kg (i.e. 30%) compared to outcomes under any of the incentive schemes. Second, it appears that overall tied workers were more productive, regardless of the pay scheme in place, both when average productivity is calculated across workers and over the period of the field experiment. While both tied and untied workers report similar levels of kinnow picking experience (the average number of years of experience is 4 and 4.5 years, respectively), this could be explained by underlying (positive) differences in ability between the two groups, or the fact that tied workers had started the fruit picking season earlier, before moving to this particular district – whilst this difference is discernible, it is not central to our ensuing results.

¹¹Notice that our piece-rate scheme differs from that in Bandiera et al's various experiments, since their base pay is also a function of effort (i.e. $B = \theta(q_{it}^P), \theta' > 0$). As explained, our base pay was chosen to replicate typical working conditions (and thus minimize disruption).

It is also interesting to note that the tournament structure led to more dispersion in workers' performance, as can be seen from the standard deviations in Table 1 and the normalized interquartile range (IQR) over the course of the experiment. These are depicted in the left panel of Figure 1, together with the corresponding 3-day moving average, for all workers and for each group, with rank-order pay corresponding to the first 14 days. This is consistent with the idea that the competitive nature of tournaments leads to 'jostling for position' amongst participants. It is notable in the second and third rows of Figure 1 that the patterns of dispersion display some differences across tied and untied workers. In particular, there is evidence that the dispersion amongst tied workers drops towards the end of the two-week tournament stage whereas it moves up and down throughout the fortnight in the case of untied workers. This seems to be mostly driven by a drop in productivity from the less productive tied workers, as can be gleaned from the quintile plots on the right side panel of Figure 1. This may reflect workers learning about their prospects of winning a prize and choosing efforts accordingly. We return to this in Section 4.

< Table 1 >

< Figure 1 >

Table 2 presents unconditional worker productivity by pay schemes and by group, including p -values of a t -test for difference in means, both across incentive schemes and groups. We observe that (unconditional) average aggregate productivity is higher under piece rates by nearly 5 kg/hour, this difference being supported by the non-overlapping confidence intervals of both averages and the difference-in-means test.

< Table 2 >

Equally noticeable distinctions emerge at the group level. Indeed, productivity under piece rates for tied workers is substantially higher than under rank-order pay, and it is also above that for untied workers' under piece rates. For the latter group, average productivity is higher under rank-order pay, though the non-overlapping confidence intervals do not accord with the difference-in-means test. These results are further supported by inspecting Figure 2, which shows kernel density estimates of workers' productivity for each

pay scheme. While there appear to be no substantial differences in aggregate terms for untied workers, there is a clear positive shift in the tied workers' distribution when piece rates are introduced. Also, Figure 3 suggests that aggregating the two groups may not be valid, as they appear to behave distinctively under the different incentive schemes.

< Figures 2-3 >

Table 2 further checks whether the above results could arise from differences in hours worked (the numbers of workers involved in each group was roughly the same). We see that although hours worked vary across pay schemes and groups, this cannot explain differences in productivity reported above. For instance, average hours worked under different incentive schemes in the tied workers group is practically the same, while average productivity is markedly different. On the other hand, under rank-order pay tied workers appear to have worked longer than untied workers, but their average productivity is similar, so one needs to look elsewhere for an explanation for differences in productivity. Naturally, the actual quantity of fruit picked (not reported here) reflects some of the contrasts noticed above.

Next, we provide econometric analysis that takes into account heterogeneity among workers, while also controlling for other specific factors that, besides pay schemes, may be influencing workers' productivity.

3.2 Econometric Analysis

3.2.1 Estimation Framework

In this section, we estimate the impact of different pay schemes on (log) individual worker productivity. The latter is estimated using the regression

$$y_{it} = \alpha_i + \beta ROP_t + \delta C_t + \eta S_t + \tau H_{it} + \lambda \mathbf{D}_t + \gamma t + u_{it} \quad (2)$$

where y_{it} denotes log of worker i 's productivity (kilograms per hours worked) on date t , α_i captures worker fixed effects and ROP_t is a dummy equal to one when rank-order pay is in place and zero when piece rates are paid, so that the parameter β reflects the effect

on productivity of the different pay schemes.¹²

We control for time-varying effects in several ways. First, through the dummy variable C_t , which takes the value 1 when the workforce first starts cutting at a given section in the fields – it is easier to cut the fruit in the periphery of the tree, so we expect productivity to be higher at these times. Second, in the course of the experiment, the two groups worked almost side-by-side for a short period (the first five days of work under rank-order pay). Thus, we define the dummy S_t , which takes the value 1 for these days, therefore capturing potential ‘competition’ effects. Third, we also gathered data on atmospheric conditions such as air temperature, relative humidity and luminosity, collected in $\mathbf{D}_t$, while other unmeasured time-varying factors are captured by a time trend (through γt). Finally, we control for worker stress levels in H_{it} by measuring workers’ heart rates, as well as diastolic and systolic blood pressure, with a bio-scanner device, at the end of the day. We combine these variables into a single stress measure, obtained as their standardized principal component.¹³ Thus, the disturbance term reflects unobservable individual determinants of productivity.

Our panel displays some attrition, as not all individuals worked under the different pay schemes for the duration of the experiment. This may favour balancing the panel but this, in turn, would imply a reduction in the sample size – for example, only 26 workers out of the 72 analyzed here worked for the full period under all pay schemes (the attrition relating mainly to extenuating personal circumstances of the workers). We ran regressions with a balanced panel to check for robustness and also varied the minimum number of days for which workers had to have worked under a pay scheme to be included in the analysis. In both cases, results were qualitatively similar to those presented below. Given these remarks, the results presented here are based on least squares estimation of the full, unbalanced, panel.

¹²One may also employ daily production (in kg) as a measure of workers’ performance; the results are qualitatively similar – see the Table A3 in the Supplementary Appendix.

¹³Using these ‘autonomic’ stress measures separately as controls – either ‘raw’ or relative to cohort or worker average levels – does not change results qualitatively. Regarding H_{it} in (2), it could be argued that there are potential endogeneity issues or that causation might run in the opposite direction – stress levels are an outcome rather than a control. As discussed below, the results suggest these concerns are unfounded.

3.2.2 Baseline results

The top panel of Table 3 displays estimation results for pooled regressions, in which group effects are not separated, while fixed-effects estimation is shown in the lower panel of Table 3, which takes into consideration individuals' heterogeneity, given that poolability tests support the presence of individual effects. When all workers are pooled together, the impact of rank-order pay is negative and always statistically significant, regardless of the control variables used in the estimation. Adding the latter improves estimation fit and leads to an estimated reduction in productivity under rank-order pay of around 25%, from an 8% estimated effect without controls. The 'cut' dummy is significant in both pooled and fixed-effects estimations, while the 'stress' variable helps explain variation in productivity in pooled estimations, but not when heterogeneity is accounted for, as anticipated in footnote 13 (the reverse is true for the time trend and atmospheric controls).

< Table 3 >

A different picture emerges if differences between untied and tied workers are taken into account. Estimating separate regressions for each group confirms the findings in Table 2. Indeed, the productivity of tied workers is negatively affected by rank-order pay, with estimated semi-elasticities (i.e. $\exp(\beta) - 1$) between -20% and -25% , and this variable alone seems to be explaining variations in productivity. In contrast, the pay scheme dummy is never significant (and occasionally positive) in the regressions for untied workers, while field and atmospheric controls drive fluctuations in productivity.

Interestingly, the dummy S_t appears to be capturing a competition effect. It is significant and positive in the joint regression, and the reason for this becomes clearer when the separate regressions are analysed. Indeed, working near to the tied employees raised untied workers' productivity by nearly 30%, while no effect is detected for the tied group.¹⁴ This is akin to the 'clean' peer effects found in Falk and Ichino (2006)'s team setting, in that the proximity of the teams was the only change in the working conditions and resulting productivity levels. We return to this result in the next section.

¹⁴If the (unconditional) average productivity for untied workers is computed when the 'competition' period is *excluded*, then the difference between the two groups is even larger (82.423, compared with 88.429 in Table 1), while it remains similar (93.918) for tied pickers.

3.2.3 Heterogeneous effects

We now consider the effect of rank-order pay in more detail. In particular, we are interested in examining the role played within the tied workers. To see how different workers were affected by rank-order pay, we use quantile regression methods to estimate the conditional distribution of (log) productivity at different quantiles. Defining Q_π as the π th quantile, we now have

$$Q_\pi(y_{it}|\cdot) = \alpha_i + \beta ROP_t + \delta C_t + \eta S_t + \tau H_{it} + \lambda \mathbf{D}_t + \gamma t + u_{it} \quad (3)$$

Results of simultaneous-quantile regression estimates of (3) for each group at different quantiles, with and without controls, are reported in Table 4. As in Bandiera et al. (2010 and 2013), we allow for fixed effects, with bootstrapped standard errors based on 5,000 replications (estimation without fixed effects yields similar results).

< Table 4 >

The results suggest that the average negative effects of rank-order pay uncovered previously are distinct across the conditional distribution of productivity, in particular for tied workers. More precisely, the magnitude of this pay scheme effect is larger (i.e. more negative) and very significant for workers in the lower tail of the conditional distribution, even more so when estimation includes the control variables.

The left panel of Figure 4 illustrates the changes of the effect throughout the conditional distribution of productivity for tied workers, with the associated 90% bootstrapped confidence interval; this ranks workers by their output under piece rates, then plots the percentage change in outputs achieved by these workers under rank-order pay. Rank-order pay reduces productivity of workers in the 10th quantile (of piece-rate output) by 50%, while for workers in the 25th quantile productivity is reduced by a third. Tests of differences between adjoining and extreme quantiles are always significant, but this effect is then attenuated, increasing monotonically as one moves along the conditional distribution, with the rank-order pay dummy becoming insignificant for workers whose conditional productivity is above the 75th quantile. Although rank-order pay appears to have a positive effect on the most productive workers, the estimates are imprecise and not different from zero.

< Figure 4 >

In contrast, estimates for the untied group are mostly positive, but they are not statistically significant across all quantiles, as can be seen in the right panel of Figure 4. This, in turn, largely confirms the previous result that rank-order pay, relative to piece rates, has little additive effect on untied workers' productivity: wherever they are ranked in the piece-rate output distribution, their output is similar under rank-order pay. Interestingly, unreported elements of the quantile regressions also reveal that less productive workers react more to working conditions, i.e. the magnitude of the field controls is larger the lower the quantile, in particular for the 'cut' and the 'competition' dummies. Perhaps the less able workers were more sensitive to the working environment than their more able counterparts.

3.2.4 Testing the 'discouragement' effect

The preceding results suggest the presence of a potential 'discouragement' effect of rank-order pay amongst tied workers.¹⁵ Our setup allows us to perform a simple check on this hypothesis. If the distribution of abilities is, to a certain extent, known to the workers, then they can, in principle, adjust their efforts dynamically depending on how close they are to the winning harvest. More specifically, for tied workers (who appear to have been more susceptible to this effect), we would expect to find a negative relationship between a worker's productivity in day t and the difference between this worker's collection and the prize-winning pick in the previous day, such that

$$y_{it} = f(y_{t-1}^* - y_{i,t-1}) + \alpha_i + \xi X_{it} + \nu_{it} \quad (4)$$

where y_{t-1}^* denotes the previous day's winner's productivity and X_{it} is the group of controls described in Section 3.2.1.¹⁶

< Figure 5 >

As for the shape of $f(\cdot)$ in (4), Figure 5 plots tied workers' productivity on day t against their difference in productivity compared to the winner on day $t - 1$. The plot appears to

¹⁵We thank an anonymous referee for helpful suggestions on this section.

¹⁶One could also specify a response of productivity to differences in actual pickings (in kg) in $t - 1$ rather than productivity differences, the results being qualitatively similar.

confirm the discouragement mechanism hypothesized above and hints at its nonlinearity.¹⁷ Table 5 shows the empirical implementation of this conjecture for both tied and untied workers. We initially regress productivity on the simple difference, but a RESET test suggests that a simple linear specification (with or without controls) will be misspecified, so we focus on a quadratic specification for $f(\cdot)$.¹⁸ The results in Table 5 show that the apparent nonlinear negative response of tied workers to ‘differences to yesterday’s winner’ is reasonably well captured by the quadratic specification. Note also that the sample for the untied workers’ regressions are slightly larger and the fit much better relative to the tied workers’, so the absence of a ‘discouragement’ effect for untied workers is not due to weaker sampling. Overall, this simple check reinforces our explanation for the distinct behaviour between tied and untied workers. We discuss this further in the next section.

< Table 5 >

4 Discussion

We have discovered a number of interesting results about the operation of financial incentives in our developing economy setting.

The effects of incentive pay To begin, we find that financial incentives (whether piece rates or tournaments) outperform a simple fixed pay scheme. Whilst this is consistent with an incentive story present in other studies, we need to check that it is not an artefact of our study design. In particular, perhaps workers were not fully acclimatized to our presence during the fixed pay period; or they viewed our presence as a performance monitoring device once we turned to incentive pay. In fact, both of these seem unlikely. As we have said, the fixed pay output we recorded was collected towards the end of the first phase of work, when the workers were used to our presence. The low standard deviation of fixed pay output in Table 1 indicates that output from that period did not change dramatically so there is little evidence of productivity being held back by an ‘acclimatisation’ effect immediately

¹⁷A more elaborate exercise would involve verifying the possibility of a *positive* response in productivity when deviations from the winning pickings are small (i.e. an ‘encouragement’ effect for pickers capable of competing with the best), but this would require defining an *ad-hoc* threshold to determine small vs. large differences and, more significantly, it would leave us with a sample too small to allow us to draw any meaningful conclusions.

¹⁸The RESET test is carried out by adding polynomials of the fitted values of (4) (including the fixed effects).

following our arrival. Turning to monitoring, our experiment design sought to make our assistants as unobtrusive as possible – they mainly recorded weights at some distance from the picking itself. In addition, if the workers felt monitored, it is perhaps harder to explain why our different incentive schemes produced some differences in outcomes, and why we also observed a ‘competition effect’ when the workers were in close proximity.

Another possible concern is that time-varying effects on productivity may have played a role in the changes we see across pay schemes. Ideally, operating a fixed pay control throughout the experiment would have addressed this but, as we have seen, Widget’s preference for minimal disruption and equal treatment of workers prevented such a feature. We believe, however that we can rule out a number of time-varying effects: our fields, trees and terrain were similar across groups and time (with equal preparation of the soil for obvious commercial reasons); export market standards meant that no variation in quality requirements contaminated our results; climate remained similar throughout, as did the identity of most of the workers. We have already described how our initial period of fixed pay (and careful liaison with line managers) assisted worker acclimatisation and, finally, we have sought to strip out additional time-varying effects via inclusion of a time trend.

Taking these various remarks together, we think it is reasonable to attribute higher outputs to the incentives inherent to our piece rate and tournament schemes.

Piece rates versus Tournament pay We also find that piece rates are a more successful stimulant to productivity than tournaments (with an 8-25% range of improvement taking place across our estimates).¹⁹ The effects of the piece rates could be the result of $s > 0$ raising the marginal cost of leisure (and hence stimulating work – as in Bandiera et al, 2005). Alternatively, given the possible presence of moral hazard in the absence of incentives, the presence of $s > 0$ may create an efficiency wage effect that overcomes shirking behaviour to the benefit of output. The relative role of each mechanism would be interesting future research.

Effects of incentive pay on different worker groups Perhaps our most distinctive results arise when we try to examine the relative difference between the performance under piece rates and tournaments by distinguishing amongst worker-types. The contractual

¹⁹In the Supplementary Appendix, we show that this increase in productivity took place without increasing the cost of the piece-rate scheme relative to the rank-order tournament, so it is not the result of different budgets being available across the two schemes; see Row (a) of Table A2.

arrangements of the workers could have had at least two effects on our results: tied contracts may attract better workers, and the features of their contractual set-up may induce different reactions to the financial incentives on offer; these effects are not inconsistent with each other. The first is borne out by the fact that, overall, tied workers appear to have had higher productivity. Interestingly, however, the effect is larger under piece pay, which perhaps suggests a different magnitude of response to the incentives on offer. Our quantile regressions confirm that this is the case, with only the most productive of the tied workers maintaining their productivity across the two schemes; lower quantile tied workers actually reduced their productivity in the tournament setting. The same is not true of the untied workers, whose productivity remained similar across the schemes for each quantile.²⁰ We now consider how this difference might be explained, and draw some implications for tournament literature.

4.1 Explanations

4.1.1 Dynamic tournaments with learning

As stated earlier, the tournament we ran was a repeated simultaneous-move tournament in which the players discovered the identity of the (three) winners on a next-day basis. In general, such tournaments do not always maximize output/effort because of the effects of the feedback the players receive (implicitly and explicitly) on their performance. For example, Krämer (2007) presents a model where players do not know their innate abilities and in which it is possible for the ‘better’ player to read too much into a bad outcome, assume s/he is weak and, thus, stop trying to win. This can become self-enforcing and the resulting equilibrium sees the least productive player winning in every period, and output consequently below its potential. In contrast, Münster (2009) describes a ‘sandbagging’ equilibrium where the better party tries to lose the first round tournament for fear of giving away too much information to a planner in the second. Konrad (2009) describes a ‘discouragement effect’ whereby players compete over several rounds for an ultimate prize (e.g. a league champions’ trophy). Here it is possible that the efforts supplied in each round dissipate the value of winning and, hence, reduce effort incentives at the

²⁰Our Supplementary Appendix shows that this effect led the introduction of piece rates to be profitable for Widget, while the drop in tied workers’ output under rank pay saw overall revenue drop just below costs for this scheme; see Row (b2) of Table A2.

margin. Gürtler and Harbring (2010) demonstrate that principals may optimally respond by feeding back only limited performance data to tournament participants in order to limit discouragement. Various empirical studies have also examined repeated tournaments, with, for instance, Bandiera et al. (2013) and Berger and Pope (2011) finding positive effects in *team* settings²¹, but Gill and Prowse (2012) finding more negative results in a sequential tournament where the first player enjoys a first mover advantage over the opposing player. Delfgaauw et al. (2014) also find that feedback of poor performance harms final tournament results in a setting where prizes are awarded after the final round of ‘play’.

Although these are indicative of potential difficulties in repeated tournaments, the above settings do not quite match ours: for example, we have *individual* players and simultaneous moves; arguably, our players know something about their abilities (they all have prior picking experience); there is no social planner and there are repeated prizes on offer. Nonetheless, like the above, it is possible that our results reveal an interesting feature of repeated tournaments. It seems reasonable that our tied workers are more familiar with the distribution of abilities within their group than the untied ones: they have worked together for a much longer period. As such, having observed the identities of the winners, the weaker workers in the tied group may quickly realize that their chances of winning a prize are small, with the consequent effect that their output soon drops; in contrast, a group of high ability workers win the prizes on a regular basis and are therefore incentivized to compete. This is, to a certain extent, similar to the results in Brown (2011), who documents an ‘adverse’ incentive effect in tournaments whereby the benefits of competition depend critically on workers’ relative abilities, with large differences in skill resulting in effort reduction (see also Singh and Masters, 2018). In contrast, the workers in our untied group, having no common work experience, are potentially less familiar with each others’ abilities, and this may render the identities of winners a noisier signal of relative ability. As such, less able pickers may take longer to realize that there is little point in competing for the prizes.²²

²¹The team mechanisms studied by these papers are the effects of trading players on the basis of observed performance (Bandiera et al.) and the need to specialize and support each other in order to achieve a single goal (Berger and Pope, 2011).

²²Recall that the rank-order tournament stage preceded the piece-rate one, so these workers could only have learnt about abilities from the fixed pay stage. Yet, we collected no individual data there and, as already argued, self-monitoring of co-workers was difficult to achieve. Thus, the workers had little opportunity to learn about the distribution of talents in the untied group until the tournament stage started.

This interpretation is consistent with our quantile results. It is also consistent with Figure 1's suggestion that the variance of output under rank-order pay may eventually have dropped for tied workers but not for untied ones – perhaps reflecting fewer effective tied competitors over time. To explore the explanation further, we considered the dispersion of ‘winners’ in the two incentive schemes (i.e. prize-winners in the tournaments and those who exceeded the threshold, $\bar{q}$, under piece rates).²³ The results support the hypothesis that winners (from both groups) are densely distributed, but are especially concentrated for rank-order pay, and for tied workers.²⁴ This is readily seen from the Herfindahl-Hirschman indices of winner concentration presented in Table 6.²⁵

< Table 6 >

This interpretation of our results suggests a direction for future research on dynamic tournaments. In particular, it suggests that the role of individual workers’ uncertainty about their rivals’ productivities in tournaments which involve repetition of simultaneous moves and regular prizes warrants further investigation. Krämer (2007) presents a relevant model but our evidence is not consistent with his explanation for tournament underperformance.²⁶ An implication is that situations where abilities are better known amongst workers may suffer negative feedback effects that lead to greater reductions in productivity. As a corollary, we would also expect that such reductions will take longer to arise in situations where abilities are not well known amongst the rivals.²⁷

²³We calculated winning frequencies for workers who were present for at least three days under each scheme, so as to rule out potential outliers and minimize any selection bias arising from the few untied workers who did not work under both schemes.

²⁴This result may also suggest that, under piece rates, our tied workers did not appear to suffer from a ‘ratchet effect’ where the best tied workers hold back because they are worried about facing stiffer targets in future if they reveal their ability. The fact that the same tied people also win the tournaments suggests that the best workers were not holding back (for ratchet reasons) in the piece-rate scheme.

²⁵These are calculated as follows. Let the number of times worker i from group $j = \{T, U\}$ ‘wins’ under incentive scheme $u = \{\text{Rank, Piece}\}$ be N_{ij}^u . Then the Herfindahl-Hirschman index is calculated as $HH_j^u = \sum_i (N_{ij}^u / \sum_i N_{ij}^u)^2$. Higher values of HH_j^u imply a more concentrated set of winners.

²⁶Krämer (2007) suggests that the most able workers underperform to avoid drawing adverse inference from a loss. We can examine this in our data if we assume (reasonably) that our piece-rate experiment gives no incentives for the best workers to shirk. If the top piece-rate workers are also typically the winners of the tournaments, this suggests that the better workers are winning the tournaments in contrast to the mechanisms in Krämer (2007). This is what our quantile regressions found.

²⁷On this reading of our results, we would eventually have seen a reduction in the untied workers’ outputs had we run our tournaments for more days.

4.1.2 Alternative explanations

Of course, there may be other explanations for the reductions in tied workers' output under rank-order pay. We consider five of these (in addition to Krähmer (2007)'s discussed above).

One possibility is that the tied workers, being more familiar with each other having worked together for longer, collude to allow weaker workers to win on occasions and that this lowers output. Yet this would suggest quite a wide pool of winners, which Table 6 confirms that we do not see. Collusion could also manifest itself if workers have a strong enough incentive to deviate jointly to low effort levels (while still keeping their relative effort levels unchanged). However, as pointed out by Gürtler (2010), collusion is less likely when there is sufficient skills heterogeneity, such that very productive workers gain from deviating to higher effort levels, or when workers are equally productive, but the winner's prize is relatively large – as it happens, both of these are features of our field experiment.

In fact, our quantile regression results suggest that the most productive workers maintain their productivity even under rank-pay, which implicitly rules out collusive behaviour. An additional simple check can be carried out by comparing the top outputs under rank-pay and piece-rates. Under collusion with rank-pay, we would expect the average of top pickings to be significantly lower than under piece rates. However, this is not the case: t -tests of the null hypothesis that the (unconditional) means of top outputs for rank-pay and piece-rates are the same delivers a p -value of 0.27, if one considers the top two prizes, the p -value is 0.32, while for the mean of the top three winners the p -value is 0.26. Thus, there is no significant difference in effort (or productivity) at the top end under each pay scheme, a finding that is inconsistent with the presence of collusion under rank-pay. It would be interesting, however, to know whether social concerns lead to prize sharing through *ex post* division of the gains: this would be valuable future research.²⁸

An alternative explanation for the differences between tied and untied pickers' tournament performance would be that the latter are maximising income for the (relatively short) period of time they are employed, regardless of the pay scheme in place. This would then imply that variation in per-worker daily productivities should be lower, as reported in Table 1 (a test of equal variances rejects the null at 5%). Nevertheless, such behaviour is

²⁸Bandiera et al (2005) also look for evidence of 'social concerns' amongst workers paid by incentive pay, though in the strategies of the players during the game rather than in the *ex post* division of surplus.

hard to reconcile with the results for the competition dummy (S_t) reported above. Apparently, competition provided a real incentive (productivity rose by nearly 30% over a very short period), suggesting that at least some untied workers were not always operating at their maximum effort.

Perhaps instead untied workers take a break after winning a prize. This would tend towards a wider variation in the number of prize winners when compared to tied workers, even if the untied workers knew each others' abilities (contrary to our suggestion) – see Table 6. Our data do not support this as the prize winners amongst the tied workers remained throughout the two weeks of tournament pay.

A fourth possibility is that the best tied workers expect their less able counterparts to dampen their efforts (eventually) and, as such, the best ones can also afford to lower their outputs at the margin. This would then see tied outputs fall under rank-order pay. If this were the case, however, *all* tied workers' quantiles would exhibit negative coefficients on the *ROP* dummy in our quantile regressions (assuming that the outputs of the top workers under piece rates reflect their best efforts): we do not observe this. In fact, following Szymanski and Valletti (2005), the set-up of our tournament also mitigates against this because the presence of additional prizes has the effect of incentivising the best workers to compete hard amongst themselves (see also Moldovanu and Sela, 2001). In this respect, it is intuitive that we appear to have five or six tied workers who compete throughout for the three prizes on offer.

Finally, another possibility is that there is a higher proportion of talented untied workers (relative to their peers) than tied ones. Thus, even if the untied workers knew each others' abilities (contrary to our suggestion) there would be many who could win and this would keep productivity high. In fact, the unconditional distributions of productivity under piece rates (where workers are only competing against themselves) do not support this: the right panel of Figure 3 presents no evidence of a thicker righthand tail for untied workers than for tied ones.

4.2 Incentives and welfare

An interesting implication of our results relates to the possible welfare effects of rank-order tournaments. Layard and Nickell (2011), for example, have raised the possibility that

excessive competition can have a demotivating effect on workers, causing both output and their own happiness and self-esteem to suffer.²⁹

To consider this, we interviewed workers about their preferences over pay schemes at the end of the experiment. We asked them to rank the schemes from most preferred (first) schemes to least preferred (third). Figure 6 shows preferences over pay schemes across workers, where the numbers represent the inverse of the average score by scheme and worker-type, so that unity denotes top preference. We find that, on average, tied workers clearly preferred the piece-rate over the tournament system, while untied ones slightly preferred the former. We also ran simple regressions to see whether our basic measure of worker stress H_{it} could be linked to the various incentive schemes we operated: see Table 7. The rank-order tournaments led to higher levels of stress amongst both groups of workers (though, interestingly, given their preferences, especially amongst untied ones), a result that is consistent with Layard and Nickell’s arguments above. Clearly, more research is needed in understanding the interactions between pay schemes, workers’ preferences and stress, but, overall, our workers appear to have been less comfortable with a tournament-type structure that placed them in direct competition with each other.

< Figure 6 >

< Table 7 >

5 Conclusion

This paper provides evidence on the impact of rank-order tournament and piece-rate incentive schemes on the productivity of fruit pickers in Pakistan. Such incentives are not commonly observed, yet, given the economic significance of this activity, it is important to explore whether they have any effect on performance here, and across many settings. We find that incentive pay (of both kinds) is, indeed, effective in raising productivity compared to a simple fixed pay alternative and that this may have appreciable effects on the speed (and quality) of picking, and revenues. We also find, however, that the effects of such pay is nuanced. Thus, there are differences between the effects of the schemes and between their

²⁹An interesting literature from sports contests also suggests that competition can raise stress levels amongst participants – see Mellalieu et al. (2009).

impacts on two groups of workers distinguished by the types of contract that are common in our setting. In particular, tournament pay appears to discourage lower-ranked workers despite incentivising their ‘top’ counterparts; in turn this dampens the effects of rank-pay on overall productivity and may reduce its benefits relative to the cost of operating the scheme. This happens when workers are ‘tied’ to the employer for a longer period of time. We suggest that such workers might know more about each others’ abilities than workers employed on a shorter term (‘untied’) basis, enabling them to learn quickly about their chances of winning the simultaneous-move, repeated tournaments that we run. As a result, lower ability tied workers may be discouraged from working hard. Given this finding, we argue that our results support recent theoretical and empirical literature that emphasizes the problems caused by negative feedback in tournament settings featuring individuals (as opposed to teams), though we also note that our results are not readily explained by existing examples of such models. Instead, we suggest that a direction for further research on dynamic tournaments could usefully examine the ways in which participants’ uncertainty about relative abilities might affect individual incentives.

In addition to these results we also find potentially important evidence of a role for non-financial motivation amongst our less specialized (‘untied’) workers in the sense that their productivity improved when they were in close proximity to (though not direct competition with) their more experienced (‘tied’) counterparts. We suggest that there is valuable future research to be carried out into what causes such motivation and whether it arises in other settings.

Our results may also shed light on possible reasons for the low take-up of incentive pay in such settings. These are to do with recognition by owners that they do not possess suitable skills and equipment to implement such schemes, and potential worries that some owners have about the ways in which their business environment could change if performance pay was widely dispersed across their industry – this despite apparent gains in output that we observed under incentive pay. It is clearly important to understand the basis for these perceptions if incentives are to become a regular feature; it is possible that the positive business case we find may help here.

Of course, the finding that incentive pay can raise output is not the same as finding that it has positive welfare benefits (at least, not in an imperfect market where compensation

for hardship suffered is not necessarily paid from the productivity gains). A number of authors have worried about the demoralising effects of competition in the workplace and, indeed, it is possible that some of our workers felt this. Interestingly, such problems may extend beyond those workers directly affected: the field manager in our study reported to us after the experiments that during the rank system there was greater animosity towards him, as workers felt that they were being unfairly treated – a ‘fairer’ tournament structure could include the handicap system proposed in Lazear and Rosen (1981), see also Gürtler and Kräkel (2010). These issues are also central to the actual take-up of performance pay, and the desirability of doing so, and we see them as an important next step for our research agenda.

Supplementary material

Supplementary material is available online at the OUP website. These are the data and replication files and the supplementary appendix.

Funding

This work was supported by the Nuffield Foundation [SGS/33719].

Acknowledgements

We are grateful to Abhijit Banerjee, Jo Blanden, Alex Gershkov, Anil Jain, Mushtaq Khan, Khartik Muralidharan, Imran Rasul, Paul Schweinzer, Roman Sheremeta and Chris Woodruff for valuable comments.

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**, 917-962.
- Bandiera, O., Barankay, I. and Rasul, I.** (2010) Social Incentives in the Workplace, *Review of Economic Studies*, **77**, 417-458.

- Bandiera, O., Barankay, I. and Rasul, I.** (2011) Field Experiments with Firms, *Journal of Economic Perspectives*, **25**, 63-82.
- Bandiera, O., Barankay, I. and Rasul, I.** (2013) Team Incentives: Evidence from a Firm Level Experiment, *mimeo*, University of Pennsylvania.
- Banerjee, A. and Duflo, E.** (2011) *Poor Economics: A Radical Rethinking of the Way to Fight Global Poverty*, PublicAffairs, New York.
- Bardhan, P.** (1979) Wages and Unemployment in a Poor Agrarian Economy: A Theoretical and Empirical Analysis, *Journal of Political Economy*, **87**, 479-500.
- Bardhan, P.** (1980) Interlocking Factor Markets and Agrarian Development: A Review of Issues, *Oxford Economics Papers*, **32**, 82-98.
- Berger, J. and Pope, D.** (2011) Can Losing Lead to Winning? *Management Science*, **3**, 86-99.
- Brown, J.** (2011) Quitters Never Win: The (Adverse) Incentive Effects of Competing with Superstars, *Journal of Political Economy*, **119**, 982-1013.
- CABI** (2008) Kinnow Value Chain: Knowledge Gaps and ICT Prevalence in the Chain, *mimeo*, CABI South Asia, Ref: FR/CABI-PK/DR60007/1109, December.
- Cason, T., Masters, W. and Sheremeta, R.** (2010) Entry into Winner-Take-All and Proportional-Prize Contests: An Experimental Study, *Journal of Public Economics*, **94**, 604-611.
- Delfgaauw, J., Dur, R., Non, A. and Verbeke, W.** (2014) Dynamic Incentive Effects of Relative Performance Pay: A Field Experiment, *Labour Economics*, **28**, 1-13.
- Dechenaux, E., Kovenock, D. and Sheremeta, R.** (2015) A Survey of Experimental Research on Contests, All-Pay Auctions and Tournaments, *Experimental Economics*, **18**, 609-669.
- Duflo, E., Kremer, M. and Robinson, J.** (2011) Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya, *American Economic Review*, **101**, 2350-2390.

- Engelland, A. and Riphahn, R.** (2005) Temporary contracts and employee effort, *Labour Economics*, **12**, 281-229.
- Falk, A. and Ichino, A.** (2006) Clean Evidence on Peer Effects, *Journal of Labor Economics*, **24**, 39-57.
- Gill, D. and Prowse, V.** (2012) A Structural Analysis of Disappointment Aversion in a Real Effort Competition, *American Economic Review*, **102**, 469-503.
- Grund, C. and Sliwka, D.** (2005) Envy and Compassion in Tournaments, *Journal of Economics and Management Strategy*, **14**, 187-207.
- Gürtler, O.** (2010) Collusion in homogeneous and heterogeneous tournaments, *Journal of Economics*, **100**, 265-280.
- Gürtler, O. and Harbring, C.** (2010) Feedback in Tournaments under Commitment Problems: Experimental Evidence, *Journal of Economics and Management Strategy*, **19**, 771-810.
- Gürtler, O. and Kräkel, M.** (2010) Optimal tournament contracts for heterogeneous workers, *Journal of Economic Behavior and Organization*, **75**, 180-191.
- Konrad, K.** (2009) *Strategy and Dynamics in Contests*, Oxford University Press.
- Krähmer, D.** (2007) Equilibrium Learning in Simple Contests, *Games and Economic Behaviour*, **59**, 105-131.
- Layard, R. and Nickell, S.** (2011) *Combating Unemployment*, Oxford University Press, New York.
- Lazear, E. and Rosen, S.** (1981) Rank-Order Tournaments as Optimum Labor Contracts, *Journal of Political Economy*, **89**, 841-864.
- Macchiavello, R. and Morjaria, A.** (2015) The Value of Relationships: Evidence from a Supply Shock to Kenya Flower Exports, *American Economic Review*, **105**, 2911-2945.
- Mellalieu, S., Neil, R., Hanton, S. and Fletcher, D.** (2009) Competitions Stress in Sport Performers: Stressors Experienced in the Competition Environment, *Journal of Sports Sciences*, **27**, 729-744.

- Moldovanu, B. and Sela, A.** (2001) The Optimal Allocation of Prizes in Contests, *American Economic Review*, **91**, 542-558.
- Muralidharan, K. and Sundararaman, V.** (2011) Teacher Performance Pay: Experimental Evidence from India, *Journal of Political Economy*, **119**, 39-77.
- Münster, J.** (2009) Repeated Contests with Asymmetric Information, *Journal of Public Economic Theory*, **11**, 89-118.
- Naseer, M.** (2010) Report on Export of Kinnow, *mimeo*, Trade and Development Authority of Pakistan, Government of Pakistan, January 14.
- Prendergast, C.** (1999) The Provision of Incentives in Firms, *Journal of Economic Literature*, **37**, 7-63.
- Qureshi, K., Laghari, M., Khokhar, K. and Khan, M.** (1993) Performance of Mandarin (*Citrus Reticulata Blanco*) Varieties in Islamabad, *Pakistan Journal of Agricultural Science*, **30**, 424-427.
- Sheremeta, R.** (2016) The Pros and Cons of Workplace Tournaments, *IZA World of Labor*, **302**, 1-10.
- Singh, P. and Masters, W. A.** (2018) Performance bonuses in the public sector: Winner-take-all prizes versus proportional payments to reduce child malnutrition in India, *Journal of Development Economics*, in press.
- Szymanski, S. and Valletti, T.** (2005) Incentive Effects of Second Prizes, *European Journal of Political Economy*, **21**, 467-481.

Tables

Table 1: Descriptive statistics for daily workers' productivity (kg/hour)

		Mean	Std. Dev.	Min.	Max.	Obs.
Fixed	All workers	70.55	24.06	43.71	103.9	14
Rank-order + piece	All workers	91.88	44.93	3.18	310.05	1076
Rank-order pay	All workers	89.62	47.33	3.18	310.05	531
Piece-rate	All workers	94.07	42.43	18.38	255.70	545
Overall	Tied	96.46	47.41	17.91	295.50	462
	Untied	88.43	42.69	3.18	310.05	614
per worker	Tied	98.75	35.30	35.81	206.08	30
	Untied	80.32	32.32	26.60	165.56	42
per day	Tied	95.52	10.99	76.41	120.33	28
	Untied	88.59	16.09	68.58	125.14	28

Source: Authors' calculations.

Note: For fixed-pay we have data on total daily collections.

Table 2: Differences in unconditional workers' productivity (kg/hour) per pay scheme

		Tied	Untied	Diff-in-mean (<i>p</i> -value)	All workers
Productivity	Rank-order pay	88.26 (3.19)	90.56 (2.69)	0.291	89.62 (2.05)
	Piece-rate	103.78 (2.98)	86.21 (2.13)	0.000	94.07 (1.82)
	<i>p</i> -value t-test	0.000	0.103		0.052
Hours worked per day	Rank-order pay	4.97 (0.10)	4.35 (0.08)	0.000	4.60 (0.06)
	Piece-rate	4.98 (0.07)	4.90 (0.05)	0.201	4.94 (0.04)
	<i>p</i> -value <i>t</i> -test	0.467	0.070		0.000

Source: Authors' calculations.

Note: standard errors in brackets.

Table 3: Impact of Pay Schemes on workers' productivity (pooled and fixed effects)

	All workers with controls		Tied workers with controls		Untied workers with controls	
<i>Pooled regressions</i>						
Rank-order pay	−0.084** (0.042)	−0.162* (0.097)	−0.200*** (0.042)	−0.286*** (0.098)	0.008 (0.057)	0.018 (0.123)
Cut dummy		0.107*** (0.034)		0.034 (0.049)		0.211*** (0.049)
Stress variable		0.073*** (0.025)		0.047 (0.040)		0.086*** (0.027)
Temperature		0.063 (0.038)		0.042 (0.208)		0.051** (0.025)
Light		−0.008 (0.006)		0.024 (0.036)		−0.011** (0.005)
Humidity		0.003 (0.002)		−0.005 (0.018)		−0.004 (0.003)
Time trend		0.000 (0.006)		−0.004 (0.008)		0.008 (0.007)
Competition dummy		0.057 (0.078)		0.107 (0.299)		0.132 (0.096)
Adj. R^2	0.007	0.059	0.040	0.036	0.000	0.083
No. observations	1076	1060	462	459	614	601
No. workers	72	70	30	29	42	41
<i>Fixed effects</i>						
Rank-order pay	−0.075*** (0.023)	−0.254*** (0.065)	−0.196*** (0.041)	−0.297*** (0.089)	0.018 (0.034)	0.104 (0.135)
Cut dummy		0.118*** (0.029)		0.060 (0.048)		0.221*** (0.041)
Stress variable		0.012 (0.016)		−0.006 (0.021)		0.023 (0.024)
Temperature		0.096*** (0.187)		−0.146 (0.155)		0.032 (0.030)
Light		−0.011*** (0.004)		0.012 (0.041)		−0.011** (0.004)
Humidity		0.001 (0.002)		0.013 (0.012)		−0.007** (0.003)
Time trend		−0.004 (0.004)		−0.009 (0.006)		0.011 (0.007)
Competition dummy		0.111** (0.048)		−0.120 (0.269)		0.250*** (0.072)
Adj. R^2	0.007	0.047	0.040	0.035	0.000	0.071
No. observations	1076	1060	462	459	614	601
No. workers	72	70	30	29	42	41

Source: Authors' calculations.

Note: *** indicates significance at 1%, ** at 5% and * at 10%; some observations are dropped when health variables are used, as some measurements were occasionally not taken.

Table 4: Quantile regressions, rank-order pay dummy (β) estimates

<u>Tied workers</u>	Quantiles					$Q_{10} = Q_{90}$ (p -value)
	10th	25th	50th	75th	90th	
No controls	−0.357*** (0.069)	−0.186*** (0.049)	−0.153*** (0.039)	−0.129*** (0.039)	−0.061 (0.047)	0.000
With controls	−0.686*** (0.105)	−0.410*** (0.096)	−0.188** (0.083)	−0.017 (0.091)	0.078 (0.161)	0.000
<u>Untied workers</u>						
No controls	−0.035 (0.070)	−0.026 (0.050)	0.003 (0.037)	0.048 (0.038)	0.048 (0.041)	0.287
With controls	0.039 (0.216)	0.014 (0.156)	0.022 (0.170)	0.128 (0.178)	0.075 (0.255)	0.734

Source: Authors' calculations.

See notes to Table 3; the reported results are obtained by estimating simultaneously each quantile π as $Q_{\pi}(y_{it}|\cdot) = \alpha_i + \beta ROP_t + \delta C_t + \eta S_t + \tau H_{it} + \lambda \mathbf{D}_t + \gamma t + u_{it}$; bootstrapped standard errors based on 5000 replications; the last column displays the p -value of a test of an equal pay dummy coefficient at quantiles 10 and 90.

Table 5: Estimates of the ‘discouragement’ effect

	Tied	Untied	Tied	Untied	Tied	Untied	Tied	Untied
$y_{t-1}^* - y_{i,t-1}$	-0.033** (0.014)	0.019 (0.012)	-0.090*** (0.023)	0.012 (0.028)	-0.059** (0.028)	-0.105 (0.066)	-0.077** (0.027)	-0.060 (0.086)
$(y_{t-1}^* - y_{i,t-1})^2$					$3.7 \times 10^{-5**}$ (1.4×10^{-5})	$8.4 \times 10^{-5**}$ (3.78×10^{-5})	$5.1 \times 10^{-5***}$ (1.4×10^{-5})	4.8×10^{-5} (4.7×10^{-5})
Controls	No	No	Yes	Yes	No	No	Yes	Yes
RESET	0.000	0.000	0.000	0.169	0.765	0.604	0.178	0.191
Adj. R^2	0.010	0.018	0.021	0.193	0.015	0.070	0.047	0.206
No. Obs.	190	246	187	241	190	246	187	241

Source: Authors’ calculations.

See notes to Table 3; The dependent variable is daily workers’ productivity y_{it} ; p -values reported for the RESET test with two polynomial terms.

Table 6: Herfindhal-Hirschman Index for the Share of Wins Frequency

		Index
Tied	Rank-order pay	0.29
	Piece-rate	0.09
Untied	Rank-order pay	0.17
	Piece-rate	0.15

Source: Authors’ calculations.

Note: Index computed assuming full participation by pickers in each team and includes workers with a minimum of 3 days worked in each scheme.

Table 7: Impact of pay schemes on workers' stress levels

	All workers	Tied workers	Untied workers
Rank-order pay dummy	0.475*** (0.143)	0.342** (0.159)	1.488*** (0.266)
Cut dummy	-0.002 (0.063)	0.105 (0.098)	-0.103 (0.089)
Temperature	0.205*** (0.044)	-0.012 (0.186)	-0.020 (0.059)
Light	0.021** (0.009)	0.021** (0.010)	-0.073 (0.064)
Humidity	0.013** (0.005)	-0.021 (0.016)	0.003** (0.012)
Time trend	0.003 (0.007)	0.009 (0.011)	0.026 (0.012)
Adj. R^2	0.081	0.045	0.226
No. observations	1060	459	601
No. workers	70	29	41

Source: Authors' calculations.

See notes to Table 3; The dependent variable is the stress variable H_{it} discussed in section 3.2.1.

Figures

Figure 1: Dispersion of Productivity (IQR/Mean) and Bottom vs Top Quintiles: All workers, Tied and Untied.

Source: Authors' calculations.

Figure 2: Productivity: kernel density estimates.
Source: Authors' calculations.

Figure 3: Productivity kernel density estimates: Rank-order pay vs. Piece rates; Tied vs. Untied workers.
Source: Authors' calculations.

Figure 4: Quantile regressions estimates of the pay scheme effect: Tied (left) vs. Untied (right) workers.
Source: Authors' calculations.

Figure 5: Discouragement effect.
Source: Authors' calculations.

Figure 6: Workers' Preferences over pay schemes.
Source: Authors' calculations.