Peers and Motivation at Work:
Evidence from a Firm Experiment in Malawi

Lasse Brune
Northwestern University

Eric Chyn
University of Virginia

Jason Kerwin
University of Minnesota

December 5, 2017

Abstract: This paper sheds light on the nature of workplace peer effects by analyzing an experiment with a tea estate in Malawi. We randomly allocate tea-harvesting workers to locations on fields to estimate the impact of peers on worker performance. Using data on daily productivity, we find strong evidence of positive effects from working near higher-ability peers. Our estimates show that increasing the average of co-worker ability by 10 percent increases own-productivity by about 0.3 percent. Since workers receive piece-rates and do not work in teams, neither production nor compensation externalities drive peer effects in our setting. In additional analysis, we find evidence against learning or worker socialization as mechanisms. Rather, results from an incentivized choice experiment suggest that workers view their co-workers as a source of "motivation." When given a choice to be re-assigned, the majority of workers want to work near fast (high-ability) coworkers. In open-ended survey responses, workers with demand for high-ability peers state that working near faster peers provides motivation to work harder.

Acknowledgements: Brune (corresponding author): Center for the Study of Development Economics, Northwestern University, Evanston, IL 60208 (lasse.brune@yale.edu). Chyn: Department of Economics, University of Virginia, P.O. Box 400182, Charlottesville, VA 22904 (ericchyn@virginia.edu). Kerwin: Department of Applied Economics, University of Minnesota, 1994 Buford Avenue, St. Paul, MN 55108 (jkerwin@umn.edu).We are grateful for feedback and guidance from Martha Bailey, Charlie Brown, Brian Jacob and Jeff Smith. We also received insightful comments from Emily Breza, Dean Karlan, Dan Keniston, Supreet Kaur, Chris Udry, Tavneet Suri and from seminar participants at the University of Michigan, the Minnesota Population Center, the University of Minnesota, Yale University, NEUDC and CSAE. Data collection for this project was supported by grants from the Michigan Institute for Teaching and Research in Economics (MITRE), Population Studies Center, Center for Education of Women and Rackham Graduate School at the University of Michigan. Chyn also acknowledges support from a NICHD training grant to the Population Studies Center at the University of Michigan (T32 HD0077339). Chyn and Kerwin are both grateful for use of services and facilities at the Population Studies Center which is funded by a NICHD Center Grant (R24 HD041028).

I. Introduction

Social scientists and policymakers have long-standing interest in understanding how peers shape an individual's behavior. A key question is whether peers affect productivity in the workplace. The answer to this question is important for determining the optimal allocation of labor and designing firm incentives.

While an emerging literature provides compelling evidence that peer effects exist in workplace settings, the mechanisms behind these peer effects are less clear.[1] Several studies show worker effort is sensitive to the social pressure that arises in settings where there are externalities from effort due to joint production and team compensation (Mas and Moretti, 2009; Gould and Winter, 2009; Bandiera et al., 2013; Babock et al., 2015; Cornelissen et al. 2017; Arcidiacono et al., 2017).[2] Fewer studies test whether peer effects on productivity may also arise from mechanisms such as motivation or norms – channels not directly controlled by firms (Kaur et al., 2010).

This paper provides new evidence on the mechanisms that drive workplace peer effects by conducting a unique field experiment with an agricultural firm. We collaborated with a tea estate in Malawi and randomly allocated about 1,000 piece-rate workers to locations on tea fields. Each worker is assigned a specific plot area to pick tea leaves each day, and our design created exogenous, within-worker variation in the composition of plot neighbors. We focus on estimating the effect of the average of peers' ability (permanent productivity) on workers' output.

Importantly, several aspects of this setting allow us to test for the existence of social influences on worker's performance that are unrelated to spillovers in the production process or in

[1] See Herbst and Mas (2015) for additional discussion and a list of previous studies on peer effects in the workplace.
[2] Another well-studied channel for workplace peer effects is knowledge spillovers (i.e., learning). Notable studies testing for this form of peer effect include Waldinger (2012), Azoulay et al. (2010), Jackson and Bruegmann (2009), and Guryan et al. (2009).

the compensation scheme. Unlike much of the previous work examining peer effects, workers in our setting are paid piece rates, and there is no cooperation in the process of collecting tea. Hence, any impact of peers on productivity in our setting is likely to be due to learning or psychological mechanisms related to "motivation" (e.g., self-control or norms).[3]

We find that the average ability (i.e. permanent productivity) of coworkers affects a worker's own daily volume of tea collected. Specifically, increasing the average ability of nearby coworkers by 10 percent raises a tea worker's productivity by about 0.3 percent (p=0.028). Studying very different contexts, Mas and Moretti (2009) and Falk and Ichino (2006) find effects that are about three times as large.

In additional analysis, we test for heterogeneity in peer effects that varies with individual characteristics. In our sample, we find that the impact of mean peer ability is largest for women, while the effect for men is smaller and not statistically significant. This pattern of results is notable because it suggests that there is potential for aggregate productivity gains from resorting workers to ensure that men are near women.[4] This is because men in our sample tend to have higher permanent productivity for tea leaf plucking.

Notably, our results contrast with previous studies that estimate the impact of working with friends (Bandiera et al.. 2010; Park, 2017). We measured social connections in our sample and exploit the fact that our randomization scheme ensured that workers sometimes worked on a plot where a friend was immediately nearby. We find small and statistically insignificant impacts of working near a friend on a worker's own productivity. Moreover, the ability of a worker's friends

[3] In our setting, it is also possible that workers may synchronize their productivity and effort in an attempt to socialize with other workers nearby. We discuss and provide evidence against this hypothesis in greater detail in our analysis.
[4] Studies of education contexts find evidence of heterogeneous peer effects that imply there may be gains to resorting students (Sacerdote, 2001; Carrell et al., 2009, Carrell et al., 2013, Booji et al., 2017).

has no influence on his or her output. Instead, we find significant and positive impacts of working near higher ability *non-friends*.

To learn more about the mechanisms driving the peer effects in our setting, we conducted an incentivized choice experiment to measure demand for working next to specific peers. Specifically, we surveyed a subset of employees in the season following our first experiment and found that 71 percent of respondents stated they would prefer having a high-productivity co-worker nearby if re-assignment were possible. When asked for the reason for their choices in an open-ended question, 79 percent workers stated that higher-productivity peers provide motivation.

Moreover, we also find that workers are willing to pay for faster peers. We give workers the option to be re-assigned next to a high-productivity co-worker if they are willing to give up bars of soap that that they received for taking the survey. Among workers who stated a preference for having a fast peer nearby, 71 percent were willing to give up one bar and 55 percent were willing to give up two bars of soap – equivalent to 9 and 18 percent of average daily earnings respectively. Our choice experiment also allows workers to pay to work next to any other worker of their choice, but almost none of them do so. This indicates that workers place value specifically on having high-ability coworkers, rather than on other traits correlated with ability.

This paper's main contribution is to provide evidence that workplace peer effects can be driven by a particular class of psychological mechanisms. Our setting rules out effects driven by externalities in the production process or the terms of the financial contract. Supplementary analysis also suggests that learning does not drive the estimates in our study because there no evidence that peer effects vary by experience of workers. Further, the choice experiment that we conduct shows that workers in our sample are willing to pay to work near high-productivity peers, a pattern that is inconsistent with standard models of rank preferences, last-place aversion, shame or reputational concerns (Tincani, 2015; Kuziemko et al., 2014; Kandel and Lazear, 1992; Breza and Chandrasekhar, 2015). Rather, our results are consistent with models of contagious enthusiasm or limited self-control (Mas and Moretti, 2009; Kaur et al., 2014).

Finally, our study also sheds light on solutions to several methodological issues concerning the estimation of peer effects models. We make three main contributions in this regard. First, we demonstrate the usefulness of utilizing explicit random assignment of peers in our workplace setting. This is necessary because we find that workers choose to work near co-workers with similar levels of ability when given the opportunity. Second, our randomization scheme allows us to eliminate a bias inherent common to many peer effect settings. As noted by Guryan et al. (2009), Angrist (2014), and Caeyers and Fafchamps (2016), there is a mechanical negative correlation between a worker's own ability and their peers' ability. This correlation exists even if there is random assignment because a worker cannot be assigned to be her own peer. To address this issue, we randomly assigned workers to a different set of peers for each day in a work cycle. This feature of our design allows us to eliminate any correlation between own and peer ability by estimating models with individual fixed effects. Third, we provide guidance on how to estimate ability when pre-intervention measures are not available. As in Mas and Moretti (2009), we measure ability as estimated permanent productivity using data from the same period as our intervention. We build on their approach to estimating permanent productivity by using a novel leave-one-out estimator of own ability that eliminates spatially-correlated productivity shocks that would otherwise cause upward bias estimates peer effects.

II. Background

To conduct our study, we partner with Lujeri Tea Estates, a large agricultural firm in Malawi. Our sample is a group of roughly 1,000 employees who hand-pick ("pluck") leaves from tea bushes (hereafter, we refer to these workers as pluckers). Workers temporarily store plucked leaves in baskets and empty their baskets at a central weighing station. There is no explicit cooperation involved in this process, and pay is a constant piece rate for each kilogram of plucked tea.

Production at the firm is organized by assigning workers to “gangs” which are each managed by a supervisor. The size of a gang is typically around 45 pluckers, but the sizes range from 29 on the low end to 76 on the high end. Each gang is responsible for plucking tea from a pre-determined set of fields over the course of a harvesting “cycle” (7 to 12 calendar days). In our analysis sample, there are 78 fields for the 22 gangs we study.

On each tea field for a gang, the supervisor assigns workers to pluck tea from a specific set of plots (between 1 to 3 per day depending on the characteristics of the field and day of the week).[5] Each field has between 30 and 120 plots, and workers must pluck on their assigned plots before moving on to other plots.[6,7] At the completion of a harvesting cycle, the gang returns back to the initial field for a new round of plucking – unlike other crops that are harvested once or a few times, tea bushes grow continuously throughout the season.

Figure 1 illustrates the general assignment of workers to plots on a given field and the rotation of workers throughout the harvesting cycle.[8] Panel A shows that each worker is assigned two contingent plots (blue squares). The example highlights three workers who are colored red, green and yellow. The illustration shows that workers B and C are the immediate plot neighbors of worker A. Panel B provides an illustration showing how workers change assignments across fields covered during a 6-working-day harvesting cycle. On each day of the harvesting cycle a given worker has an assigned set of plots for that day’s specific field. Across days in the harvesting

[5] Plots vary somewhat in size and shape but are typically about 25 meters across.
[6] Pluckers are expected to finish their assigned plots in accordance with the cycle schedule. Pluckers who finish their plots can ask for additional plots that are not assigned to other workers. If a worker is absent on a given day, the supervisor will assign the plot to other workers for the day.
[7] Fixed plot assignment is done so that workers internalize the negative effects of over- and under-plucking bushes on their plots.
[8] Note that in reality the fields and plots are often not evenly-sized rectangles.

cycle, a worker will have different neighbors. In the example, the three hypothetical workers are sometimes separated as shown for Cycle Days 3, 5 and 6. On these days, the workers will have different plot neighbors.

III. Random Assignment of Workplace Peers

We designed our experimental intervention to randomly assign workers to plots on tea fields in order to generate exogenous variation in exposure to workplace peers. To implement this, we obtained the roster of workers in each gang and a "plucking program" for each gang. The plucking program is a predetermined list of which field (or fields) a gang works during each day of its cycle and the number of pluckers that should be assigned to each field. In the simplest case, there is one field on each cycle day with all the pluckers working on it.[9] We use this information to generate randomly ordered lists of pluckers for each day of a gang's harvesting cycle. On cycle days where a gang works on multiple fields, we also randomly determine which workers are on each field.

We used these randomized lists to determine the order in which pluckers were assigned to plots on each field. The random assignment took advantage of the usual assignment process in which pluckers stand in a queue and receive plot assignments in the order in which they were standing. The supervisor makes the assignments by "snaking" back and forth across the field and taking the next plucker from the queue for each plot. Our random assignment scheme altered this system by giving the supervisors a randomly-ordered list to use in this snake pattern.[10] Each gang supervisor was responsible for assigning workers using the randomly generated list of worker assignments in February 2015. We verified compliance with these assignments by having our

[9] Many gangs have more complicated schedules, spending multiple cycle days on some fields, and splitting the gang across more than one field on certain days of their cycle.
[10] An exception to our randomization is the first work day ("Cycle Day 1") in a gang's cycle. We intentionally did not randomize work assignments in this data. On these days, supervisors assigned workers using the usual method, in which the plots are still assigned using the snaking pattern across the field, but the order of the pluckers comes from the order in which they stand in the queue. We use this non-random assignment on the first work day to test for endogenous assignment and sorting of workers to locations on a field in Section 5.

project managers visit each gang in the week after randomization. As a result of our intervention, workers are randomly assigned to plots within a field for different cycle days as illustrated in Panel B of Figure 1.

IV. Data

To study the impact of workplace peers, we use three main sources of data. First, we rely on administrative data from the firm on worker productivity. Productivity is defined as kilograms of tea plucked per day and is electronically recorded by the firm for the purpose of paying employees. As a result, it is measured with minimal error. This data on worker productivity is available from December 2014 to April 2015 (the beginning and end of the main tea harvest season, respectively). Second, we hired project staff to record information on the plot neighbors assigned to each worker as a result of the randomized assignment that we implemented. Third, we collected survey data to obtain measures of worker characteristics such as background demographics and social networks.

III.A Main Analysis Sample

Our study centers on 1,046 pluckers who worked during the main season after we implemented our randomized work assignments in February 2015. Table 1 provides summary statistics based on the survey and administrative data. The average age for workers is about 37 years and about 43 percent of the sample is female. Only 7 percent of workers are new (with zero previous experience at the firm) and average experience is nearly 8 years. Over the course of our study period, the average daily output for each worker is 69 kilograms of plucked tea leaves and workers have on average about 5 assigned neighbors on any given day of work.

Our study focuses on studying how working alongside peers of different ability affects daily output. As detailed in Section 5, we measure a worker's ability by estimating their permanent productivity. Table 1 shows that the average ability estimate for workers in our sample is 62.19

kilograms. To provide a sense of the "treatment", Table 1 also shows the mean of nearby co-worker ability on each day. Across workers and days in our sample, the standard deviation of peer ability is nearly 13 kilograms.

V. Empirical Strategy

The main question in this paper is whether working in close proximity to higher-ability co-workers increases productivity in our sample of tea pluckers. To address this question, we estimate the following linear model of peer effects for the productivity of worker *i*:

$$y_{ift} = \mu_i + \beta\overline{Ability}_{-ift} + \delta_{tf} + \epsilon_{ift} \tag{1}$$

where y_{ift} is the (logged) total kilograms of tea plucked on field *f* and date *t*. The key variable in Equation 1 is $\overline{Ability}_{-ift}$ which is the mean of ability of all co-workers who are assigned to work adjacent to the plots that worker *i* is assigned. The model also includes date-by-field fixed effects, δ_{tf}, to control for variation in harvest conditions over the course of the season and across the tea estate. Finally, we also control for time-invariant determinants of productivity – such as the worker's own plucking ability – by including individual-level fixed effects μ_i. We cluster all standard errors at the level of the treatment, which is the combination of workers and cycle days. We also cluster by the combination of field and date to account for correlated shocks that might affect entire fields. Because we estimate the treatment variable, we correct for the sampling error in the ability measure using the Bayesian parametric bootstrap technique from Mas and Moretti (2009).

To measure ability in our sample of tea pluckers, we rely on an approach pioneered by Mas and Moretti (2009) which uses estimates of worker fixed effects as a measure of ability ("permanent productivity").[11] Specifically, we use the plucking data and estimate:

$$y_{ift} = \mu_i + \boldsymbol{M_{ift}}\gamma' + \delta_{tf} + \tau_{ift} \tag{2}$$

where the term $\boldsymbol{M_{ift}}$ is a vector of dummy variables which indicate whether worker j is working next to worker i in field f on date t.[12] The idea is that the vector γ contains a set of parameters that absorb any possible peer effects and allows us to obtain unbiased estimates of worker fixed effects μ_i, under the assumption that each individual worker can have any effect on his or her coworkers.[13] For Equation 1, we use these estimates to define $\overline{Ability}_{-ift} = \bar{\hat{\mu}}_{-ift}$ as our measure of peer influence.[14]

The resulting ability distribution is shown in Panel A of Figure 2. Ability is appears to be approximately log-normally distributed. A Kolmogorov-Smirnov test fails to reject the null hypothesis that the log of ability is normally distributed. This is consistent with the kernel density of log ability (Panel B). Appendix Table A1 shows that the ability measure is correlated with known determinants of productivity in our sample. Ability is positively correlated with experience, and these effects are highly nonlinear. Women are less productive than otherwise-similar men because physical strength is an important determinant of tea-plucking efficiency. The more weight you can carry the less often you have to bring your tea to the weighing station, and the higher the share of work time you can spend on plucking leaves.

[11] Bandiera et al. (2010) and Park (2017) use similar approaches to estimating ability as permanent productivity.
[12] To be clear, the set of possible co-workers is based on the gang for worker *i* so that $\boldsymbol{M_{ift}}$ is a vector of J_i -1 dummy variables. Here, J_i is the total number of pluckers in *i*'s gang.
[13] One additional assumption for identification is that the form of any coworker peer effects is additively separable across workers.
[14] Since $\overline{Ability}_{-ift}$ is based on estimated quantities, the correct standard errors for Equation (1) need to be adjusted for this additional source of sampling variability.

In models of peer effects such as Equation 1, there are three main concerns for identification. First, the key assumption for identification of β is that there is no correlation between the average ability of one's peers and the unobserved determinants of individual productivity: $cov(\overline{Ability}_{-ift}, \epsilon_{ift}) = 0$. One way this assumption could be violated is if supervisors assign workers with higher ability to work on particularly productive areas of a field that are physically close together. Our intervention eliminates this possibility by randomly assigning workers to plots within a field and makes it possible to purge estimates β of any endogenous sorting effects.

Table 2 shows that this random assignment is key for producing causal estimates of peer effects by presenting results from a series of regressions of worker' own ability on the mean of their co-workers ability.[15] Column (1) shows that the there is a slight positive correlation between own ability and peer ability on the sample of plucking days that correspond to "Cycle Day 1" of each gang's work cycle. These are days that we explicitly did not randomize workers and which gang supervisors implemented plot assignments through the status quo system. In line with our random assignment intervention, the results in Columns (3) and (4) show that this correlation does not exist for the remainder of the sample, which supports the identifying assumption in our linear-in-means model.

A second threat to identification in Equation 1 is the fact that a worker cannot be assigned to be her own neighbor. As noted in Guryan et al. (2009) and Angrist (2014), there is a mechanical negative correlation between a worker's own ability and that of her neighbors. Consider a worker who is at the top of the ability distribution. Her neighbors will necessarily be lower ability than

[15] Note that we follow the recommendation of Guryan et al. (2009) and include the leave-one-out mean in our test of random assignment. The inclusion of this term corrects for exclusion bias in tests for random assignment, but only eliminates the bias in the case of non-overlapping peer groups (Caeyers and Fafchamps 2016).

her, and likewise for a worker at the bottom of the distribution. Caeyers and Fafchamps (2016) call this phenomenon "exclusion bias": since the worker's ability appears in the error term of the regression, there is a mechanical negative correlation between peer ability and the error term and hence coefficient estimates are downward-biased. Unlike classical measurement error, this bias can push estimates through zero and into negative values. Since we cannot perfectly measure worker ability, even adding our estimated ability measure as a control will leave some component of ability in the error term, and estimated peer effects will be negatively biased. The small negative correlations between own ability and peer ability in Column 4 of Table 2 are consistent with the existence of the exclusion bias problem.

Our research design allows us to address exclusion bias in a straightforward way. Specifically, the within-worker random assignment means that workers face different peers throughout the course of a work cycle. This allows us to implement a simple solution to address exclusion bias: we include individual fixed effects μ_i in our regression model. These worker fixed effects break any potential correlation between the fixed component of the error term and the ability of a worker's peers.

Third, spatial correlations in output within a field can generate correlations between the output of coworkers and individuals that bias estimates of $\overline{Ability}_{-ift}$ from Equation 2. For example, suppose that one corner of a specific field has higher productivity – maybe due to better sun exposure or an uneven distribution of fertilizer. This type of spatial correlation between plots will raise the output of all the workers located in that corner on each day, and also increase their estimated ability. We address this issue by estimating Equation 2 using a double leave-one-out process that is similar to a jackknife estimator. In addition to the standard approach of leaving the worker herself out of the calculation of the peer-group mean, we also exclude all data from the same field when estimating her peers' ability levels. Thus process estimates ability at the field level by restricting the sample for Equation 2 to all *other* fields for a specific gang. This implies that, for example, the ability estimates we use for Field 5 for Gang 3 leave out Field 5. As a result, we always estimate Equation 1 using a measure of mean peer ability that *excludes* both data for the same date for which we observe output and any other data for the same field. This procedure ensures that spatial correlation in plot quality, or spatially correlated shocks, do not cause violations of the assumption that $cov(\overline{Ability}_{-ift}, \epsilon_{ift})$ = 0.

Finally, we are also interested in testing whether peer effects vary with a worker's characteristics. To explore this, we augment Equation 1 by interacting $\overline{Ability}_{-ift}$ with dummy variables for characteristics such as sex or a worker's age. In addition, we also create a series of dummies for an individual's own ability quartile and interact these with $\overline{Ability}_{-ift}$.[16] Previous research has used this type of specification and found evidence of notable heterogeneity in peer effects across the distribution of student ability (Hoxby and Weingarth, 2005; Carrell et al., 2009; Imberman et al., 2012; Carrell et al., 2013; Booji et al. 2017) and worker ability (Mas and Moretti, 2009; Cornelissen et al. 2017).

VI. Main Results

To test whether the average ability of co-workers affects productivity, Table 3 reports estimates from Equation 1. Column (1) shows that there is a positive and significant effect of the mean ability of peers on worker productivity. Specifically, a 10 percent increase in mean ability of peers

---

[16] Specifically, we estimate the following more general model of peer effects:

$$y_{ift} = \mu_i + \sum_{q=1}^{q=4} \theta_q \, D_i^q \times \overline{Ability}_{-ift} + \delta_t + \lambda_f + \epsilon_{ift}$$

where the terms D_i^q are indicators which equal one if a person is in the *q* quartile of the distribution of worker ability.

is associated with a 0.3 percent increase in the daily amount of kilograms of tea plucked for each worker. Column (2) shows that our estimates are essentially unchanged when we condition on date-by-location fixed effects. Relative to the literature, these estimates are about one third of the size of estimates produced in a lab setting by Falk and Inchino (2006) and studying supermarket cashiers by Mas and Moretti (2009).

Selective attendance at work in response to changes in peer quality does not drive these results. We show this by creating a panel of observations for all days over the course of the season and creating an indicator for whether not a worker was at work and plucking tea.[17] Appendix Table 1 provides results from estimating Equation 1 where the dependent variable is attendance (Column 1) and plucking tea (Column 2). The point estimates are not significant and very small in magnitude, which suggests there is no impact of peers on work attendance. This lack of effects on attendance is consistent with the idea that the incentive to attend work is very strong, because the firm fires workers for irregular attendance; the overall attendance rate is roughly 87 percent.

The double-leave-one-out estimator matters substantively for our results, suggesting that correlated shocks would otherwise cause upward bias in our estimates of peer effects. Appendix Table 3 presents the results of estimating Equation 1 without making the double-leave-one-out correction. That is, the ability estimates in that table are constructed using data that does include the same field that is used to measure output. Similar to our main results, the estimates are positive and significant. Notably, these estimates are between 43 percent and 89 percent larger in magnitude depending on the specification used. This suggests that without our correction, spatially correlated productivity shocks would cause us to overestimate the magnitude of the peer effects in this context.

[17] To be clear, for our main productivity analysis we use the subset of observations where a worker was present at work and also plucked tea (as opposed to being assigned to other tasks for that day).

Finally, we also test for the existence of peer effects that are vary across workers with different individual characteristics. Table 4 shows treatment effect heterogeneity by gender, age, and workers' own ability. We see no evidence of heterogeneity in peer effects by workers' ability levels. There is some evidence that younger workers experience larger peer effects, but these differences are not statistically significant.

In contrast with the results for other characteristics, there are stark differences in the magnitudes of the peer effects experienced by men and women. Women's output rises by 0.6 percent for every 10 percent increase in coworker ability – an effect twice as large as what we see for the overall sample. This effect is strongly statistically significant (p=0.007). Men, on the other hand, experience essentially zero peer effects. The across-gender difference in the magnitudes of the peer effects is significant at the 10 percent level. In addition to differing in the magnitudes of the peer effects they experience, men and women differ in terms of their estimated ability level. Figure 3 shows kernel densities of worker ability by gender. The male distribution is shifted to the right relative to the female distribution. Appendix Table 3 shows summary statistics for ability by gender and shows that men have an underlying productivity level that is 8.4 kilograms of tea higher than women.

The heterogeneity in peer effects and ability levels by gender is important because it allows for the possibility of raising aggregate productivity by rearranging workers. If peer effects were constant across individuals, then re-assigning a high-ability peer from one group to another would have equal and offsetting effects. Because men do not experience peer effects in our sample, we can raise the productivity of low-ability female workers by placing them next to high-ability men without affecting men's productivity. Moreover, because men tend to be more productive than women in this context, creating matches between high-ability men and low-ability women does not necessitate creating an equal number of matches between low-ability men and high-ability women.

Thus, our results suggest there is the potential for creating more output gains than losses. The gender differences in ability suggest simple rules for optimizing output that could be implemented based just on easily-observed characteristics. Since men have a higher distribution of ability than women, an alternating male-female pattern would lead to higher output than mixing workers purely randomly or placing workers next to colleagues that share their gender.

VII. Mechanisms:

The evidence presented thus far shows that mean co-worker ability has an impact on productivity. A range of mechanisms could generate positive peer effects in general, but our setting allows us to rule out two of these immediately. First, unlike in many previously studied settings, externalities in the production process are not present in our setting since there is no cooperation and no need for workers to coordinate. Second, the compensation scheme does not generate peer effects because workers receive piece-rates. With this in mind, this section proceeds to consider three other types of mechanisms that could be driving our estimates of peer effects.

VII. A. Socialization

One leading mechanism for workplace peer effects is socialization between workers. In setting similar to ours, Bandiera et al., (2010) studied workers who picked fruit at a large agricultural firm in the UK and estimated the impact of working physically near a friend. Their analysis suggests that socialization between friends affects worker productivity. When slow fruit pickers worked near friends who were typically fast, they work harder to catch up. Similarly, relatively fast pickers slow down for their slower friends. Further evidence on the impact of friends also comes from Park (2017). He studied workers at a seafood processing plant and found that a worker's productivity drops by six percent when working near a friend.

Using data on social networks, Table 7 provides evidence that suggests socialization and interactions between friends do not drive peer effects in our sample. Specifically, we use self-reported friendship between pluckers (measured at baseline) to identify when workers are plucking on plots near their friends. We then compute the average ability of nearby co-workers who are friends. Similarly, we calculate the average ability of nearby co-workers who are not friends. On the average day in our sample, a worker has around three plot neighbors that are friends. We use these two separate measures of average co-worker ability in our basic linear-in-means specification (Equation 1) and report the results in Column (3). The results show that a 10 percent increase in the mean ability of non-friends increases worker productivity by 0.28 percent (p=0.028), which is nearly identical to impact that we obtain from our main specification in Table 3. In contrast to these effects for non-friends, the point estimate on the effect of increasing ability of friends is much smaller and not statistically significant.[18]

VII. B. Learning

Another potential mechanism to explain our findings is learning ("knowledge spillovers"). It is conceivable that plot neighbors learn from observing each other work, thereby generating the positive effects that we observe.[19] To explore this possibility, we perform two tests. First, we examine whether peer effects in our setting are heterogeneous with respect to workers' past experience. Under the learning hypothesis, we would expect the effects of average peer ability

---

[18] While these point estimates appear quite different, we cannot reject a test of the null hypothesis that the effects of non-friends and friends are equal (p-value = 0.24).

[19] Among previous studies testing for the existence of knowledge spillovers, Jackson and Bruegmann (2009) find evidence of knowledge spillovers among teachers while Waldinger (2012) finds no evidence among university scientists.

would be largest for workers who have no prior or relatively less experience. Second, we test whether lagged measures of peer ability appear have any affect a worker's current productivity. If workers learn from their co-workers, we believe that any lagged measure of co-worker ability should have detectable effects.

Table 6 provides evidence against the idea that learning drives our peer effect results by reporting estimates from augmented versions of Equation 1 in which we add measures of worker experience. The results in Column (1) replicate the estimate from our baseline specification for the sample of workers for whom we have self-reported experience. Column (2) builds on our main specification by adding an interaction between a dummy indicating status as a new worker (no prior experience) and our measure of peer ability. Notably, the estimate for this estimate for this interaction is not significant and implies that the effect of higher-ability peers is actually *negative* for new workers. As an alternative test for heterogeneity in effects by experience level, we create dummies based on the quartiles of worker experience observed in our sample. We interact these dummies with our measure of average peer ability and present the results for these terms in Column (3). Although the results for this specification are not precise, the point estimates for the least experience and most experienced workers are remarkably similar. Overall, the results in Table 6 provide no evidence that workers with less experience benefit more from working near higher ability co-workers.

We also find that results from models that include lagged measures of co-worker ability do not suggest there is any learning between co-workers. Table 7 reports estimates from an augmented version of Equation 1 which includes measures of co-worker ability measured one cycle day ago ("t-1"), two cycle days ago ("t-2") and three cycle days ago ("t-3"). Our preferred specification is a model which includes current and all lagged measures and is reported in Column (3). These results show that current peers have a positive impact on productivity while there is no detectable impact of any lagged measure.

VII. C. Psychological Mechanisms: Motivation vs. Shame

A final prominent possibility that we consider is that the peer effects we measure might operate through psychological channels. Based on Kandel and Lazear (1992), Equation 3 shows a stylized utility as a function of effort (denoted as e)::

$$u(e, \theta) = \begin{cases} w(e) - c(e) & if\ \theta = \theta^L \quad (Low - Ability\ Peer) \\ w(e) - c(e) + p(e) & if\ \theta = \theta^H \quad (High - Ability\ Peers) \end{cases}$$

where the functions $w(.)$, $c(.)$and $p(.)$ are the wage, cost and "peer pressure" functions, respectively. Assume the wage function increases with worker productivity, which is a function of effort. In the case that an individual has low ability peers, workers choose an optimal effort level e^* based on setting the marginal cost of effort equal to marginal payoff in wages. When an individual has fast peers, there is an additional peer pressure term in the utility function. The existence of positive peer effects implies the peer pressure function has a positive first derivative (i.e., $\frac{\partial p}{\partial e} > 0$) so that workers choose higher effort level when surrounded by high-ability peers.

A common characterization of psychological peer effects ascribes them to shame or last-place aversion. In this case, we can imagine that the function $p(.)$ is always negative: having high-ability peers lowers a worker's utility. In this case, workers will increase effort as a way of minimizing the utility loss. Alternatively, there are various other psychological mechanisms that would suggest that the function $p(.)$ is always positive. These channels could include motivation or "contagious enthusiasm" effects.[20] When $p(.)$ is always positive, high-ability peers *raise* a worker's utility level, and effort increases to maximize this benefit.

[20] They are also akin to the benefits that runners, cyclists, and other athletes receive from pacing against other competitors.

Both classes of psychological mechanisms $p(.)$ have identical predictions for the impact of peers on effort, but distinct predictions for welfare. Effort always increases regardless of whether high-ability peers cause shame or contagious enthusiasm because this choice only depends on the slope of $p(.)$. Yet, the existence of shame-based peer effects imply that workers are worse off if they have high ability peers. This type of psychological mechanism could make attempts to optimize output and profits through the use of peer effects unsustainable: workers would tend to quit or demand higher wages, undermining any potential gains. In contrast, if peer pressure is a good thing, then rearranging workers to exploit peer effects would have the side effect of making them happier as well, making it a more-sustainable strategy.

A key prediction of the model of motivation as an explanation for peer effects is that exposure to faster coworkers is beneficial. Workers should therefore be willing to pay for higher-ability peers. To test this prediction, we conducted a supplementary survey and incentivized choice experiment for a subset of tea workers during the harvest season following the one used in our main experiment (2015-2016). In this experiment, we asked workers whether they wanted higher-ability peers, and whether they would be willing to give up part of the compensation that they received for taking part in the survey (workers were each given two bars of soap as a token of thanks for taking the survey). Workers were informed that one worker per gang would have their choices implemented for real (a chance of about 1 in 50).

Panel A of Table 8 reports that 71 percent of workers would like to be assigned next to a fast (high ability) peer in their gang. Further, Panel B shows that these workers seeking re-assignment are *willing to pay* for these peers: 71 percent of workers who want a fast peer would be willing to give up one bar of soap while 55 percent would be willing to give up two bars of soap. When asked for the main reason for their choices in an open-ended question, 83 percent workers state that faster peers provide motivation. Only 15 percent state learning as a reason for wanting higher-ability peers.

We also find some evidence that the demand for peers is correlated with the magnitude of peer effects. Appendix Table 5 provides results from a specification where we interact mean co-worker ability with four indicator variables based on an individual's willingness to pay for peers. Workers who want faster peers but are not willing to pay experience the largest peer effects. Peer effects are close to zero for workers who are willing to pay soap for better peers. This gap could be driven by differences in income or liquidity across groups, which would also affect measured WTP. There are two reasons for caution in interpreting these results, however. First, we record the WTP measure in the season after the peer effects took place and so reverse causality is possible. Second, this sample includes only workers who were in the follow-up experiment and present in our sample for the main experiment (N=442). This latter fact is problematic because there is some evidence of differential attrition: workers who did not want faster peers are 29 percent of the overall incentivized choice sample, but just 18 percent of the sample for this heterogeneity analysis.

Overall, the results from our willingness to pay experiment strongly suggest that motivation is the key driver of the peer effects we measure in this study. They rule out a range of other potential mechanisms posited in the literature, such as shame or a desire to avoid being last (Kandel and Lazear, 1992; Kuziemko, 2014). Since workers are willing to pay for faster peers, shame-type mechanisms can only be the operative mechanism inasmuch as it serves as a commitment device, inducing workers to reach a higher level of effort that they truly would like to achieve.

VIII. Conclusion

This paper provides evidence on workplace peer effects by conducting an experimental intervention with an agricultural firm in Malawi from February 2015 to April 2015. We randomly assigned tea pluckers to plot assignments on fields and use this variation in peer composition to examine the effect of mean coworker ability (permanent productivity) on workers' output.

Using administrative data on daily productivity, we find that the average of coworker ability has a positive and significant effect: increasing the average of coworker ability by 10 percent increases own-productivity by about 0.3 percent. Furthermore, supplementary analysis suggests that these peer effects vary based on a worker's characteristics. Specifically, we find that the mean of peer ability has larger effects for women in our sample. This finding is notable because it implies that re-sorting workers on gender could generate gains in aggregate productivity. This is possible because we find that the average male in our sample has higher productivity than the average female.

To shed light on the mechanisms driving our peer effect estimates, we conducted a choice experiment in the next harvesting season that allowed workers to choose new coworkers as plot neighbors. In this experiment, we find that 71 percent of workers wanted to be assigned to a fast (high-ability) coworker. Moreover, workers were willing to pay for faster coworkers: 55 percent of workers were willing to give up two bars of soap (worth 18 percent of daily wages) that we had given them as a gift for survey participation. In open-ended follow-up questions, 83 percent of workers state that working near faster peers motivates them.

Overall, our analysis provide strong evidence that workplace peer effects in our setting stem from the effect that co-workers have on motivation. In additional analysis, we do not find evidence that learning or worker socialization drive our results. Finally, the fact that workers receive piece rates and there is no cooperation in tea plucking rules out that the effects in our setting stem from production externalities or incentives of the firm’s compensation scheme.

References:

Arcidiacono, Peter, Josh Kinsler, and Joseph Price. 2017. "Productivity Spillovers in Team Production: Evidence from Professional Basketball." *Journal of Labor Economics* 35 (1):191–225.

Angrist, J. D. (2014). The Perils of Peer Effects. *Labour Economics*, 30, 98–108.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2010. "Social Incentives in the Workplace." *The Review of Economic Studies* 77 (2): 417–58.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2013. "Team Incentives: Evidence from a Firm Level Experiment." *Journal of the European Economic Association* 11 (5): 1079–1114.

Booij, Adam S., Edwin Leuven, and Hessel Oosterbeek. 2017. "Ability Peer Effects in University: Evidence from a Randomized Experiment." The Review of Economic Studies 84 (2):547–78.

Caeyers, Bet, and Fafchamps, Marcel. 2016). Exclusion Bias in the Estimation of Peer Effects (Working Paper No. No. 22565). NBER.

Carrell, Scott E., Richard L. Fullerton, and James E. West. 2009. "Does Your Cohort Matter? Measuring Peer Effects in College Achievement." *Journal of Labor Economics* 27 (3): 439–64.

Carrell, Scott E., Bruce I. Sacerdote, and James E. West. 2013. "From Natural Variation to Optimal Policy? The Importance of Endogenous Peer Group Formation." *Econometrica* 81 (3): 855–82.

Cornelissen, Thomas, Christian Dustmann, and Uta Schönberg. 2017. "Peer Effects in the Workplace." American Economic Review 107 (2):425–56.

Falk, Armin, and Andrea Ichino. 2006. "Clean Evidence on Peer Effects." *Journal of Labor Economics* 24 (1): 39–57.

Gould, Eric D, and Eyal Winter. 2009. "Interactions between Workers and the Technology of Production: Evidence from Professional Baseball." *Review of Economics and Statistics* 91 (1): 188–200.

Guryan, Jonathan, Kory Kroft, and Matthew J Notowidigdo. 2009. "Peer Effects in the Workplace: Evidence from Random Groupings in Professional Golf Tournaments." *American Economic Journal: Applied Economics* 1 (4): 34–68.

Herbst, Daniel, and Alexandre Mas. 2015. "Peer Effects on Worker Output in the Laboratory Generalize to the Field." *Science (New York, N.Y.)* 350 (6260): 545–49.

Hoxby, Caroline M., and Gretchen Weingarth. 2005. “Taking Race out of the Equation: School Reassignment and the Structure of Peer Effects.” *Unpublished Manuscript.*

Imberman, Scott A., Adriana D. Kugler, and Bruce I. Sacerdote. 2012. “Katrina’s Children: Evidence on the Structure of Peer Effects from Hurricane Evacuees.” *American Economic Review* 102 (5): 2048–82.

Jackson, C. Kirabo, and Elias Bruegmann. 2009. “Teaching Students and Teaching Each Other: The Importance of Peer Learning for Teachers.” *American Economic Journal: Applied Economics* 1 (4): 85–108.

Kandel, Eugene, and Edward P. Lazear. 1992. “Peer Pressure and Partnerships.” *Journal of Political Economy*, 801–817.

Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan. 2010. “Self-Control and the Development of Work Arrangements.” *American Economic Review* 100 (2): 624–28.

Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan. 2014. “Self Control at Work.” *Journal of Political Economy*.

Kuziemko, I., R. W. Buell, T. Reich, and M. I. Norton. 2014. “Last-Place Aversion: Evidence and Redistributive Implications.” *The Quarterly Journal of Economics* 129 (1): 105–49.

Mas, Alexandre, and Enrico Moretti. 2009. “Peers at Work.” *American Economic Review* 99 (1): 112–45.

Park, Sangyoon. 2017. “Socializing at Work: Evidence from a Field Experiment with Manufacturing Workers.” *Unpublished Manuscript.*

Sacerdote, Bruce. 2001. “Peer Effects with Random Assignment: Results for Dartmouth Roommates.” *The Quarterly Journal of Economics* 116 (2): 681–704.

Waldinger, Fabian. 2012. “Peer Effects in Science: Evidence from the Dismissal of Scientists in Nazi Germany.” *The Review of Economic Studies* 79 (2): 838–61.

Figure 1. Tea Worker Field Assignment Illustrations

Panel A. Hypothetical Assignment for Three Tea Workers

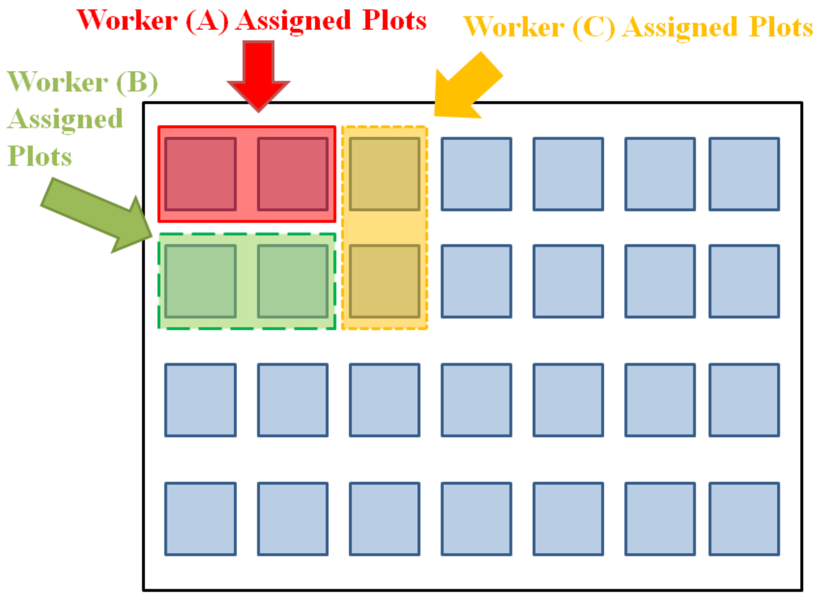

Panel B. Plot Assignments Change Over Days in Harvesting Cycle

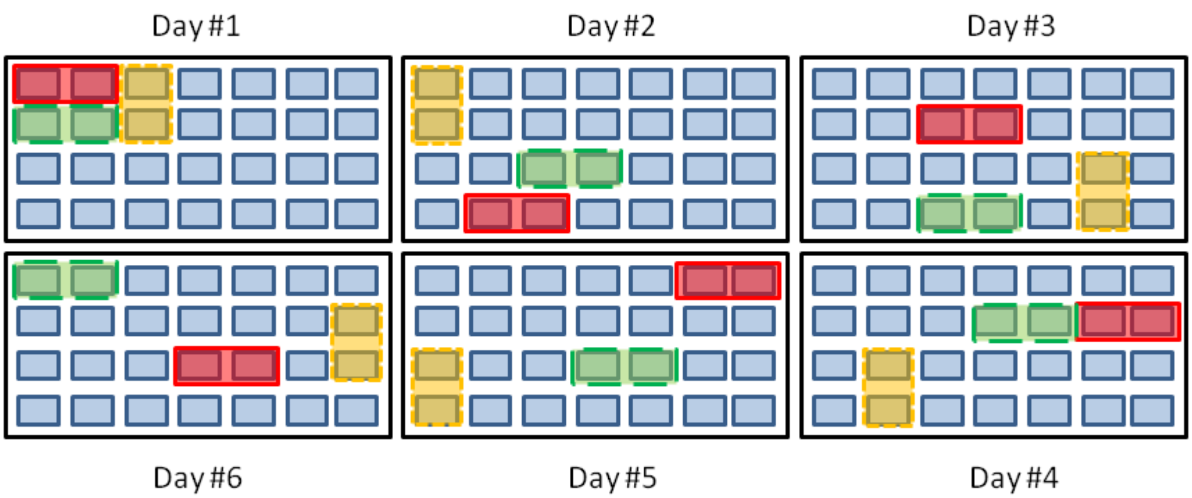

Notes: The two panels illustrate work assignments for tea workers at the Lujeri Tea Estates in Malawi. Panel A shows how three workers would be assigned two plots each. For our analysis, all workers A, B and C would be neighboring co-workers. Panel B shows how plot work assignments change over the course of a harvest cycle that lasts 6 calendar days and visits distinctly different fields. On some days and fields, workers A, B and C are neighbors. Yet, there are also cases where they are not neighbors: for example, on Day #3, #5 and #6, workers A, B and C are not assigned to work in neighboring plots.

Figure 2: Distribution of Worker Ability

Panel A. Kernel Density of Ability

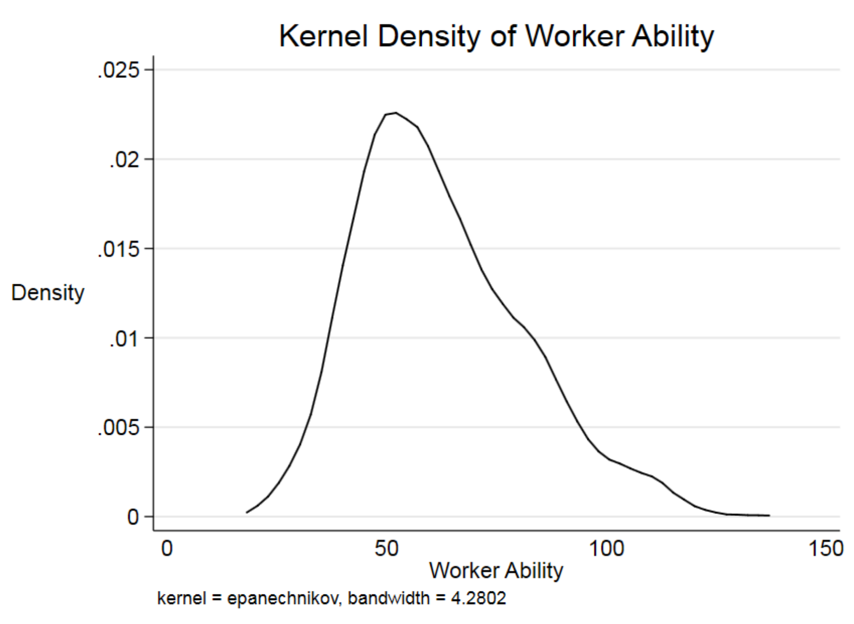

Panel B. Kernel Density of log(Ability)

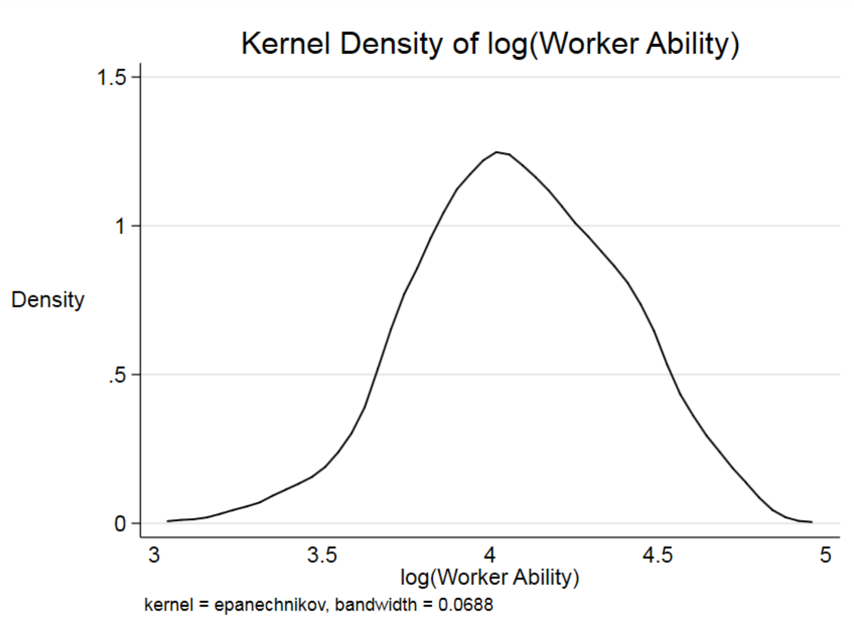

Figure 3: Distribution of Worker Ability by Gender

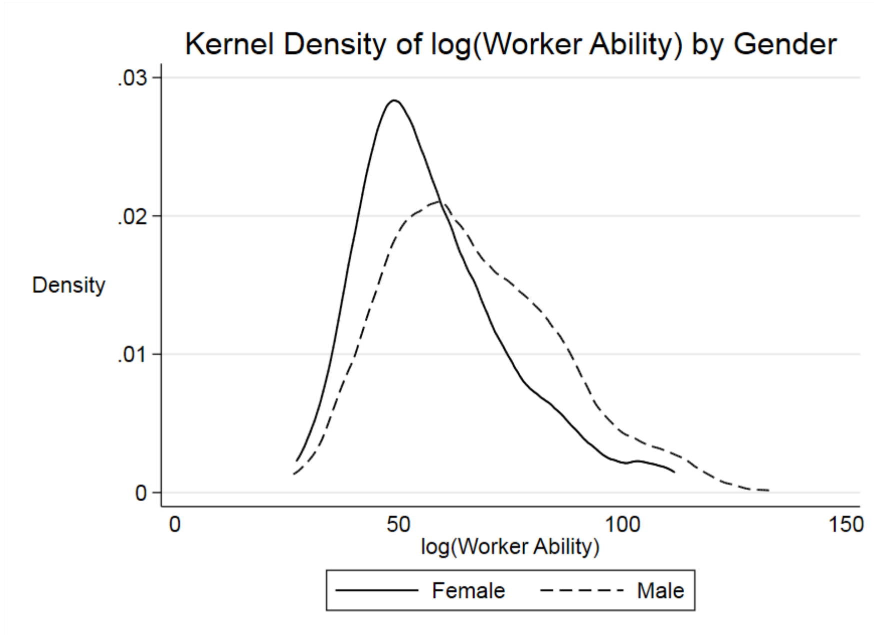

Table 1. Summary Statistics, Lujeri Worker Sample

	(1) Average	(2) Std. Deviation	(3) 10th Percentile	(4) 90th Percentile	(5) Obs (N)
Age	37.43	10.64	25.00	52.00	944
Female (=1)	0.43	0.50	0.00	1.00	944
Married (=1)	0.63	0.48	0.00	1.00	944
New Worker (=1)	0.07	0.26	0.00	0.00	944
Experience (Yrs.)	7.72	8.31	0.08	15.50	944
Ability (Estimate)	62.19	18.93	40.83	88.48	999
# Neighbors	4.69	1.82	2.00	8.00	35,460
Mean Peer Ability	61.44	12.92	47.21	79.46	35,644
Output (kgs.)	69.21	36.11	27.00	118.00	38,034

Notes: This table presents descriptive statistics based on survey data we collected for a sample of tea pluckers at the Lujeri Tea Estates in Malawi.

Table 2. Testing for Random Assignment

	Dependent Variable: Log(Own Ability)			
	(1)	(2)	(3)	(4)
Log(Mean Peer Ability)	0.062	-0.020	-0.039	-0.046
	(0.074)	(0.030)	(0.033)	(0.028)
Log(Leave-One-Out Gang Mean Ability)	0.860***	0.945***	0.963***	-8.922***
	(0.089)	(0.037)	(0.041)	(0.505)
Cycle Day 1	Yes	Yes	No	No
Remaining Cycle Days	No	Yes	Yes	Yes
Worker Fixed Effects	No	No	No	No
Date by Location Fixed Effects	No	No	No	Yes
Observations	9,313	44,858	35,545	35,449
Adjusted R-squared	0.246	0.233	0.230	0.397

Notes: This table presents results from a regression of our measure of a worker's own ability on the mean ability of physically nearby co-workers. The underlying data is a panel at the worker and day level. The results in Columns (1) are from the sample of "Cycle Day 1" days which did not have random assignment of workers to plot assignments at the tea estate. Column (2) presents results using the full sample of all dates and cycle dates in our data. Columns (3) and (4) use the sample of all non Cycle 1 days -- this is the sample for which we randomly assigned workers to locations on fields. Standard errors are two-way clustered at the level of a worker-by-cycle-day and date-by-field and adjusted using the Bayesian parametric bootstrap of Mas and Moretti (2009).

Table 3. Effects of Workplace Peers, Linear Model

	Dependent Variable: Log of Daily Output	
	(1)	(2)
Log(Mean Peer Ability)	0.030**	0.028**
	(0.013)	(0.013)
Worker Fixed Effects	Yes	Yes
Date Fixed Effects	Yes	No
Location (Field) Fixed Effects	Yes	No
Date by Location Fixed Effects	No	Yes
Observations	35,545	35,641
Adjusted R-squared	0.715	0.396

Notes: This table presents results from a regression of daily output (kilograms of tea plucked) on the mean ability of physically nearby co-workers. The underlying data is a panel at the worker and day level. The results in Columns (1) and (2) use two different approaches to control for date and location effects. TThe measure of daily output comes from administrative data obtained from Lujeri Tea Estates; information on neighbors was recorded and collected staff for this project Standard errors are two-way clustered at the level of a worker-by-cycle-day and date-by-field and adjusted using the Bayesian parametric bootstrap of Mas and Moretti (2009).

Table 4. Heterogeneous Peer Effects

	Dependent Variable: Log of Daily Output		
	(1)	(2)	(3)
Male X Log(Mean Peer Ability)	0.006		
	(0.019)		
Female X Log(Mean Peer Ability)	0.060***		
	(0.022)		
Quartiles of Age			
[Age 20 to 29] X [Log(Mean Peer Ability)]		0.057**	
		(0.029)	
[Age 30 to 35] X [Log(Mean Peer Ability)]		0.020	
		(0.029)	
[Age 36 to 44] X [Log(Mean Peer Ability)]		0.013	
		(0.028)	
[Age 44 to 72] X [Log(Mean Peer Ability)]		0.035	
		(0.032)	
Quartiles of Own Ability			
[Own Ability Quartile 1] X [Log(Mean Peer Ability)]			0.029
			(0.026)
[Own Ability Quartile 2] X [Log(Mean Peer Ability)]			0.022
			(0.025)
[Own Ability Quartile 3] X [Log(Mean Peer Ability)]			0.042
			(0.030)
[Own Ability Quartile 4] X [Log(Mean Peer Ability)]			0.030
			(0.025)
Worker Fixed Effects	Yes	Yes	Yes
Date by Location Fixed Effects	Yes	Yes	Yes
Observations	33,010	33,010	35,545
Adjusted R-squared	0.725	0.725	0.715

Notes: This table presents results from a regression of daily output (kilograms of tea plucked) on the mean ability of physically nearby co-workers, interacted with worker characteristics. The underlying data is a panel at the worker and day level. The measure of daily output comes from administrative data obtained from Lujeri Tea Estates; information on neighbors was recorded and collected by staff for this project. Standard errors are two-way clustered at the level of a worker-by-cycle-day and date-by-field and adjusted using the Bayesian parametric bootstrap of Mas and Moretti (2009).

Table 5. Effects of Friends and Non-Friends in the Workplace

	Dependent Variable: Log of Daily Output		
	(1)	(2)	(3)
Log(Mean Peer Ability), Non-Friends	0.028**		0.028**
	(0.013)		(0.013)
Any Non-friends (=1)	-0.068		-0.072
	(0.061)		(0.061)
Log(Mean Peer Ability), Friends		0.006	0.006
		(0.013)	(0.013)
Any Friends (=1)		-0.033	-0.035
		(0.055)	(0.055)
Worker Fixed Effects	Yes	Yes	Yes
Date by Location Fixed Effects	Yes	Yes	Yes
Observations	35,583	35,583	35,583
Adjusted R-squared	0.715	0.715	0.715

Notes: This table presents results from a regression of daily output (kilograms of tea plucked) on measures of the mean ability of nearby co-workers who are friends and non-friends. The underlying data is a panel at the worker and day level. The measure of daily output comes from administrative data obtained from Lujeri Tea Estates; information on neighbors was recorded and collected by staff for this project. Standard errors are two-way clustered at the level of a worker-by-cycle-day and date-by-field and adjusted using the Bayesian parametric bootstrap of Mas and Moretti (2009).

Table 6. Peer Effects by Experience Level

	Dependent Variable: Log of Daily Output		
	(1)	(2)	(3)
Log(Mean Peer Ability)	0.030**	0.032**	
	(0.013)	(0.015)	
New Worker (=1) X Log(Mean Peer Ability)		-0.022	
		(0.068)	
Quartile 1 Exp. X Log(Mean Peer Ability)			0.046
			(0.031)
Quartile 2 Exp. X Log(Mean Peer Ability)			0.018
			(0.030)
Quartile 3 Exp. X Log(Mean Peer Ability)			0.028
			(0.029)
Quartile 4 Exp. X Log(Mean Peer Ability)			0.029
			(0.029)
Worker Fixed Effects	Yes	Yes	Yes
Date by Location Fixed Effects	Yes	Yes	Yes
Observations	35,545	33,010	33,010
Adjusted R-squared	0.715	0.725	0.725

Notes: This table presents results from a regression of daily output (kilograms of tea plucked) on the mean ability of physically nearby co-workers, broken down by workers' experience at the firm. The underlying data is a panel at the worker and day level. The results in Column (2) and (3) are from specifications that include additional interactions based on the worker's self-reported experience at Lujeri. The measure of daily output comes from administrative data obtained from Lujeri Tea Estates; information on neighbors was recorded and collected staff for this project. Standard errors are two-way clustered at the level of a worker-by-cycle-day and date-by-field and adjusted using the Bayesian parametric bootstrap of Mas and Moretti (2009).

Table 7. Effects of Previous Days' Peers

	Dependent Variable: Log of Daily Output					
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Mean Peer Ability)	0.036***	0.041***	0.035**			
	(0.014)	(0.014)	(0.015)			
Log(Mean Peer Ability), t-1	0.024*	0.020	0.009	0.021	0.016	0.005
	(0.014)	(0.015)	(0.016)	(0.014)	(0.015)	(0.016)
Log(Mean Peer Ability), t-2		-0.020	-0.025		-0.025	-0.029*
		(0.016)	(0.017)		(0.016)	(0.017)
Log(Mean Peer Ability), t-3			-0.016			-0.022
			(0.017)			(0.017)
Worker Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Date by Location Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	33,120	30,221	27,505	33,146	30,245	27,528
Adjusted R-squared	0.717	0.715	0.709	0.717	0.715	0.709

Notes: This table presents results from a regression of daily output (kilograms of tea plucked) on the mean ability of physically nearby co-workers, including lagged as well as contemporaneous peer ability. The underlying data is a panel at the worker and day level. The measure of daily output comes from administrative data obtained from Lujeri Tea Estates; information on neighbors was recorded and collected staff for this project. Standard errors are two-way clustered at the level of a worker-by-cycle-day and date-by-field and adjusted using the Bayesian parametric bootstrap of Mas and Moretti (2009).

Table 8. Willingness to Pay for Fast Peers

	(1) Pct.	(2) Obs.
Panel A. All Survey Respondents		
Who do you want to be re-assigned next to?		
A fast plucker in your gang	0.71	724
A slow plucker in your gang	0.05	724
Any person of your choosing	0.11	724
No-reassignment	0.14	724
Panel B. Respondents who want to be next to fast pluckers		
If you could switch to be near a fast plucker…		
…would you be willing to give up 1 bar of soap?	0.71	515
…would you be willing to give up 2 bar of soap?	0.55	515

Notes: This table presents statistics from survey data that we collected for tea pluckers at the Lujeri Tea Estates. The full sample for our survey is 620 individuals and we performed the choice experiment (in Panel B) with a subset of respondents due to logistical and administrative costs. For the choice experiment, respondents were given a gift of two bars of soap (18 percent of average daily wages) and asked if they would be willing to give up soap in exchange for being re-assigned.

Appendix Table 1. Regression of Worker Ability on Worker Attributes

	Dependent Variable: Worker Ability	
	(1)	(2)
Female	-5.676***	-6.088***
	(2.046)	(2.034)
Married	2.127	2.214
	(2.124)	(2.112)
Household Size	0.852	0.701
	(0.549)	(0.553)
Household Spending per Capita	0.000	0.000
	(0.000)	(0.000)
Age	-0.015	
	(0.067)	
Quartiles of Age		
[Age 30 to 35] X [Log(Mean Peer Ability)]		2.113
		(1.727)
[Age 36 to 44] X [Log(Mean Peer Ability)]		2.049
		(1.827)
[Age 44 to 72] X [Log(Mean Peer Ability)]		1.766
		(1.896)
Experience	0.272***	
	(0.086)	
Quartiles of Experience		
[2.1 to 5 Years] X [Log(Mean Peer Ability)]		5.694***
		(1.698)
[5.1 to 10.7 Years] X [Log(Mean Peer Ability)]		6.258***
		(1.716)
[10.8 to 49.5 Years] X [Log(Mean Peer Ability)]		7.742***
		(1.846)
Worker Fixed Effects	Yes	Yes
Date by Location Fixed Effects	Yes	Yes
Observations	909	909
Adjusted R-squared	0.062	0.075

Notes: This table presents results from a regression of workers' ability levels, as measured in predicted kilograms of tea plucked per day, on various exogenous covariates. The underlying data is a cross-section at the worker level. Standard errors are heteroskedasticity-robust.

Appendix Table 2. Effects of Assigned Peers on Attendance and Tea Plucking

	Dependent Variable: Attendance	*Dependent Variable: Tea Plucking*
	(1)	(2)
Log(Mean Peer Ability)	0.006	-0.000
	(0.011)	(0.012)
Worker Fixed Effects	Yes	Yes
Date Fixed Effects	Yes	Yes
Observations	48,000	48,000
Adjusted R-squared	0.106	0.188

Notes: This table presents results from a regressions of an indicator for the worker being present at work (column 1) or being engaged in tea plucking (column 2) on the mean ability of the physically nearby co-workers for their assigned field for the day. The underlying data is a panel at the worker and day level. The regressions control for date instead of date-by-location because workers do not have an assigned location unless their gang is engaged in plucking. The measure of daily attendance and plucking comes from administrative data obtained from Lujeri Tea Estates; information on neighbors was recorded and collected staff for this project. Standard errors are two-way clustered at the level of a worker-by-cycle-day and date-by-field and adjusted using the Bayesian parametric bootstrap of Mas and Moretti (2009).

Appendix Table 3. Effects of Workplace Peers without Double Leave-One-Out Correction

	Dependent Variable: Log of Daily Output	
	(1)	(2)
Log(Mean Peer Ability)	0.043***	0.053***
	(0.014)	(0.014)
Worker Fixed Effects	Yes	Yes
Date Fixed Effects	Yes	No
Location (Field) Fixed Effects	Yes	No
Date by Location Fixed Effects	No	Yes
Observations	35,545	35,641
Adjusted R-squared	0.715	0.396

Notes: This table presents results from a regression of daily output (kilograms of tea plucked) on the mean ability of physically nearby co-workers. The underlying data is a panel at the worker and day level. The results in Columns (1) and (2) use two different approaches to control for date and location effects. TThe measure of daily output comes from administrative data obtained from Lujeri Tea Estates; information on neighbors was recorded and collected staff for this project Standard errors are two-way clustered at the level of a worker-by-cycle-day and date-by-field and adjusted using the Bayesian parametric bootstrap of Mas and Moretti (2009).

Appendix Table 3. Summary Statistics for Worker Ability by Gender

	(1) Average	(2) Std. Deviation	(3) 10th Percentile	(4) 25th Percentile	(4) 50th Percentile	(4) 75th Percentile	(4) 90th Percentile	(5) Obs (N)
Overall	62.83	18.75	41.94	49.00	59.29	73.98	88.98	909
Females	58.07	16.87	40.25	45.76	54.26	67.21	81.93	393
Males	66.46	19.31	43.51	52.04	63.94	79.65	92.58	516

Notes: This table presents descriptive statistics about worker ability for the subset of workers who have gender information from our survey data (909 of the overall total of 999 workers in our sample). The ability measure is estimated using Equation 2.

Appendix Table 4. Peer Effects vs WTP for Fast Peers

	Dependent Variable: Log of Daily Output	
	(1)	(2)
[Does Not Want Fast Peers] X [Log(Mean Peer Ability)]	0.061	0.061
	(0.048)	(0.048)
[Wants Fast Peers, Any Price] X [Log(Mean Peer Ability)]	0.018	
	(0.021)	
[Wants Fast Peers for 0 Bars] X [Log(Mean Peer Ability)]		0.109***
		(0.040)
[Wants Fast Peers for 1 Bar] X [Log(Mean Peer Ability)]		-0.010
		(0.047)
[Wants Fast Peers for 2 Bar] X [Log(Mean Peer Ability)]		-0.030
		(0.033)
Worker Fixed Effects	Yes	Yes
Date by Location Fixed Effects	Yes	Yes
Observations	17,517	17,517
Adjusted R-squared	0.722	0.721

Notes: This table presents results from a regression of daily output (kilograms of tea plucked) on the mean ability of physically nearby co-workers interacted with dummies for the worker's willingness to pay for faster peers. The underlying data is a panel at the worker and day level. The measure of daily output comes from administrative data obtained from Lujeri Tea Estates; information on neighbors was recorded and collected staff for this project. Standard errors are two-way clustered at the level of a worker-by-cycle-day and date-by-field and adjusted using the Bayesian parametric bootstrap of Mas and Moretti (2009).