Peers and Motivation at Work:

Evidence from a Firm Experiment in Malawi*

Lasse Brune
Eric Chyn
Jason Kerwin

May 7, 2020

Abstract

This paper studies workplace peer effects by randomly varying work assignments at a tea estate in Malawi. We find that increasing mean peer ability by 10 percent raises productivity by 0.3 percent. This effect is driven by the responses of women. Neither production nor compensation externalities cause the effect because workers receive piece rates and do not work in teams. Additional analyses provide no support for learning or socialization as mechanisms. Instead, peer effects appear to operate through "motivation": given the choice to be reassigned, most workers prefer working near high-ability co-workers because these peers motivate them to work harder.

Keywords: Peer effects, firm productivity, field experiment.
JEL: J24, J33, M11, M54.

*Lasse Brune is a postdoctoral fellow at the Global Policy Research Lab of the Buffett Institute for Global Studies at Northwestern University. Eric Chyn is an assistant professor of economics at Dartmouth College and a faculty research fellow at NBER (email: eric.t.chyn@dartmouth.edu). Jason Kerwin is an assistant professor of applied economics at the University of Minnesota and an affiliated professor at J-PAL. The authors declare that we have no relevant or material financial interests related to the research described in this paper. We would like to thank three anonymous referees for their detailed comments and suggestions. We are grateful for feedback from Martha Bailey, Charlie Brown, John DiNardo, Brian Jacob and Jeff Smith. We also received insightful comments from Emily Breza, Esther Duflo, Dean Karlan, Supreet Kaur, Dan Keniston, Erzo F.P. Luttmer, Sangyoon Park, Joseph Ritter, Bryan Stuart, Tavneet Suri, Chris Udry, and seminar participants at the University of Michigan, the Minnesota Population Center, the University of Minnesota, Yale University, ABCA, CSAE, IUSSP, MIEDC, NEUDC, and SOLE. Data collection for this project was supported by grants from the Michigan Institute for Teaching and Research in Economics (MITRE), the Population Studies Center, the Center for Education of Women, and the Rackham Graduate School at the University of Michigan. Chyn acknowledges support from a NICHD training grant to the Population Studies Center at the University of Michigan (T32 HD0077339). Chyn and Kerwin are grateful for the use of services and facilities at the Population Studies Center, which is funded by a NICHD Center Grant (R24 HD041028). We registered this study at the AEA RCT Registry as AEARCTR-0001105. The data used in this article are available online: Brune, Lasse, Eric Chyn, and Jason Kerwin. 2020. Replication Data for: Peers and Motivation at Work: Evidence from a Firm Experiment in Malawi. Harvard Dataverse. doi:10.7910/DVN/CAXO8Y. All appendices can be found in the online appendix to the paper, available at http://jhr.uwpress.org/. All errors and omissions are our own.

1 Introduction

A large literature provides compelling evidence that a worker's own performance depends on her peers and social interactions (Herbst and Mas 2015). Several studies show that worker effort is sensitive to the social pressure that arises when there are externalities from effort due to joint production and team compensation (Mas and Moretti 2009; Gould and Winter 2009; Bandiera, Barankay, and Rasul 2013; Babcock et al. 2015; Arcidiacono, Kinsler, and Price 2016; Battisti 2017; Cornelissen, Dustmann, and Schonberg 2017; Jiang 2020; Amodio and Martinez-Carrasco 2018; Silver 2019). For example, Mas and Moretti (2009) find that retail workers in teams appear to engage in monitoring and free-riding behavior that affects productivity.[1] Another well-studied channel for workplace peer effects is knowledge spillovers (i.e., learning). Jackson and Bruegmann (2009) and Azoulay, Zivin, and Wang (2010) find evidence of learning among teachers and medical researchers, respectively.[2]

Yet few studies provide evidence on workplace peer effects when jobs do not directly incentivize them through production externalities or team incentives. For example, some psychological mechanisms such as motivation or norms may drive peer effects even when workers do not work in teams or receive joint compensation.[3] Gneezy and Rustichini (2004) provide evidence consistent with this type of peer effect by showing that a child runs faster when running alongside a peer than when running alone. Falk and Ichino (2006) study a laboratory experiment and find that students fill envelopes faster when they share a room with a peer.[4] Bandiera, Barankay, and Rasul (2010) use a natural experiment to study

[1] Mas and Moretti (2009) study the productivity of cashiers in a national supermarket chain. As they note, this is an environment that is characterized by group production and prone to free-riding. This is because customers are not committed to a single aisle: if one cashier is working slowly, other cashiers will have a greater workload.

[2] De Grip and Sauermann (2012) also find that training programs have spillover effects on co-workers. Their finding is consistent with the existence of knowledge spillovers. At the same time, Waldinger (2012) and Guryan, Kroft, and Notowidigdo (2009) find little evidence of learning-based peer effects.

[3] Battaglini, Benabou, and Tirole (2005) develop a theory of self-control and peer effects to explain the impact of self-help groups and role models.

[4] Similarly, Kaur, Kremer, and Mullainathan (2010) use pilot data from an experiment to show that piece-rate data entry workers increase their own productivity if they sit near a peer with above-average productivity.

the impact of working with friends, and they find that workers increase or decrease their productivity to match the output of their social ties.

This paper provides new evidence on the role of psychological peer effects by conducting a unique field experiment at 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 day the firm assigns specific plots for workers to pick tea leaves, and our design created exogenous, within-worker variation in the composition of nearby co-workers. We focus on estimating the effect of the average of peer ability (i.e., permanent productivity) on the worker's own output.

Several features of this setting allow our analysis to provide a relatively clean examination of the mechanisms that drive workplace peer effects in a real-world context. First, workers in our setting are paid piece rates, and there is no cooperation in the process of collecting tea. The implication of this is that any impact of peers is not due to shirking or the pressure that arises when workers attempt to counteract free-riding problems in joint production (Kandel and Lazear 1992). This distinguishes our work from prior studies that focus on contexts where peer pressure is a key attribute (e.g., Mas and Moretti 2009). Second, we use detailed data to conduct a wide range of supplementary analyses to quantify the extent to which peer effects are due to socialization or learning between workers.

Our main finding is that the average ability of co-workers affects a worker's own daily volume of tea collected. Specifically, increasing the average ability of nearby co-workers by 10 percent raises a tea worker's productivity by about 0.3 percent (p-value=0.028).[5] Notably, these estimates are much smaller than the results obtained in previous studies where joint production induces peer effects. For example, Mas and Moretti (2009) study retail workers who received fixed wages and engaged in group production, finding peer effects that are about five times as large. We can reject that our estimates are equal to the effects detected in their work at the one percent significance level.

[5] As an additional interpretation, a one standard deviation increase in mean peer ability would increase a worker's productivity by 0.6 percent.

The analysis also reveals heterogeneous responses to peer ability in our sample. We find that peer effects are large for women, while there is a small and statistically insignificant estimate for men. This pattern differs from Gneezy and Rustichini (2004), who found that peers do not improve the performance of young girls running a race, but do affect the performance of boys. At the same time, our results are more in line with Hahn et al. (2017), Lavy, Silva, and Weinhardt (2012) and Stinebrickner and Stinebrickner (2006) who provide evidence that girls distinctly benefit from peers in educational settings. More generally, our finding on heterogeneity is notable because it suggests that there is potential for productivity gains from re-sorting workers to ensure that men are near women.[6] This is because men in our sample have higher average permanent productivity for tea leaf plucking, and peer effects are driven by peer productivity rather than by gender directly.

We also show that our results contrast with previous studies estimating the impact of working with friends (Bandiera, Barankay, and Rasul 2010; Park 2019). We measure social connections in our sample and exploit the fact that our randomization scheme ensured that workers sometimes worked adjacent to friends. 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 has no influence on his or her output. Instead, we find significant and positive impacts of working near higher-ability non-friends. One potential reason that friends have different effects in our setting than in other studies is that tea workers are often not close enough to communicate with their peers while working. That said, the workers in our setting are sufficiently close to see adjacent co-workers and observe their productivity.

Additional analyses provide several pieces of evidence against the idea that learning drives our results. We find no evidence that peer effects vary by the experience level of workers. Furthermore, we find no evidence that lagged measures of peer ability have an impact on a worker's current productivity.[7] A simple model of knowledge spillovers would suggest that

[6]Studies in education contexts find evidence of heterogeneous peer effects that imply there may be gains to re-sorting students (Sacerdote 2001; Carrell, Fullerton, and West 2009; Carrell, Sacerdote, and West 2013; Booij, Leuven, and Oosterbeek 2017).

[7]This analysis exploits the fact that the experimental design insured that workers have new peers on

lagged measures of co-worker ability should affect current productivity. The lack of evidence of learning in our sample is consistent with the idea that effort and inherent physical ability are the main determinants of productivity in our setting.

Overall, both the setting that we study and the pattern of results suggest that psychological mechanisms drive the peer effects that we detect. Supplementary survey data provide additional evidence in support of this argument. During the following agricultural season, we asked workers about their preferences for working next to specific peers in an incentivized choice exercise.[8] In this sample, we find that 72 percent of respondents preferred having a high-productivity co-worker nearby if reassignment were possible. When asked for the reason for their choices in an open-ended question, 74 percent of workers said higher-productivity peers provide motivation.

This paper's main contribution is to provide experimental estimates of workplace peer effects in an environment that has negligible production complementarities and high incentives for productivity. As discussed in Herbst and Mas (2015), teams that feature joint compensation are susceptible to shirking due to free-riding problems. In these settings, spillover estimates may be negative due to the shirking or positive because workers engage in monitoring and peer pressure to counteract shirking. We study a setting that removes both forces. Our results can be considered as a step forward in terms of understanding workplace peer effect mechanisms. In addition, our study provides relatively clearer evidence on the types of psychological mechanisms driving the effects that we detect. The fact that workers view high ability peers as a source of motivation is consistent with models of contagious enthusiasm and limited self-control (Battaglini, Benabou, and Tirole 2005; Mas and Moretti 2009; Kaur, Kremer, and Mullainathan 2010, 2015). Our analysis contrasts with models of rank preferences, including last-place aversion, shame or reputational concerns (Kuziemko et al. 2014; Breza and Chandrasekhar 2015; Tincani 2015).[9] This distinction is important since

different days within a harvest cycle.

[8]Workers were informed that there was a one in ten chance of being selected to have one choice implemented.

[9]Similarly, DellaVigna, List, and Malmendier (2012) study charitable giving and provide evidence that

these latter mechanisms imply that workers will have a lower level of utility when exposed to high-performing peers. Mechanisms of this type suggest that workers may resist policies that seek to use peer effects to enhance firm output.

Finally, this paper provides solutions to methodological issues concerning the estimation of peer effects models. We make three contributions in this regard. First, we demonstrate explicitly the value of random assignment of peers in our workplace setting. Random assignment is necessary because we find that workers choose to work near peers with similar ability levels when given the opportunity. Second, our within-worker randomization scheme allows us to eliminate a bias common to many peer effect settings. As noted by Guryan, Kroft, and Notowidigdo (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 design allows us to eliminate any correlation between own and peer ability by estimating models with worker 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 double leave-one-out estimator that eliminates spatially-correlated productivity shocks that would otherwise bias estimates of peer effects.

2 Background

To conduct our study, we partnered 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

individuals may respond to disutility associated with social pressure.

plucked leaves in baskets and empty their baskets at a central weighing station.[10] There is no explicit cooperation involved in this process, and pay is a constant piece rate for each kilogram of plucked tea.[11]

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 and 3 per harvest cycle day, depending on the characteristics of the field).[12,13] The assignment of workers to plots for given field is done at the beginning of the main season and generally remains in place throughout the season. Each field has between 30 and 120 plots. Workers are expected to complete full (i.e., eight-hour) workdays, and workers who finish plucking their assigned plots early are sent to additional plots to pluck those as well.[14,15] At the completion of a harvesting cycle (most commonly 6 work days, or

[10]The locations of weighing stations are fixed throughout the season for a number of logistical reasons such as coordination with the trucks that pick up plucked leaves. The weighing stations are also often located under trees to provide shade and to hang the scale. There is usually one weighing station per field.

[11]Lujeri pays workers their earnings every two weeks.

[12]Supervisors have a number of responsibilities in addition to handling work assignments. These responsibilities include monitoring that bushes are plucked to the right height (to avoid over-plucking), managing leaf quality inspection, coordinating water, tea and food allocation, preparing weighing stations, coordination of weighing station clerks, and working with tractor drivers. Supervisors also must request additional temporary workers from the head office if there are work absences.

[13]The modal number of assigned plots per plucker in our data is two: workers are assigned one plot on 19.5 percent of plucker-days, two plots on 53 percent of plucker-days, and three plots on 27.5 percent of plucker-days. The data for our sample do not allow us to observe the actual number of plots a plucker works on any given day.

[14]If a plucker finishes their assigned plots for a given day, there are two ways that pluckers are assigned additional plots on that day. One is that they are sent to work on additional plots that are not assigned to any worker; these plots exist because fields are typically not evenly divisible by the total number of workers in a gang. Another is that when workers are absent on a given day, their plots are given to other workers after those workers are finished with their own plots. While worker absences are rare (i.e., the attendance rate is 87 percent in our sample), the size of both the gang and field imply that there are many additional plots for workers to use if they finish early. This limits the scope for crowd-out among the set of workers who finish early.

[15]Fixed plot assignment is done so that workers internalize the negative effects of over-plucking bushes on their plots. Specifically, the concern is that over-plucking could reduce the future productivity of a plot. Note supervisors conduct plot inspections to further minimize the risk of over-plucking.

7 calendar days, since Sunday is a day off), the gang returns 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 typical pattern of assignment of workers to plots on a given field and the rotation of workers throughout the harvesting cycle.[16,17] Panel A shows 28 hypothetical square plots. The example highlights three workers who are each assigned two contiguous plots (each assignment is indicated by the larger rectangles with solid, dashed, and thin-dashed lines, respectively). 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 the fields covered during a six-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 cycle, a worker will have different neighbors. In the example, the three hypothetical workers are separated at times, as shown for cycle days 3, 5 and 6 (where the larger rectangles representing contiguous assignments for workers are all separated). On these days, the workers will have different plot neighbors.

In this workplace context that features neither joint compensation nor teamwork in production, peer effects can still occur because workers observe co-worker productivity in two main ways. First, the plots are approximately 25 meters from one edge to another, implying that workers are about 25 meters from their closest peer on average. This means that a worker is close enough to see a peer's performance and speed (ability) by observing their movement through the field; tea plants do not block lines of sight since the plants are pruned to (roughly) waist height.[18] Second, workers regularly travel to the weighing station to drop off tea. In this case, seeing your neighbors go to the weighing station provides an

[16]While square plots are the most common shape, in reality the fields and plots are often not evenly-sized. In some cases, workers may share more or less of a plot boundary depending on the field geography. Because we do not have precise measures of plot boundaries, we are unable to test whether peer effects vary based on the amount of a plot boundary that is shared.

[17]To provide a better sense of the size and shape of the plots, Appendix Figure A1 shows a photograph of a tea field at Lujeri Tea Estates.

[18]This distance between workers also implies that workers are often not close enough for communication to be easy.

easily-observable measure of peer productivity.

3 Experimental Design

We designed our experimental intervention to randomly assign workers to plots on tea fields 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 on 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.[19] We used this information to generate randomly-ordered lists of pluckers for each day of a gang's harvesting cycle.[20] On cycle days where a gang works on multiple fields, we also randomly determined which workers were on each field.

We used these randomized lists to determine the order in which supervisors assign pluckers to plots on each field. The random assignment took advantage of the usual assignment process, wherein pluckers stand in a queue and receive plot assignments in the order in which they are 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.[21] Each gang supervisor assigned workers using the randomly generated list of worker assignments in February 2015. We verified compliance with these assignments by having our project managers visit each gang in the week after randomization. In addition,

[19]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.

[20]We implemented the random assignments by collecting lists of the members of all tea-plucking gangs in five divisions at the tea estate. No demographic or other restrictions were applied in determining who was included in the sample.

[21]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 on this work day. On cycle day 1, 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 choose to stand in the queue, giving them some degree of control over who their co-workers are. In Section 6, we use this non-random assignment on the first work day to test for endogenous sorting of workers to locations on a field.

project staff confirmed compliance with the assignment via random spot checks several weeks after the initial assignment. As a result of our intervention, workers were assigned randomly to plots within a field for different cycle days as illustrated in Panel B of Figure 1.

4 Data

To study the impact of workplace peers, we use three main sources of data.[22] 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. These data on worker productivity are available from the beginning of the season in December 2014 to end of the main tea harvest season in April 2015. 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 baseline social networks. For the social network data, we asked respondents to identify friends, and there was no maximum on the number of friends that could be listed.[23]

4.1 Main Analysis Sample

Our study centers on 999 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.[24] 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 per worker is 69 kilograms of

[22]All data and code used for the analyses in this paper are available via the Harvard Dataverse: https://doi.org/10.7910/DVN/CAXO8Y (Brune, Chyn, and Kerwin 2020).

[23]See Appendix E for details of the social network data collection.

[24]Due to survey non-response, we are missing demographic information for 5 percent of the sample (55 workers).

plucked tea leaves. At the average output per day, the daily wage implied by the piece-rate of MK 19.32/kg is equivalent to \$7.73 in PPP terms at the time of the study. 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.

5 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 model of peer effects for the productivity of worker i:

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

where y_{ift} is the (logged) total kilograms of tea plucked on field field f and date t. The key variable in Equation 1 is $\overline{Ability}_{-i-f,t}$, which is the mean of ability of all co-workers who are assigned to work adjacent to the plots that worker i is assigned.[25,26,27] The model

[25]The subscript $-f$ indicates that a measure excludes data from field f. As detailed further below, we estimate $\overline{Ability}_{-i-f,t}$ using data excluding field f to address concerns over spatial spillovers.

[26]We use the mean of peer ability in Equation 1 since this measure is standard within the literature. In addition, this metric appears to be the best measure based on our analysis of peer effects in this context. When we use the maximum of peer ability as the measure of peer influence, we find statistically significant and positive impacts that are somewhat smaller than estimates based on Equation 1. We find statistically insignificant and positive estimates when we use the minimum of peer ability as the measure.

[27]An alternative model of peer effects is discussed in Silver (2019). This work innovates relative to the literature on workplace peer effects by focusing on group match effects that capture the influence of working with a particular co-worker group. As an exploratory exercise, we have implemented an approach based on Silver (2019) and found there are important worker-by-peer group match effects in our setting. Specifically, workers are 24.3 percent more productive when working in a two-standard deviation faster peer group environment.

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.[28,29] 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 at the worker-by-cycle-day level.[30] We also cluster by the combination of field and date to account for correlated shocks that might affect entire fields.[31] 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).[32]

To measure ability for each tea plucker in our sample, we rely on an approach pioneered by Mas and Moretti (2009), which uses estimates of worker fixed effects as a measure of ability (i.e., permanent productivity).[33] Specifically, we use the plucking data and estimate:

$$y_{igt} = \mu_{i-f} + \mathbf{M}_{igt}\gamma' + \delta_{tg} + \tau_{igt} \tag{2}$$

where the term $\mathbf{M}_{igt}$ is a vector of dummy variables which indicate whether worker j is working next to worker i in field g on date t.[34] 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

[28] We randomized plot assignment within the combination of a field and a cycle day. Since date-by-field fixed effects are a subset of the field-by-cycle-day fixed effects, this implies that our experiment yields causal impacts after conditioning on δ_{tf}. Appendix D examines simulated data and shows that peer effect estimates are biased upward if we omit the field-by-date fixed effects from Equation 1.

[29] Note that the field-by-date fixed effects absorb any level differences in ability across gangs.

[30] We do not cluster at a more general level (such as by worker or by gang) because workers are assigned to independently-randomized sets of co-workers on each cycle day. This approach follows Mas and Moretti (2009), who cluster their standard errors at the level of a worker-by-date rather than by worker or by store.

[31] Gangs sometimes spend multiple cycle days on the same field, which implies that clustering at the worker-by-cycle-day level is not the same as clustering at the field-by-date level.

[32] Appendix B details how we construct the standard errors.

[33] Bandiera, Barankay, and Rasul (2010) and Park (2019) use similar approaches to estimating ability as permanent productivity. We prefer this measure of ability to the use of pre-experiment mean output for two reasons. First, this variable is available for all workers in our sample, while worker turnover means some workers will not exist in the pre-experiment data. Second, any measure of output from the tea estate will be affected by peer effects, and thus will not represent the worker's true underlying ability level. In the pre-experiment period, we lack the data on workplace peers needed to correct the ability measures for peer effects. This implies that pre-experiment output has measurement error of unknown magnitude and sign for each worker.

[34] To be clear, the set of possible co-workers is based on the gang for worker i so that $\mathbf{M}_{igt}$ is a vector of $J_i - 1$ dummy variables, where J_i is the total number of pluckers in worker i's gang.

of the worker fixed effects μ_{i-f} under the assumption that each individual worker can have any effect on his or her co-workers.[35] We use the index g to denote all fields except for f in Equation 1. As detailed below, the term μ_{i-f} is the ability measure for worker i using all fields except field f, and we rely on this measure to address concerns over spatial spillovers. Using the results from Equation 2, we define $\overline{Ability}_{-i-f,t} = \overline{\mu}_{-i-ft}$ as our measure of peer influence.[36]

The resulting ability measure has a well-behaved distribution and also correlates well with known determinants of productivity in our sample. The kernel density of ability is shown in Panel A of Figure 2. Ability appears to be approximately log-normally distributed, and a Kolmogorov-Smirnov test fails to reject this null hypothesis. This is consistent with the kernel density of log ability (Panel B). Appendix Table A1 shows a linear regression of ability on a vector of determinants of productivity. Ability is positively correlated with experience, and this relationship is highly nonlinear. Women have lower productivity on average than otherwise-similar men. This is likely because physical strength determines how much tea a plucker can carry at one time and how quickly they can pull leaves off the bushes.

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: $cov(\overline{Ability}_{-i-f,t}, \epsilon_{ift}) = 0$. A violation of this assumption would occur if supervisors assign workers with higher ability to work on particularly productive areas of a field. Our intervention eliminates this possibility by randomly assigning workers to plots within a field; hence, we can purge estimates β of any endogenous sorting effects.

In Table 2, we provide evidence to support the assumption that there is no correlation between peer ability and unobserved determinants of individual productivity on days when

[35] One additional assumption for identification is that the functional form of any co-worker peer effects is additively separable across workers.

[36] Note that the peer ability measure in Equation 1 does not take into account co-worker absences (non-compliance). This implies that our model is an intent-to-treat specification. Absences are very rare in our sample: the work attendance rate is 87 percent.

peers are randomly assigned. Specifically, Table 2 shows a series of regressions of workers' own ability on the mean of their co-workers' ability.[37] The results in Column (1) provide some evidence on the importance of our randomization of workers: 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.[38] These are days on which we explicitly did not randomize workers; instead, gang supervisors implemented plot assignments through the status quo system. In line with our random assignment procedure, 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.[39] Appendix Table A3 provides additional tests of balance and shows that there are no statistically significant correlations between (baseline) worker characteristics (e.g., age or experience) on the mean of peer ability.

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, Kroft, and Notowidigdo (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 her, and vice versa for a worker at the bottom of the distribution. Caeyers and Fafchamps (2016) call this phenomenon "exclusion bias": since the

[37]We follow the recommendation of Guryan, Kroft, and Notowidigdo (2009) and include the leave-one-out gang mean of ability in our test of random assignment. The inclusion of this term corrects for exclusion bias in tests for random assignment, but only completely eliminates the bias in the case of non-overlapping peer groups (Caeyers and Fafchamps 2016).

[38]Further support for the importance of randomization comes from an analysis of whether workers sort based on their social networks. Specifically, we use the daily worker panel to examine whether workers are more likely to work near their friends on cycle day 1 (relative to other work days where plot assignments were randomized). Appendix Table A2 reports estimates where we regress an indicator for working near at least one friend on a dummy for cycle day 1. The results show that workers are about 8.2 percentage points more likely to work near any friend on cycle day 1 relative to other days when peers were randomized; this estimate represents a 40 percent increase relative to the mean on the cycle days with random assignment. These estimates support the hypothesis that workers purposely sort on social ties when they are free to do so.

[39]Columns (1)-(3) of Table 2 show a positive coefficient on the gang leave-one-out mean, which indicates that there is positive assortative matching into gangs: some gangs have systematically higher-ability workers. In Column (4), the date-by-field fixed effects hold gangs constant, thereby giving the coefficient for the leave-one-out mean the usual negative sign, in line with Guryan, Kroft, and Notowidigdo (2009).

worker's ability appears in the error term of the regression, there is a mechanical negative correlation between peer ability and the error term. This results in coefficient estimates that are downward-biased. Unlike classical measurement error, this bias can push estimates through zero and into wrong-signed values. Since we cannot perfectly measure worker ability, even controlling for the worker's own estimated ability will leave some component of ability in the error term, and estimated peer effects will be negatively biased. The small (and insignificant) negative correlations between own ability and peer ability in Column (4) of Table 2 are consistent with the existence of exclusion bias.[40]

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 Equation 1. These worker fixed effects break any potential correlation between the fixed component of the error term and the ability of a worker's peers. Intuitively, the worker fixed effects difference out *all* fixed worker-level contributions to output, which solves the exclusion bias problem. We conduct simulations to confirm that our coefficient estimates are substantially downward-biased if we omit the worker fixed effects, but approximately correct if we include them (see Appendix D).

Third, an additional concern in estimating peer effect models is that spatial correlations in output can generate correlations between the output of co-workers and individuals. For example, suppose that one area 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 area on each day, and also increase their estimated ability. Such spatial correlations in plot quality are a potential concern in our setting because our workers return to the same randomly-assigned plots each time they come back to the same field, and plot locations drive the random variation in peer

[40]We can interpret the negative sign as an indication of exclusion bias because we have overlapping peer groups in this context.

composition from our experiment. We therefore cannot control for plot fixed effects, which would address this problem.[41]

Our specification in Equation 2 addresses this issue by using a double leave-one-out approach 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 field f when computing the estimated peer ability levels for use in Equation 1. We do this by restricting the sample to the set of g fields other than f when estimating Equation 2.[42] For example, this allows us to construct ability estimates for workers when they are on Field 5 that exclude Field 5 observations. As a result, we always estimate Equation 1 using a measure of mean peer ability that excludes 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, cannot cause violations of the assumption that $cov(\overline{Ability}_{-i-f,t}, \epsilon_{ift}) = 0$.[43,44]

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 $\overline{Ability}_{-i-f,t}$ 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 $\overline{Ability}_{-i-f,t}$.[45] Previous research has used this type of specification and found evidence of notable hetero-

[41] In settings where peers vary independently from work locations, it is standard to control for location fixed effects, in part to address exactly this issue. For example, Mas and Moretti (2009) include cash register fixed effects in their regressions.

[42] Appendix C provides further details on how the double leave-one-out approach addresses spatial correlations when estimating the permanent productivity of workers.

[43] The double leave-one-out approach implies that we use less data to estimate peer ability. Due to measurement error, this will attenuate estimated peer ability relative to approaches that use all of the daily data. Appendix C provides a detailed discussion of measurement error in the double leave-one-out approach.

[44] Appendix D assesses how estimates of the impact of mean peer ability based on the double leave-one-out approach compare to estimates based on an approach which uses all daily data, and shows that the latter is sharply upward-biased if there are spatially-correlated shocks to plot quality.

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

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

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.

geneity in peer effects across the distribution of student ability (Hoxby and Weingarth 2005; Carrell, Fullerton, and West 2009; Imberman, Kugler, and Sacerdote 2012; Carrell, Sacerdote, and West 2013; Booij, Leuven, and Oosterbeek 2017) and worker ability (Mas and Moretti 2009; Cornelissen, Dustmann, and Schonberg 2017).

6 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. A 10 percent increase in mean peer ability is associated with a 0.3 percent increase in the daily kilograms of tea plucked for each worker.[46] Column (2) shows that our estimates are essentially unchanged when we condition on date-by-location fixed effects (instead of separate fixed effects for date and location). Figure 3 presents our main results graphically, as a binned scatterplot that controls for worker and date-by-location fixed effects. There is a positive, linear relationship between the log of mean peer ability and the log of output. In Appendix Table A4, we test whether our peer effect estimates are sensitive to the inclusion of a variety of other measures of peer characteristics (e.g., mean peer age) in Equation 1. Across these specifications, the estimated impact of mean peer ability is consistently positive and statistically significant.[47,48]

Relative to the literature, these estimates are smaller in magnitude than what is found in contexts where individuals are engaged in joint production. For example, our estimate is about one fifth of the size of estimates from Mas and Moretti (2009) in their study of

[46] Note that due to data limitations, we cannot examine to what degree the effect is driven by changes in a worker's pace versus time spent at work. However, the work environment has a number of restrictions that constrain time spent working. A few examples are as follows. The start and end of the work day are fixed. Workers who arrive late in the morning are also sent home for the day. Finally, workers are both expected to stay until the end of the day, and to take longer breaks only during designated break times.

[47] As an additional check, we find that we cannot reject the hypothesis that peer effects are equal in the first and second half of the agricultural season.

[48] Appendix Table A5 builds on our main analysis by estimating models that include a measure of the ability of strictly second-order co-workers (workers who are adjacent to a focal worker's neighbors and *not* directly adjacent to the focal worker). The results in Appendix Table A5 show that there are no detectable peer effects stemming from strictly second-order neighbors.

supermarket cashiers. We can reject the hypothesis that our estimates are equal to the effects detected in their work at the one percent significance level. As an additional benchmark, it is worth noting that our main point estimate is half as large as the result reported in Cornelissen, Dustmann, and Schonberg (2017), who study workplace peer effects using a fixed effects approach and administrative data for all workers and firms in one large labor market. For workers in low-skill occupations, their estimate implies that a 10 percent increase in peer ability increases wages by 0.6 percent. We can reject the hypothesis of equal effects with their study at the one percent level.[49]

We also verify that our results are not driven by selective attendance at work due to changes in peer quality. We show this by creating a panel of observations for all days over the course of the season and creating an indicator for whether or not a worker was at work and plucking tea.[50] Appendix Table A6 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. The lack of effects on attendance is consistent with the idea that the incentive to attend work is strong in general. Moreover, Brune (2015) shows that the rate of attendance does not easily move from its high baseline level of roughly 87 percent unless there is an explicit incentive.[51]

The double leave-one-out estimator matters for our results, suggesting that correlated shocks would otherwise cause upward bias in our estimates of peer effects. Appendix Table A7 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 include the same field that is used to measure output. Similar to our main results, the estimates

[49]The positive impact that we detect also contrasts with Guryan, Kroft, and Notowidigdo (2009) who study peer effects in a setting without joint production by focusing on golfers. Their estimates imply that a 10 percent increase in peer ability *reduces* own golfing performance by 0.01 percent.

[50]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).

[51]As one point of comparison, studies have shown that the work attendance for teachers in developing countries is around 75 percent (Kremer et al. 2005; Duflo, Hanna, and Ryan 2012).

are positive and significant. However, 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.[52]

We also test whether workers have symmetric responses when they have higher or lower ability relative to their nearby peers. To conduct this test, we modify the approach used by Bandiera, Barankay, and Rasul (2010) to study whether workers respond asymmetrically to friends who have higher or lower ability. For our analysis, we compute the absolute value of the difference between a worker's own ability and the mean of peer ability. We interact the log of this measure with indicators for whether a worker's own ability is higher or lower than the mean of peer ability. Using these measures, we estimate a model of heterogeneous peer effects. Appendix Table A12 shows that peer effects appear to be symmetric: worker productivity increases and decreases by similar magnitudes when a worker is less and more able than their peers, respectively.

Finally, we test for variation in peer effects across workers with different individual characteristics. Table 4 shows treatment effect heterogeneity by gender, age, and a 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 the differences across age categories are not statistically significant (the p-value for a test of the null hypothesis of equal effects is 0.216) .

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 co-worker ability—an effect twice as large as what we see for the overall sample. This effect is strongly statistically significant (p-value=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. The effects for

[52] Appendix D also uses simulated data to assess the performance of the double leave-one-out approach.

women are not due to a correlation between other attributes of women and heterogeneous responses to peers. We demonstrate this in Appendix Table A8 by showing that controlling for several characteristics of workers at the same time leaves the magnitude and standard error of the male-female difference in treatment effects essentially unchanged.[53,54] Appendix Table A10 reinforces this conclusion by showing consistent evidence that females appear to be more sensitive to peer influence relative to men within the same age-range, own ability, or experience level.[55]

In addition to differing in the magnitudes of the peer effects they experience, men and women differ in terms of their estimated ability levels. Figure 4 shows kernel densities of worker ability by gender. The male distribution is further to the right than the female distribution. Appendix Table A11 shows summary statistics for ability by gender and shows that men have an underlying productivity level that is 8.4 kilograms of tea higher (on average) than women.

To the best of our knowledge, only a handful of studies examine peer effects by gender in either workplace or educational settings. The studies most relevant to our analysis of peer effects by gender include Harmon, Fisman, and Kamenica (2019), Beugnot et al. (2019), Hahn et al. (2017), Lavy, Silva, and Weinhardt (2012), and Stinebrickner and Stinebrickner (2006). These studies vary substantially in terms of setting, design, and the parameters estimated. That said, all consistently provide evidence that peers have different impacts based on the focal individual's gender. The parameters studied in Hahn et al. (2017), Lavy,

[53]Note that Appendix Table A8 includes all interaction terms at once (unlike the model used for Table 4). This limits us to including fully-saturated terms in just one interaction (gender). The other terms show the difference from either the male or female treatment effect.

[54]Appendix Table A9 extends our gender analysis by reporting results from models that include interaction terms for own gender and gender-specific peer ability. The results in Column (1) show that separate measures of male and female co-worker ability have similar estimated peer impacts. Column (2) shows that female workers respond to both male and female co-worker ability. Note that the sample for this analysis requires non-missing information on gender for all neighbors. This condition reduces the sample size for this analysis relative to Table 4.

[55]Appendix Table A10 demonstrates this by augmenting several specifications to include interactions for gender, mean peer ability and worker characteristics such as age group, own ability, or experience level. Each pair of columns reports results point estimates from the gender-specific interaction terms that are estimated from a single regression.

Silva, and Weinhardt (2012), and Stinebrickner and Stinebrickner (2006) are arguably the most closely-related to our study, although they both focus on education settings rather than workplaces. Hahn et al. (2017) conduct an experimental evaluation where treated and control students study in groups with friends or non-friends, respectively. They find that low-ability female students benefit from studying with friends with no corresponding significant effects for males. Lavy, Silva, and Weinhardt (2012) find that girls, particularly those in the bottom half of the ability distribution, benefit from having high-ability peers. Stinebrickner and Stinebrickner (2006) study college roommates, finding that the academic performance of girls is more sensitive to their roommate's high school grade point average. We build on these prior results by showing evidence that women also respond more strongly to peers in a real-world workplace setting.[56,57]

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 reassigning a high-ability peer from one group to another would have equal and offsetting effects.[58] Because men do not experience peer effects in our sample, in principle 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. On average, surrounding women with only male peers would raise their mean peer ability by 8.4 kilograms per day, and would

[56] Harmon, Fisman, and Kamenica (2019) study a distinct type of peer effect from what we consider. They examine peer effects for politicians that arise from sitting near each other. Exploiting quasi-random variation in seating assignment, they find that sitting next to each other reduces the probability that two politicians from the same party differ in their vote. They also find that this type of peer effect is larger for female politicians.

[57] Why might peer effects vary by gender? The psychological literature suggests peers may matter more for females because they are more positively and cooperatively influenced by others (Cross and Madson 1997). The mechanism described in their research is based on gender differences in self-definition and identity.

[58] Note that Appendix Table A12 suggests that peer effects in our setting are symmetric by ability level: the magnitude of peer effects is similar whether workers are paired with faster or slower co-workers. This finding is important to keep in mind when considering the gains from re-allocating workers.

raise the log of their mean peer ability by 0.14. This would imply an increase in productivity of 0.8 percent (=0.14 x 0.6).[59]

7 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 individual 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. To preview our results, we find no evidence that standard explanations of peer effects such as socialization or learning can explain our findings. Rather, we find evidence that suggests the estimates are driven by the impact of peers on worker motivation.[60]

7.1 Socialization

One leading mechanism for workplace peer effects is socialization between workers. In a setting similar to ours, Bandiera, Barankay, and Rasul (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 work near friends who were typically fast, they work

[59] Our study is not well-powered to detect the effect of switching women's peers from 100 percent female to 100 percent male. A regression of output on the share of peers who are male, for just the women in our sample, yields a point estimate of 0.004 (one-half of the result calculated above), with a 95 percent confidence interval ranging from -0.018 to 0.026. Our MDE at 80 percent power is 0.032—four times as large as the effect size we would expect to see.

[60] Our main analysis examines mechanisms pooling the entire sample of workers. To extend on these results, Appendix Table A10 conducts analysis of mechanisms by gender. The results align with our main estimates by providing evidence that women are more sensitive to peer effects relative to men in our sample.

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 (2019), who studies workers at a seafood processing plant and finds that a worker's productivity drops by six percent when working near a friend.

Using data on social networks, Table 5 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).[61] The results show that a 10 percent increase in the mean ability of non-friends increases worker productivity by 0.28 percent (p-value=0.028), which is nearly identical to the impact that we obtain from our main specification in Table 3. In contrast to these effects for non-friends, the point estimate of the effect of increasing the mean ability of friends is smaller and not statistically significant.[62] As robustness checks, we also report estimates of the impact of (log) mean peer ability for the subsamples of observations when individuals have no friends as peers (Column 4) and at least one friend as a peer (Column 5). In line with Column (3), the point estimate for peer effects in the no-friends sample is much larger than the estimate in the sample with at least one friend.

As an additional piece of evidence, we examine survey data and also find responses that are consistent with the idea that peer effects in our setting may not be driven by socialization.

[61] Appendix Table A13 provides summary statistics for the measures of mean friend and non-friend peer ability. These statistics show that the variation in mean peer ability is generally similar for the friend and non-friend groups, and so the difference in peer effects is not driven by differences in ability levels.

[62] Using the point estimates from Table 5, the effect of having at least one non-friend (evaluated at the mean of non-friend ability) is a 0.43 percent increase in productivity, while the effect of having at least one friend (evaluated at the mean of friend ability) is a 0.10 percent decrease in productivity. We can reject the hypothesis of equal impacts of adding friends and non-friends at the 10 percent level.

Approximately 60 percent of workers in our sample report never spending more than 10 minutes of any work day talking to co-workers. This finding is consistent with the idea that communication is difficult due to the size of plots: a plucker is typically 25 meters away from a peer working in an adjacent plot.[63]

7.2 Learning

Another potential mechanism to explain our findings is learning (i.e., knowledge spillovers). It is conceivable that plot neighbors learn from observing each other working, thereby generating the positive effects that we observe.[64] 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 to be largest for workers who have relatively less experience. Second, we test whether lagged measures of peer ability appear to have any effect on a worker's current productivity. If workers learn from their co-workers, lagged measures of co-worker ability will likely affect current productivity.[65]

Table 6 reports estimates from augmented versions of Equation 1 in which we add interactions with 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 data. Column (2) builds on our main specification by adding an interaction between a dummy indicating status as a new worker (i.e., having no prior experience) and our measure of peer ability. The point estimate for this interaction is not statistically significant

[63]While the issue of distance between plots suggests that socialization is not a likely mechanism, it does not explain why non-friends have larger impacts than friends. One possible explanation that could rationalize the pattern of results is that friends are aware of each others' productivity regardless of whether they work in close proximity. This awareness could occur because friends may try to interact with each other during break times or at the end of work. In contrast, workers may only observe non-friend productivity due to plot neighbor assignment. This difference in awareness of productivity may be key to activating a motivation mechanism that could drive workplace peer effects.

[64]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.

[65]It is possible that peers help workers learn skills that enhance productivity under specific conditions that vary at the plot and day level. This type of learning spillover would not generate lagged peer effects.

and if anything would imply smaller peer effects 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). The results for this specification are not precise, although the point estimates for the least experience and most experienced workers are relatively similar.[66] 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 mean co-worker ability from one cycle day ago ("t-1"), two cycle days ago ("t-2") and three cycle days ago ("t-3").[67] Column (3) shows results from our preferred specification, which includes current peer ability and all lagged measures. These results show that current peers have a positive impact on productivity while there is no detectable impact of any lagged measure. Figure 5 builds on this analysis by estimating a model that includes both lag and leading measures of peer ability. These results again show that only current peers have an impact of productivity, and there are no spillovers of peer effects across days. Moreover, the estimates for lead measures of peer ability serve as a test of identification: current productivity should not depend on future measures of the mean ability of randomly assigned peers.

[66] We fail to reject the null hypothesis in a test that peer effects are equal for workers of all experience levels.

[67] In the administrative data, there are cases where workers stay on the same field for multiple days. To avoid treating the same set of assigned peers as its own lag or lead, we use leads and lags in terms of *cycle days* rather than dates. The sample sizes differ across specifications for two reasons. First, we set the leads or lags to missing values once they overlap with one another at the start and end of the cycle, which matters for workers who appear on just a few fields. Second, we ensure that there can be no lags at the beginning of the season and no leads at the end of the season.

7.3 Psychological Mechanisms: Motivation vs. Shame

The institutional setting at Lujeri and the evidence so far rule out many standard explanations for peer effects. The remaining possibility is that some type of psychological channel drives the peer effects. To explore this class of mechanisms further, this section considers testable implications from a basic model of peer effects. Specifically, we follow Kandel and Lazear (1992) and consider the following stylized utility function of worker effort (e):

$$u(e,\theta) = \begin{cases} w(e) - c(e) & \text{if } \theta = \theta_L \\ w(e) - c(e) + p(e) & \text{if } \theta = \theta_H \end{cases} \tag{3}$$

where the functions $w(.)$, $c(.)$ and $p(.)$ are the wage, cost and "peer pressure" functions, respectively. The variable θ represents peer quality, which can be low (θ_L) or high (θ_H). Given that workers in our setting are paid piece rates, we assume the wage function increases monotonically with worker productivity, which is determined by 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. If a worker increases effort in the presence of high-ability peers, the peer pressure function has a positive first derivative (i.e., $\partial p/\partial e > 0$) to reflect the extra marginal return to effort.

Appendix Figure A2 illustrates two common characterizations of psychological peer effects in this model. First, peer effects could reduce the level of a worker's utility due to shame or last-place aversion.[68] As shown in Panel B of the figure, this implies that the peer pressure function $p(.)$ is negative. In this case, high-ability peers reduce total utility, and workers increase effort as a way of minimizing the utility loss. Second, other psychological mechanisms such as motivation or "contagious enthusiasm" suggest that the function $p(.)$ is

[68] Note that one type of shame-based peer effect is the mutual monitoring or threat of social sanctions that workers can use to overcome free-rider problems when workers engage in team production (Herbst and Mas 2015).

positive.[69] As shown in Panel A of the figure, this implies that $p(.)$ is positive. In this case, high-ability peers raise a worker's utility level, and workers increase their effort to maximize this benefit.

Different types of psychological mechanisms have distinct predictions for worker welfare. 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 increases utility, 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 co-workers is beneficial. Workers should therefore be willing to pay for higher-ability peers. To test this prediction, we conducted a supplementary survey for a subset of tea workers during the next harvest season (2015-2016), after we completed our main experiment. We asked workers whether they wanted higher-ability peers (i.e., a peer in the top 10 percent of the gang in terms of average kilograms collected per day), 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 there was a one in ten chance of being selected to have one choice implemented. Appendix F provides details on the survey prompt and questions that we used to collect responses.

Panel A of Table 8 reports that 72 percent of workers would like to be assigned next to a

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

fast (high ability) peer in their gang.[70,71] When asked for the main reason for their choices in an open-ended question, 74 percent workers state that faster peers provide motivation. Only 9 percent state learning as a reason for wanting higher-ability peers.[72] Furthermore, Panel B shows that these workers seeking reassignment are willing to pay for these peers: 59 percent of workers would be willing to give up one bar of soap while 46 percent would be willing to give up two bars of soap.

Overall, the results from our willingness-to-pay experiment suggest that motivation is the driver of peer effects in our sample. This interpretation is supported by the context of our study, which makes peer pressure unlikely because workers receive piece rates and do not work in teams. Further support also stems from our analysis of learning and socialization peer effects: we find no evidence that suggests these mechanisms drive impacts in our setting. The willingness-to-pay results rule out a range of other potential psychological mechanisms posited in the literature, such as shame, reputation or a desire to avoid being last (Kandel and Lazear 1992; Kuziemko et al. 2014; Tincani 2015; Breza and Chandrasekhar 2015). Since

[70]When we examine the survey results by gender in Appendix Table A14, we find that 70 and 75 percent of men and women prefer to work near a fast peer, respectively. While women do have higher demand for fast peers, this result should be interpreted with caution given that the difference is not statistically significant (p-value=0.19). Furthermore, this pattern of results is puzzling given that Table 4 shows that the point estimate for peer effects for males is small and not statistically significant. One explanation that reconciles this pattern is that there are small peer effects for men that we are unable to statistically detect. The confidence interval surrounding the point estimate for males in Table 4 suggests that we can only rule out that a 10 percent increase in peers would increase productivity by 0.45 percentage points. Another explanation is that other factors correlated with gender could reduce the apparent male-female gap. When we adjust the gender differences for other covariates, the gap widens (column 4). Notably, this contrasts with the gender gap in peer effects, which is not affected by controlling for other factors (see Appendix Tables A8 and A10).

[71]We surveyed two sets of workers in the harvest season after our experiment (2015-2016): 466 workers who had been part of our original experiment, and 256 workers who were workers who had not been part of our original experiment, but were on-site and easy to interview. Our main analysis of preferences studies the sample of 434 (93 percent) of workers who experienced the original intervention and had non-missing demographic information. We focus on this sample to ensure that we can analyze preference while controlling for observed characteristics that are correlated with preferences. In the full sample of 722 surveyed workers, the overall demand for fast peers is similar to the sample of 434 workers who were part of our main experiment and had non-missing demographic information.

[72]These results are robust to controlling for the date of the survey. Workers also prefer faster peers if they are allowed to choose other kinds of peers, such as slow peers or friends; there are no order effects in these results. We registered a pre-analysis plan for the analysis of these preference data; we deviate from it by omitting Part II, which estimates the marginal willingness-to-pay for faster peers, because the estimated marginal willingness-to-pay was unreasonably high.

workers are willing to pay for faster peers, a shame-type mechanism can only be operative inasmuch as it serves as a commitment device, inducing workers to reach a higher level of effort that they truly would like to achieve. This type of behavior would be consistent with the commitment and goal-setting behavior observed in Kaur, Kremer, and Mullainathan (2010) and Dupas, Robinson, and Saavedra (2018).

8 Conclusion

This paper provides new evidence on workplace peer effects by conducting a field experiment with an agricultural firm in Malawi. We randomly assigned tea pluckers to plots within fields and use this variation in peer composition to examine the effect of mean co-worker ability (permanent productivity) on a worker's output. Unlike prior studies such as Mas and Moretti (2009), our setting provides a unique test of peer effects because workers receive piece rates (rather than fixed wages) and production complementarities are negligible.

Using administrative data on daily productivity, we find that the average of co-worker ability has a positive and statistically-significant effect: increasing the average of co-worker ability by 10 percent increases a worker's output 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 based 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 survey in the next harvesting season in which we asked workers to choose new co-workers as plot neighbors. We find that 72 percent of workers wanted to be assigned to a fast (high-productivity) co-worker. Moreover, workers were willing to pay for faster co-workers: 46 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, 74 percent of workers state that working near faster peers motivates them.

Overall, our analysis provides new evidence on the mechanisms that drive peer effects in the workplace. A better understanding of the forces that drive peer effects helps address the question of how firms might be able to harness the power of peer effects. We provide evidence that peer effects in our setting stem from the effect that co-workers have on motivation. Our results also suggest that shame and rank preferences do not drive the detected peer effects. This finding is important since these latter mechanisms suggest that workers may resist exposure to high-performing co-workers, even if these peers enhance overall firm productivity.

References

Amodio, Francesco and Miguel Martinez-Carrasco. 2018. "Input Allocation, Workforce Management and Productivity Spillovers: Evidence from Personnel Data." *The Review of Economic Studies*, 85(4): 1937–1970.

Angrist, Joshua. 2014. "The Perils of Peer Effects." *Labour Economics*, 30: 98–108.

Arcidiacono, Peter, Joshua Kinsler, and Joseph Price. 2016. "Productivity Spillovers in Team Production: Evidence from Professional Basketball." *Journal of Labor Economics*, 35(1): 191–225.

Azoulay, Pierre, Joshua S. Graff Zivin, and Jialan Wang. 2010. "Superstar Extinction." *The Quarterly Journal of Economics*, 125(2): 549–589.

Babcock, Philip, Kelly Bedard, Gary Charness, John Hartman, and Heather Royer. 2015. "Letting Down the Team? Social Effects of Team Incentives." *Journal of the European Economic Association*, 13(5): 841–870.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2010. "Social Incentives in the Workplace." *The Review of Economic Studies*, 77(2): 417–458.

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.

Battaglini, Marco, Roland Benabou, and Jean Tirole. 2005. "Self-control in Peer Groups." *Journal of Economic Theory*, 123(2): 105–134.

Battisti, Michele. 2017. "High Wage Workers and High Wage Peers." *Labour Economics*, 46: 47–63.

Beugnot, Julie, Bernard Fortin, Guy Lacroix, and Marie Claire Villeval. 2019. "Gender and Peer Effects on Performance in Social Networks." *European Economic Review*, 113: 207–224.

Booij, Adam, Edwin Leuven, and Hessel Oosterbeek. 2017. "Ability Peer Effects in University: Evidence from a Randomized Experiment." *The Review of Economic Studies*, 84(2): 547–578.

Breza, Emily and Arun G. Chandrasekhar. 2015. "Social Networks, Reputation and Commitment: Evidence from a Savings Monitors Experiment." NBER Working Paper 21169, National Bureau of Economic Research.

Brune, Lasse. 2015. "The Effect of Lottery-Incentives on Labor Supply: A Firm Experiment in Malawi." Working Paper.

Brune, Lasse, Eric Chyn, and Jason Kerwin. 2020. "Replication Data for: Peers and Motivation at Work: Evidence from a Firm Experiment in Malawi." Harvard Dataverse. doi:10.7910/DVN/CAXO8Y.

Caeyers, Bet and Marcel Fafchamps. 2016. "Exclusion Bias in the Estimation of Peer Effects." NBER Working Paper 22565, National Bureau of Economic Research.

Carrell, Scott, Bruce Sacerdote, and James West. 2013. "From Natural Variation to Optimal Policy? The Importance of Endogenous Peer Group Formation." *Econometrica*, 81(3): 855–882.

Carrell, Scott, Richard Fullerton, and James West. 2009. "Does Your Cohort Matter? Measuring Peer Effects in College Achievement." *Journal of Labor Economics*, 27(3): 439–464.

Cornelissen, Thomas, Christian Dustmann, and Uta Schonberg. 2017. "Peer Effects in the Workplace." *American Economic Review*, 107(2): 425–456.

Cross, Susan and Laura Madson. 1997. "Models of the Self: Self-construals and Gender." *Psychological Bulletin*, 122(1): 5–37.

De Grip, Andries and Jan Sauermann. 2012. "The Effects of Training on Own and Coworker Productivity: Evidence from a Field Experiment: Training Effects on Own and Peer Productivity." *The Economic Journal*, 122(560): 376–399.

DellaVigna, Stefano, John A. List, and Ulrike Malmendier. 2012. "Testing for Altruism and Social Pressure in Charitable Giving." *The Quarterly Journal of Economics*, 127(1): 1–56.

Duflo, Esther, Rema Hanna, and Stephen P. Ryan. 2012. "Incentives Work: Getting Teachers to Come to School." *American Economic Review*, 102(4): 1241–1278.

Dupas, Pascaline, Jonathan Robinson, and Santiago Saavedra. 2018. "The Daily Grind: Cash Needs and Labor Supply." Working Paper.

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

Gneezy, Uri and Aldo Rustichini. 2004. "Gender and Competition at a Young Age." *American Economic Review*, 94(2): 377–381.

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

Hahn, Youjin, Asadul Islam, Eleonora Patacchini, and Yves Zenou. 2017. "Do Friendship Networks Improve Female Education?" IZA Discussion Paper 10674, Institute of Labor Economics (IZA).

Harmon, Nikolaj, Raymond Fisman, and Emir Kamenica. 2019. "Peer Effects in Legislative Voting." *American Economic Journal: Applied Economics*, 11(4): 156–180.

Herbst, Daniel and Alexandre Mas. 2015. “Peer Effects on Worker Output in the Laboratory Generalize to the Field.” *Science*, 350(6260): 545–549.

Hoxby, Caroline and Gretchen Weingarth. 2005. “Taking Race Out of the Equation: School Reassignment and the Structure of Peer Effects.” Working Paper.

Imberman, Scott, Adriana Kugler, and Bruce Sacerdote. 2012. “Katrina’s Children: Evidence on the Structure of Peer Effects from Hurricane Evacuees.” *American Economic Review*, 102(5): 2048–2082.

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.

Jiang, Lingqing. 2020. “Splash with a Teammate: Peer Effects in High-stakes Tournaments.” *Journal of Economic Behavior & Organization*, 171: 165–188.

Kandel, Eugene and Edward P. Lazear. 1992. “Peer Pressure and Partnerships.” *Journal of Political Economy*, 100(4): 801–817.

Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan. 2010. “Self-Control and the Development of Work Arrangements.” *American Economic Review*, 100(2): 624–628.

Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan. 2015. “Self-Control at Work.” *Journal of Political Economy*, 123(6): 1227–1277.

Kremer, Michael, Nazmul Chaudhury, F. Halsey Rogers, Karthik Muralidharan, and Jeffrey Hammer. 2005. “Teacher Absence in India: A Snapshot.” *Journal of the European Economic Association*, 3(2-3): 658–667.

Kuziemko, Ilyana, Ryan Buell, Taly Reich, and Michael Norton. 2014. “Last-Place Aversion: Evidence and Redistributive Implications.” *The Quarterly Journal of Economics*, 129(1): 105–149.

Lavy, Victor, Olmo Silva, and Felix Weinhardt. 2012. “The Good, the Bad, and the Average: Evidence on Ability Peer Effects in Schools.” *Journal of Labor Economics*, 30(2): 367–414.

Mas, Alexandre and Enrico Moretti. 2009. “Peers at Work.” *American Economic Review*, 99(1): 112–145.

Park, Sangyoon. 2019. “Socializing at Work: Evidence from a Field Experiment with Manufacturing Workers.” *American Economic Journal: Applied Economics*, 11(3): 424–455.

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

Silver, David. 2019. “Haste or Waste? Peer Pressure and Productivity in the Emergency Department.” Working Paper.

Stinebrickner, Ralph and Todd R. Stinebrickner. 2006. "What Can Be Learned about Peer Effects Using College Roommates? Evidence from New Survey Data and Students from Disadvantaged Backgrounds." *Journal of Public Economics*, 90(8): 1435–1454.

Tincani, Michela. 2015. "Heterogeneous Peer Effects and Rank Concerns: Theory and Evidence." Working Paper.

Waldinger, Fabian. 2012. "Peer Effects in Science: Evidence from the Dismissal of Scientists in Nazi Germany." *The Review of Economic Studies*, 79(2): 838–861.

9 Figures and Tables

Figure 1: Tea Worker Field Assignment Illustrations

(a) Hypothetical Assignment for Three Tea Workers

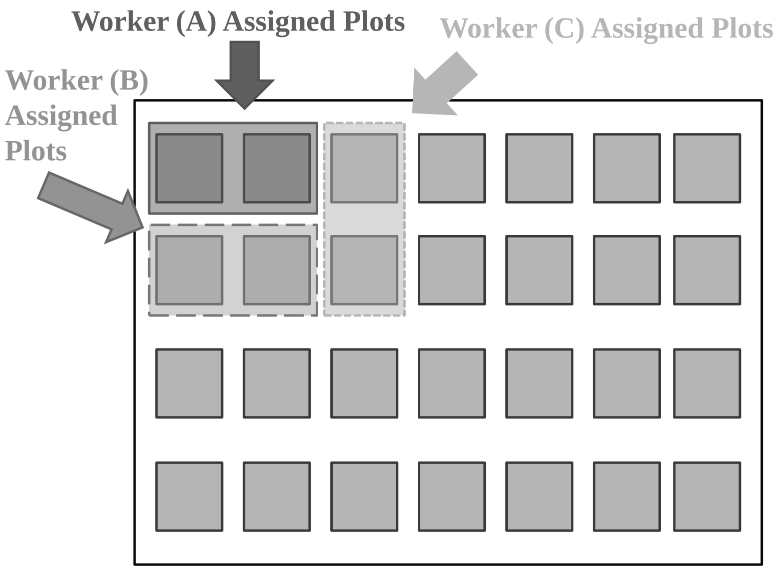

(b) Plot Assignments Change Over Days in Harvesting Cycle

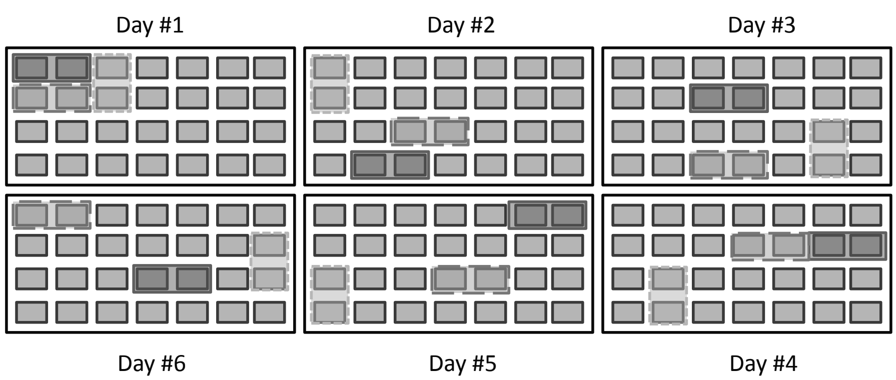

Notes: The two panels illustrate work assignments for tea workers at the Lujeri Tea Estates in Malawi. Panel A shows how three workers would be assigned two plots each. For our analysis, all workers A, B and C would be neighboring co-workers. Panel B shows how plot work assignments change over the course of a six-day harvest cycle during which the gang of workers visits six distinct 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 days #3, #5 and #6, workers A, B and C are not assigned to work in neighboring plots.

Figure 2: Distribution of Worker Ability

(a) Kernel Density of Ability

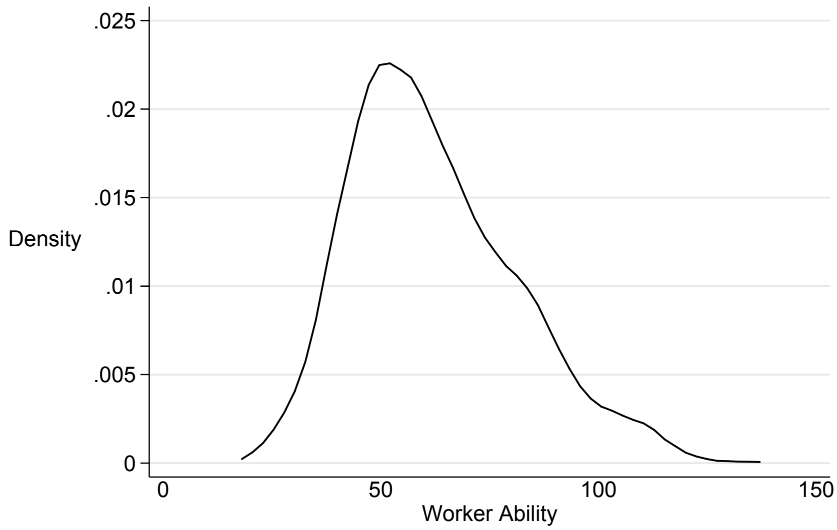

(b) Kernel Density of log(Ability)

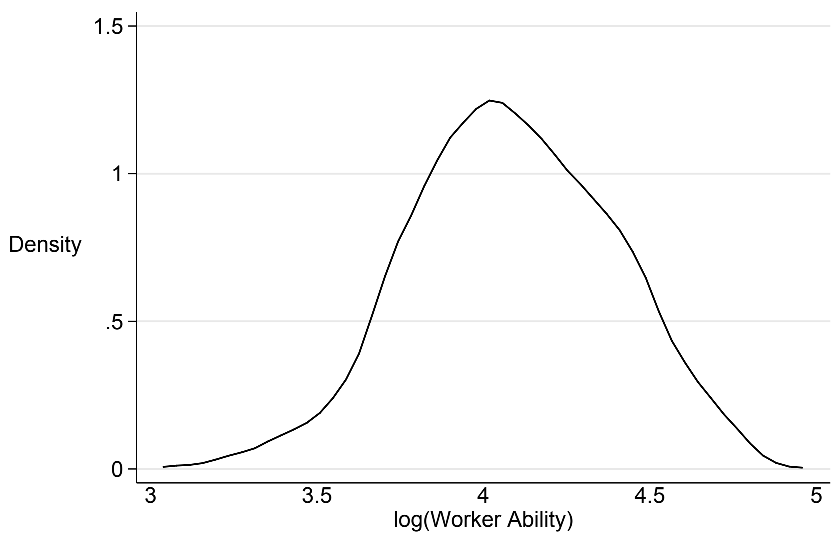

Notes: Figures present the density of estimated peer ability for a sample of tea pluckers at the Lujeri Tea Estates in Malawi. See Section 5 for the details of how we construct these estimates.

Figure 3: Binned Scatterplot of log(Output) vs. log(Mean Peer Ability), Controlling for Worker and Date-by-Location Fixed Effects

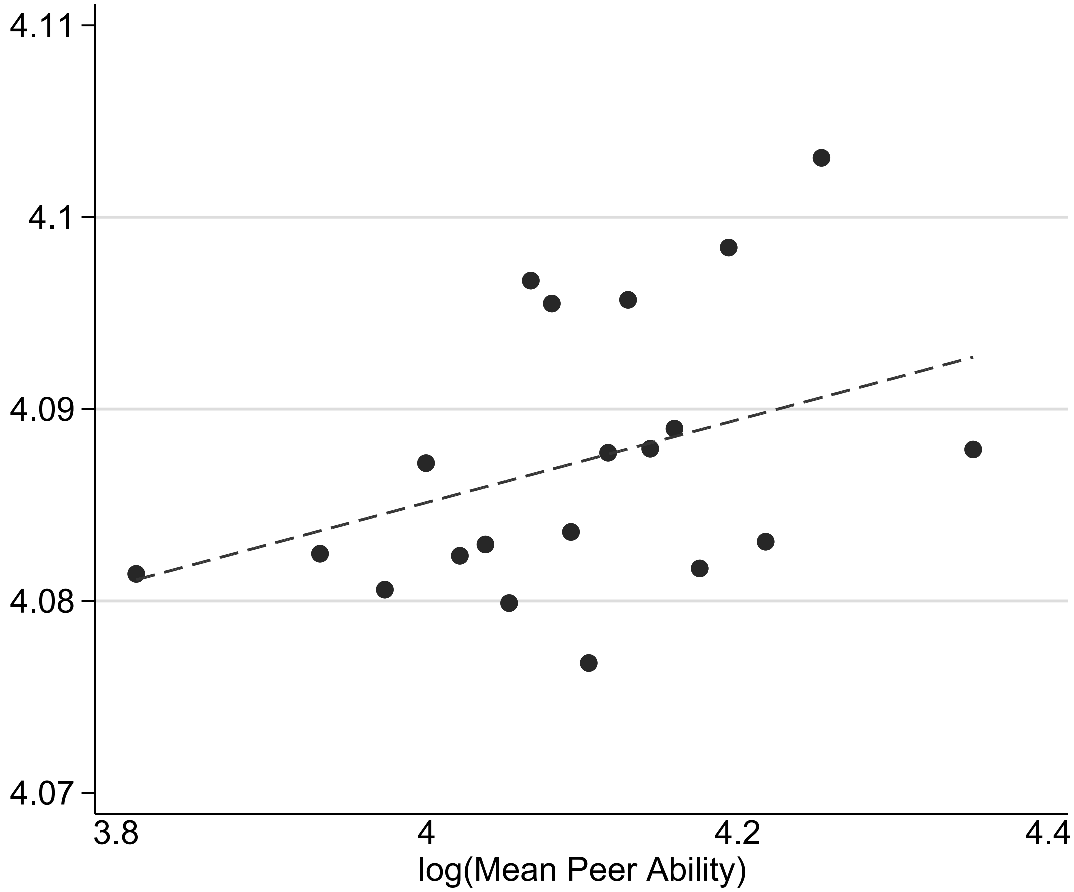

Notes: This figure plots (log) worker output (y-axis) after controlling for worker and date-by-location fixed effects by bins of (log) mean peer ability. There are 20 bins based on the ventiles of (log) mean peer ability.

Figure 4: Distribution of Worker Ability by Gender

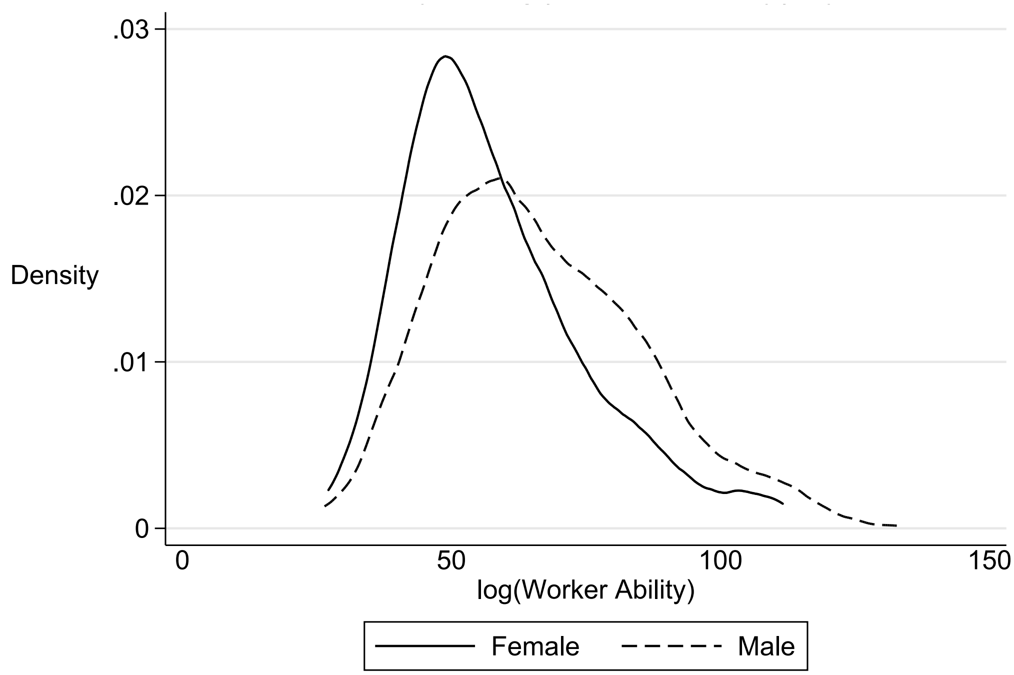

Notes: This figure presents density plots of estimated peer ability for male and female workers at Lujeri Tea Estates in Malawi. See Section 5 for the details of how we construct these estimates.

Figure 5: Estimated Lag, Lead, and Contemporaneous Effects of log(Mean Peer Ability) on log(Output)

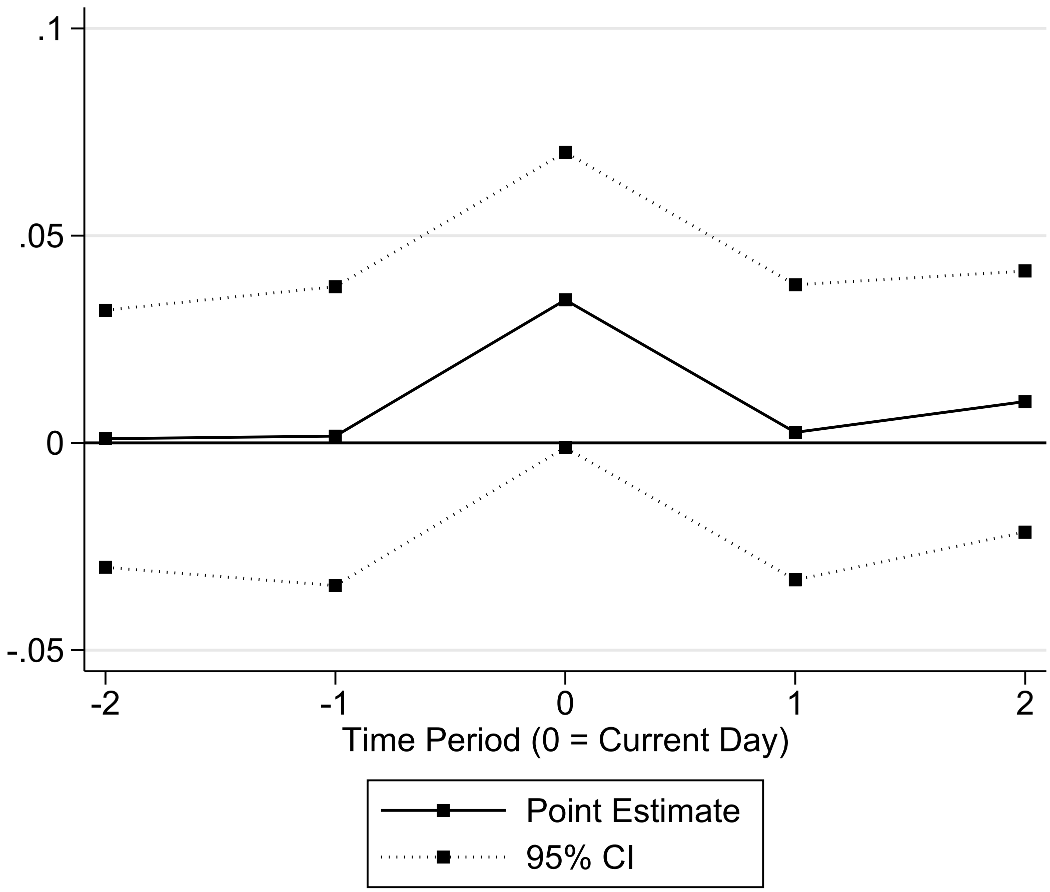

Notes: This figure plots the estimated elasticity of own output (y-axis) with respect to mean peer ability from a model which includes contemporaneous peer ability (t=0) as well as two leads (t=1, t=2) and lags (t=-1, t=-2).

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. Due to survey non-response, we are missing demographic information for 55 workers.

Table 2: Balance Test: Comparing Own and Peer Ability

	Dependent Variable: Log(Own Ability)			
	(1)	(2)	(3)	(4)
Log(Mean Peer Ability)	0.062	-0.020	-0.039	-0.046
	(0.077)	(0.034)	(0.037)	(0.032)
Log(Leave-One-Out Gang Mean Ability)	0.860***	0.945***	0.963***	-8.92***
	(0.092)	(0.041)	(0.044)	(0.967)
Cycle Day 1	Yes	Yes	No	No
Remaining Cycle Days	No	Yes	Yes	Yes
Worker Fixed Effects	No	No	No	No
Date by Location (Field) 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 dataset is a panel at the worker and day level. The results in Column (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 day 1" days—this is the sample for which we randomly assigned workers to locations within 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)	0.028**	0.030**
	(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	35,641	35,545
Adjusted R-squared	0.396	0.715

Notes: This table presents results from a regression of (log) daily output (kilograms of tea plucked) on the (log) mean ability of physically nearby co-workers. The underlying dataset 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. The measure of daily output comes from administrative data obtained from Lujeri Tea Estates; information on neighbors was recorded and collected by project staff. 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		
	(0.020)		
Female X Log(Mean Peer Ability)	0.060***		
	(0.023)		
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	
		(0.030)	
[Age 44 to 72] X [Log(Mean Peer Ability)]		0.035	
		(0.033)	
Quartiles of Own Ability			
[Own Ability Quartile 1] X [Log(Mean Peer Ability)]			0.029
			(0.032)
[Own Ability Quartile 2] X [Log(Mean Peer Ability)]			0.022
			(0.032)
[Own Ability Quartile 3] X [Log(Mean Peer Ability)]			0.042
			(0.040)
[Own Ability Quartile 4] X [Log(Mean Peer Ability)]			0.030
			(0.029)
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 (log) daily output (kilograms of tea plucked) on the (log) mean ability of physically nearby co-workers. The underlying dataset 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 project staff. 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)	(4)	(5)
Log(Non-Friends Mean Peer Ability)	0.028**		0.028**		
	(0.014)		(0.014)		
No Non-friends (=1)	0.068		0.072		
	(0.064)		(0.064)		
Log(Friends Mean Peer Ability)		0.006	0.006		
		(0.013)	(0.014)		
No Friends (=1)		0.033	0.035		
		(0.056)	(0.056)		
Log(Mean Peer Ability)				0.049***	0.010
				(0.015)	(0.058)
Worker Fixed Effects	Yes	Yes	Yes	Yes	Yes
Date by Location Fixed Effects	Yes	Yes	Yes	Yes	Yes
Sample Restriction	None	None	None	No Friends	Any Friends
Observations	35,583	35,583	35,583	28,116	7,256
Adjusted R-squared	0.715	0.715	0.715	0.711	0.728

Notes: This table presents results from a regression of (log) daily output (kilograms of tea plucked) on measures of the (log) mean ability of nearby co-workers who are friends and non-friends. The friend and non-friend peer ability measures are equal to zero when a worker has no nearby friends or non-friends, and we include indicators in the specification that equal one when this missing data problem occurs. The underlying dataset 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 project staff. 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	
	(0.014)	(0.026)	
New Worker (=1) X Log(Mean Peer Ability)		-0.022	
		(0.069)	
Quartile 1 Exp. X Log(Mean Peer Ability)			0.046
			(0.032)
Quartile 2 Exp. X Log(Mean Peer Ability)			0.018
			(0.031)
Quartile 3 Exp. X Log(Mean Peer Ability)			0.028
			(0.030)
Quartile 4 Exp. X Log(Mean Peer Ability)			0.029
			(0.030)
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 (log) daily output (kilograms of tea plucked) on the (log) mean ability of physically nearby co-workers, broken down by workers' experience at the firm. The underlying dataset 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 by project staff. 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)	0.030**	0.030**	0.028*			
	(0.014)	(0.015)	(0.016)			
Log(Mean Peer Ability), t-1	-0.001	-0.003	-0.002	-0.008	-0.012	-0.013
	(0.015)	(0.015)	(0.017)	(0.014)	(0.015)	(0.015)
Log(Mean Peer Ability), t-2		-0.002	-0.004		-0.009	-0.014
		(0.014)	(0.016)		(0.014)	(0.015)
Log(Mean Peer Ability), t-3			-0.003			-0.014
			(0.015)			(0.013)
Worker Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Date by Location Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	35,449	35,268	34,362	35,449	35,268	34,362
Adjusted R-squared	0.715	0.715	0.716	0.715	0.715	0.716

Notes: This table presents results from a regression of (log) daily output (kilograms of tea plucked) on the (log) mean ability of physically nearby co-workers, including lagged as well as contemporaneous peer ability. The underlying dataset 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 project staff. 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: Preferences for Fast Peers

	(1) Pct.	(2) Obs.
Who do you want to be reassigned next to?		
A fast plucker in your gang	0.72	434
A slow plucker in your gang	0.05	434
Any person of your choosing	0.09	434
No reassignment	0.14	434
Wants to work near a fast plucker and…		
...is willing to give up 1 bar of soap	0.59	434
...is willing to give up 2 bars of soap	0.46	434

Notes: This table presents statistics from survey data that we collected for a subset of tea pluckers at the Lujeri Tea Estates. In the survey questions, faster and slow peers were described as co-workers who are in the top or bottom 10 percent of the gang in terms of kilograms of tea plucked per day, respectively. Appendix Section F provides details of the survey prompt and questions that we used to collect responses. For the choice experiment, respondents were given a gift of two bars of soap (worth 18 percent of average daily wages) and asked if they would be willing to give up soap in exchange for being reassigned.