Peers and Motivation at Work:

Evidence from a Firm Experiment in Malawi*

Lasse Brune   Eric Chyn   Jason Kerwin

January 18, 2019

Abstract

This paper studies workplace peer effects by randomly varying work assignments at a tea estate in Malawi. We find that a 10 percent increase in mean peer ability raises productivity by 0.3 percent. In contrast to previous studies, neither production nor compensation externalities can drive the results 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": when given a choice to be re-assigned, most workers prefer working near high-ability co-workers because these peers provide motivation to work harder.

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

---

*Brune: Global Poverty Research Lab, Buffett Institute, Northwestern University, Evanston, IL 60208, USA. Email: lasse.brune@northwestern.edu; Chyn (Corresponding Author): Department of Economics, University of Virginia, P.O. Box 400182, Charlottesville, VA 22904, USA. Email: eric.chyn@virginia.edu; Kerwin: Department of Applied Economics, University of Minnesota, 1994 Buford Avenue, St. Paul, MN 55108, USA. Email: jkerwin@umn.edu. 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, Dan Keniston, Supreet Kaur, Chris Udry, Sangyoon Park, Joseph Ritter, Bryan Stuart, Tavneet Suri and seminar participants at the University of Michigan, the Minnesota Population Center, the University of Minnesota, Yale University, NEUDC and CSAE. Data collection for this project was supported by grants from the Michigan Institute for Teaching and Research in Economics (MITRE), Population Studies Center, Center for Education of Women and Rackham Graduate School at the University of Michigan. Chyn 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 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. All errors and omissions are our own.

1 Introduction

A large literature provides compelling evidence that a worker's own performance depends on his 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 et al., 2013; Babcock et al., 2015; Arcidiacono et al., 2016; Cornelissen et al., 2017; Amodio and Martinez-Carrasco, 2018). 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 et al. (2010) find evidence of learning among teachers and medical researchers, respectively.[2]

Yet, few studies provide evidence as to whether peer effects on productivity arise from psychological mechanisms such as motivation or norms. For example, a worker with self-control problems may find it easier to exert effort when surrounded by highly-productive co-workers.[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 et al. (2010) use a natural experiment to study the impact of working with friends, and they find that workers increase or decrease their productivity to match the output of their social ties.

[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 commited 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 spillovers on co-workers. Their finding is consistent with the existence of knowledge spillovers. At the same time, Waldinger (2012) and Guryan et al. (2009) find little evidence of any learning-based peer effects.

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

[4] Similarly, Kaur et al. (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.

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. This implies that any impact of peers is not due to peer pressure that arises when workers attempt to counteract free-riding problems in joint production (Kandel and Lazear, 1992). Second, we can conduct a wide-range of supplementary analyses using detailed data 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 engaged in group production and find peer effects that are about five times as large.

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 not statistically significant 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. More generally, our finding on heterogeneity is notable because it

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

suggests that there is potential for aggregate 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 et al., 2010; Park, 2017). 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.[7]

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.[8] A simple model of knowledge spillovers would suggest that 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 provides 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 this sample, we find that 71 percent of respondents preferred having a high-productivity co-worker nearby if

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

[7]As discussed in Section 2, socialization — which is a key mechanism in Bandiera et al. (2010) and Park (2017) — may be difficult in our setting because workers may not be close enough to their peers to communicate. At the same time, tea workers in our sample are sufficiently close to see adjacent co-workers and observe peer productivity.

[8]This analysis exploits the fact that the experimental design insured that workers have new peers on different days within a harvest cycle.

re-assignment were possible. When asked for the reason for their choices in an open-ended question, 83 percent of workers said higher-productivity peers provide motivation.

This paper's main contribution is to provide evidence that psychological mechanisms can drive workplace peer effects even when there are no group incentives. Furthermore, we provide insight on the type of psychological mechanisms that are important. 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 et al., 2005; Mas and Moretti, 2009; Kaur et al., 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). This distinction is important since 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.

In addition, 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 et al. (2009), Angrist (2014), and Caeyers and Fafchamps (2016), there is a mechanical negative correlation between a worker's own ability and their peers' ability. This correlation exists even if there is random assignment because a worker cannot be assigned to be her own peer. To address this issue, we randomly assigned workers to a different set of peers for each day in a work cycle. This 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 leave-one-out estimator of own ability 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 plucked leaves in baskets and empty their baskets at a central weighing station.[9] There is no explicit cooperation involved in this process, and pay is a constant piece rate for each kilogram of plucked tea.[10]

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).[11,12] 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

[9]The locations of weighing stations are fixed throughout the season for a number of logistical reasons such as coordination with the tractors 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.

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

[11]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.

[12]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 does not allow us to observe the actual number of plots a plucker works on any given day.

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.[13,14] At the completion of a harvesting cycle (most commonly 6 work days, or 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.[15,16] Panel A shows that each worker is assigned two contiguous plots (blue squares). The example highlights three workers who are colored red, green and yellow. The illustration shows that workers B and C are the immediate plot neighbors of worker A. Panel B provides an illustration showing how workers change assignments across fields covered during a 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. On these days, the workers will have different plot neighbors.

In this workplace context, peer effects could 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 they are close enough to see a peer's speed in terms of movement through the

[13]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. Note that worker absences are rare, and the attendance rate is 87 percent in our sample.

[14]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.

[15]While square plots are the most common shape, in reality the fields and plots are often not evenly-sized rectangles. 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.

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

field; tea plants do not block lines of sight since the plants are pruned to (roughly) waist height.[17] 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 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.[18] We use this information to generate randomly-ordered lists of pluckers for each day of a gang's harvesting cycle.[19] On cycle days where a gang works on multiple fields, we also randomly determine which workers are on each field.

We used these randomized lists to determine the order in which supervisors assign pluckers to plots on each field. The random assignment took advantage of the usual assignment process in which pluckers stand in a queue and receive plot assignments in the order that 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.[20] Each gang supervisor assigned workers using the randomly generated list

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

[18]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.

[19]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.

[20]An exception to our randomization is the first work day ("Cycle Day 1") in a gang's cycle. We intentionally did not randomize work assignments in this data. On these days, supervisors assigned workers

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, project staff confirmed compliance with the assignment via random spot checks several weeks after the initial assignment. As a result of our intervention, workers are 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. First, we rely on administrative data from the firm on worker productivity. Productivity is defined as kilograms of tea plucked per day and is electronically recorded by the firm for the purpose of paying employees. As a result, it is measured with minimal error. This data on worker productivity is available from 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 social networks.

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.[21] 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.

using the usual method, in which the plots are still assigned using the snaking pattern across the field, but the order of the pluckers comes from the order in which they stand in the queue. In Section 6, we use this non-random assignment on the first work day to test for endogenous assignment and sorting of workers to locations on a field.

[21]Due to survey non-response, we are missing demographic information for 5 percent of the sample ($N = 55$ individuals).

Over the course of our study period, the average daily output per worker is 69 kilograms of plucked tea leaves, and workers have on average about 5 assigned neighbors on any given day of work.

Our study focuses on studying how working alongside peers of different ability affects daily output. As detailed in Section 5, we measure a worker's ability by estimating their permanent productivity. Table 1 shows that the average ability estimate for workers in our sample is 62.19 kilograms. To provide a sense of the "treatment", Table 1 also shows the mean of nearby co-worker ability on each day. Across workers and days in our sample, the standard deviation of peer ability is nearly 13 kilograms.

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

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

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 $\overline{Ability}_{-i-ft}$, which is the mean of ability of all co-workers who are assigned to work adjacent to the plots that worker i is assigned.[22,23] The model also includes date-by-field fixed effects δ_{tf} to control for variation in harvest conditions over the

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

[23]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.

course of the season and across the tea estate.[24,25] Finally, we also control for time-invariant determinants of productivity — such as the worker's own plucking ability — by including individual-level fixed effect μ_i. We cluster all standard errors at the level of the treatment, which is at the worker-by-cycle-day level.[26] We also cluster by the combination of field and date to account for correlated shocks that might affect entire fields.[27] 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).[28]

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).[29] 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.[30] The idea is that the vector γ contains a set of parameters that absorb any possible peer effects and allows us to obtain unbiased estimates of the worker fixed effects μ_{i-f} under the assumption that each individual worker can have

[24] 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.

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

[26] We do not cluster at a more general level (such as by worker) because workers are assigned to independently-randomized treatment on each cycle day.

[27] 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.

[28] Appendix B provides details on how we construct standard errors.

[29] Bandiera et al. (2010) and Park (2017) 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.

[30] 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.

any effect on his or her co-workers.[31] We use the index g to denote all fields except for 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 $\overline{Ability}_{-i-ft} = \overline{\mu}_{-i-ft}$ as our measure of peer influence.[32]

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: $cov(\overline{Ability}_{-i-ft}, \epsilon_{ift}) = 0$. 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.

Table 2 shows a series of regressions of workers' own ability on the mean of their co-workers' ability.[33] The results in Column (1) provide some evidence on the importance of

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

[32] 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.

[33] We follow the recommendation of Guryan et al. (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

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. 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 intervention, the results in Columns (3) and (4) show that this correlation does not exist for the remainder of the sample, which supports the identifying assumption in our linear-in-means model.[34] Appendix Table A2 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 et al. (2009) and Angrist (2014), there is a mechanical negative correlation between a worker's own ability and that of her neighbors. Consider a worker who is at the top of the ability distribution. Her neighbors will necessarily be lower ability than her, and vice versa for a worker at the bottom of the distribution. Caeyers and Fafchamps (2016) call this phenomenon "exclusion bias": since the worker's ability appears in the error term of the regression, there is a mechanical negative correlation between peer ability and the error term. 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.[35]

(Caeyers and Fafchamps, 2016).

[34] 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 et al. (2009).

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

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.1).

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 composition from our experiment. We therefore cannot control for plot fixed effects, which would address this problem.[36]

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

[36]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.

Equation 2.[37] 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 $cov(\overline{Ability}_{-i-ft}, \epsilon_{ift}) = 0$.[38,39]

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 $\overline{Ability}_{-i-ft}$ 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 $\overline{Ability}_{-i-ft}$.[40] Previous research has used this type of specification and found evidence of notable heterogeneity in peer effects across the distribution of student ability (Hoxby and Weingarth, 2005; Carrell et al., 2009; Imberman et al., 2012; Carrell et al., 2013; Booij et al., 2017) and worker ability (Mas and Moretti, 2009; Cornelissen et al., 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 ability of

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

[38] 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.

[39] 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.

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

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

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.

peers is associated with a 0.3 percent increase in the daily kilograms of tea plucked for each worker.[41] Column (2) shows that our estimates are essentially unchanged when we condition on date-by-location fixed effects. 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. Relative to the literature, these estimates are smaller 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 supermarket cashiers. In Appendix Table A3, 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 changes very little.[42,43]

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.[44] Appendix Table A5 provides results from estimating Equation 1 where the dependent variable is attendance (Column 1) and plucking tea (Column 2). The point estimates are not significant and very small in magnitude, which suggests there is no impact of peers on work attendance. This lack of effects on attendance is consistent with the idea that in general the incentive to attend work is very strong, and absent an explicit incentive (as in Brune (2015)) the rate of attendance does not easily move from its high baseline level

[41] 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 the designated break times.

[42] 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.

[43] Appendix Table A4 builds on our main analysis by estimating models that include a measure of second-order co-worker ability. Second-order neighbors are defined as the workers who are adjacent to a focal worker's neighbors and *not* directly adjacent to the focal worker. The results in Appendix Table A4 show that there are no detectable peer effects stemming from second-order neighbors.

[44] 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).

of roughly 87 percent.[45]

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 A6 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 includes the same field that is used to measure output. Similar to our main results, the estimates 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.[46]

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 et al. (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 A10 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 the existence of peer effects that vary 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

[45] 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 et al., 2012).

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

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 women are not due to a correlation between other attributes of women and heterogeneous responses to peers. Appendix Table A7 shows 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.[47]

In addition to differing in the magnitudes of the peer effects they experience, men and women differ in terms of their estimated ability level. Figure 4 shows kernel densities of worker ability by gender. The male distribution is further to the right than the female distribution. Appendix Table A9 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.

The heterogeneity in peer effects and ability levels by gender is important because it allows for the possibility of raising aggregate productivity by rearranging workers. If peer effects were constant across individuals, then re-assigning a high-ability peer from one group to another would have equal and offsetting effects.[48] Because men do not experience peer effects in our sample, in principle we can raise the productivity of low-ability female workers by

[47]Appendix Table A8 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 coworker ability have similar estimated peer impacts. Column 2 shows that female workers respond to both male and female coworker 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.

[48]Note that Appendix Table A10 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.

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 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).[49]

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.

7.1 Socialization

One leading mechanism for workplace peer effects is socialization between workers. In a setting similar to ours, Bandiera et al. (2010) studied workers who picked fruit at a large agricultural firm in the UK and estimated the impact of working physically near a friend. Their analysis suggests that socialization between friends affects worker productivity. When slow fruit pickers worked near friends who were typically fast, they work harder to catch

[49] 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.

up. Similarly, relatively fast pickers slow down for their slower friends. Further evidence on the impact of friends also comes from Park (2017), who studies workers at a seafood processing plant and found that a worker's productivity drops by six percent when working near a friend.

Using data on social networks, Table 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).[50] 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 on the effect of increasing ability of friends is much smaller and not statistically significant.[51] 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. Approximately 60 percent of workers in our sample report never spending more than 10

[50] Appendix Table A11 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.

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

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.

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 work, thereby generating the positive effects that we observe.[52] 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 have any affect on a worker's current productivity. If workers learn from their co-workers, lagged measures of co-worker ability will likely affect current productivity.[53]

Table 6 reports estimates from augmented versions of Equation 1 in which we add measures of worker experience. The results in Column (1) replicate the estimate from our baseline specification for the sample of workers for whom we have self-reported experience 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 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

[52] 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.

[53] 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.

not precise, although the point estimates for the least experience and most experienced workers are relatively similar. Overall, the results in Table 6 provide no evidence that workers with less experience benefit more from working near higher ability co-workers.

We also find that results from models that include lagged measures of co-worker ability do not suggest there is any learning between co-workers. Table 7 reports estimates from an augmented version of Equation 1 which includes measures of co-worker ability measured one cycle day ago ("t-1"), two cycle days ago ("t-2") and three cycle days ago ("t-3").[54] 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. 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.

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 psychological channels drive 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

[54]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.

(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. 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 positive.[55] 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

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

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, and whether they would be willing to give up part of the compensation that they received for taking part in the survey (workers were each given two bars of soap as a token of thanks for taking the survey). Workers were informed that one worker per gang would have one of their choices implemented for real (a chance of about 1 in 50).

Panel A of Table 8 reports that 71 percent of workers would like to be assigned next to a fast (high ability) peer in their gang. Further, Panel B shows that these workers seeking re-assignment are willing to pay for these peers: 71 percent of workers who want a fast peer would be willing to give up one bar of soap while 55 percent would be willing to give up two bars of soap.[56] When asked for the main reason for their choices in an open-ended question, 83 percent workers state that faster peers provide motivation. Only 15 percent state learning as a reason for wanting higher-ability peers.[57]

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 where peer pressure is unlikely because workers receive piece rates and do not work

[56]Women have higher demand for fast peers, although the difference in the means for each sex is not statistically significant.

[57]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 this 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.

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 workers are willing to pay for faster peers, shame-type mechanisms can only be the operative mechanism inasmuch as it serves as a commitment device, inducing workers to reach a higher level of effort that they truly would like to achieve.

8 Conclusion

This paper provides novel 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.

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 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 71 percent of workers wanted to be assigned to a fast (high-productivity) co-worker. Moreover, workers were willing to pay for faster co-workers: 55 percent of workers were willing to give up two bars of soap (worth 18 percent of daily wages) that we had given them as a gift for survey participation. In open-ended follow-up questions, 83 percent of workers state that working near faster peers motivates them.

Overall, our analysis 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 Martinez-Carrasco, Miguel A. (2018). 'Input Allocation, Workforce Management and Productivity Spillovers: Evidence from Personnel Data', *The Review of Economic Studies* .

Angrist, Joshua D. (2014). 'The perils of peer effects', *Labour Economics* 30, 98–108.

Arcidiacono, Peter, Kinsler, Joshua and Price, Joseph. (2016). 'Productivity Spillovers in Team Production: Evidence from Professional Basketball', *Journal of Labor Economics* .

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

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

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

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

Battaglini, Marco, Benabou, Roland and Tirole, Jean. (2005). 'Self-control in peer groups', *Journal of Economic Theory* 123(2), 105–134.

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

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

Brune, Lasse. (2015). 'The Effect of Lottery-Incentives on Labor Supply: A Firm Experiment in Malawi', *Economic Growth Center, Yale University* .

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

Carrell, ScottE., Fullerton, RichardL. and West, JamesE. (2009). 'Does Your Cohort Matter? Measuring Peer Effects in College Achievement', *Journal of Labor Economics* 27(3), 439–464.

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

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

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

Duflo, Esther, Hanna, Rema and Ryan, Stephen P. (2012). 'Incentives Work: Getting Teachers to Come to School', *American Economic Review* 102(4), 1241–1278.
URL: *https://www.aeaweb.org/articles?id=10.1257/aer.102.4.1241*

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

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

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

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

Herbst, Daniel and Mas, Alexandre. (2015). 'Peer Effects on Worker Output in the Laboratory Generalize to the Field', *Science (New York, N.Y.)* 350(6260), 545–549.

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

Imberman, Scott A, Kugler, Adriana D and Sacerdote, Bruce I. (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 Bruegmann, Elias. (2009). 'Teaching Students and Teaching Each Other: The Importance of Peer Learning for Teachers', *American Economic Journal: Applied Economics* 1(4), 85–108.

Kandel, Eugene and Lazear, Edward P. (1992). 'Peer Pressure and Partnerships', *Journal of Political Economy* pp. 801–817.

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

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

Kremer, Michael, Chaudhury, Nazmul, Rogers, F. Halsey, Muralidharan, Karthik and Hammer, Jeffrey. (2005). 'Teacher Absence in India: A Snapshot', *Journal of the European Economic Association* 3(2-3), 658–667.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1162/jeea.2005.3.2-3.658*

Kuziemko, I., Buell, R. W., Reich, T. and Norton, M. I. (2014). 'Last-Place Aversion: Evidence and Redistributive Implications', *The Quarterly Journal of Economics* 129(1), 105–149.

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

Park, Sangyoon. (2017). 'Socializing at Work: Evidence from a Field Experiment with Manufacturing Workers'.

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

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

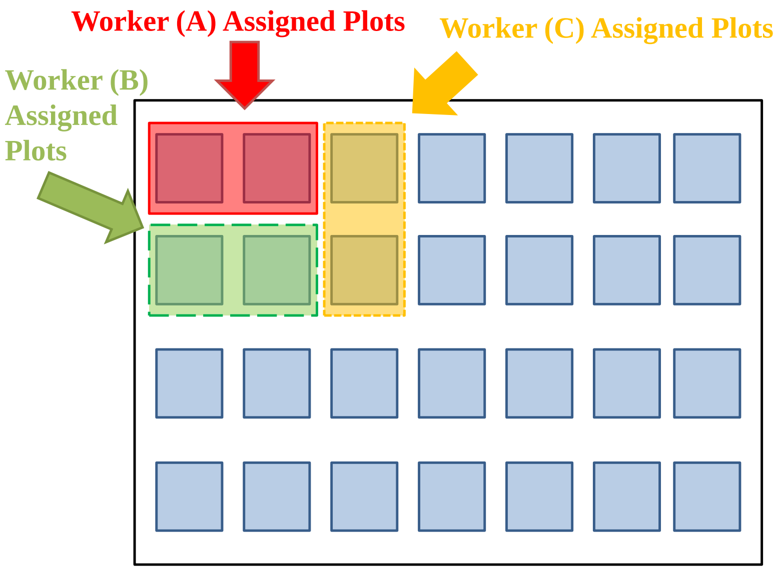

(b) Plot Assignments Change Over Days in Harvesting Cycle

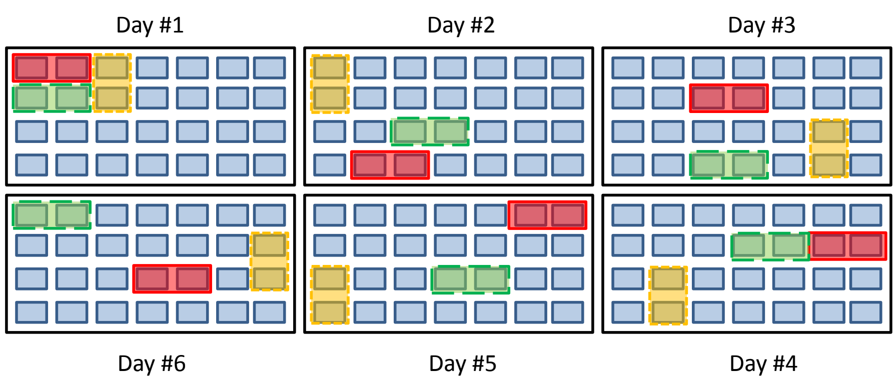

Notes: The two panels illustrate work assignments for tea workers at the Lujeri Tea Estates in Malawi. Panel A shows how three workers would be assigned two plots each. For our analysis, all workers A, B and C would be neighboring co-workers. Panel B shows how plot work assignments change over the course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

Figure 2: Distribution of Worker Ability

(a) Kernel Density of Ability

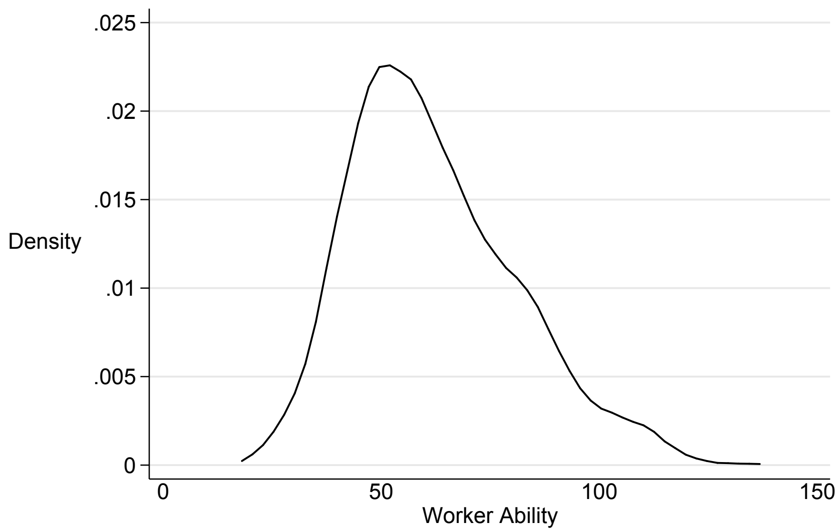

(b) Kernel Density of log(Ability)

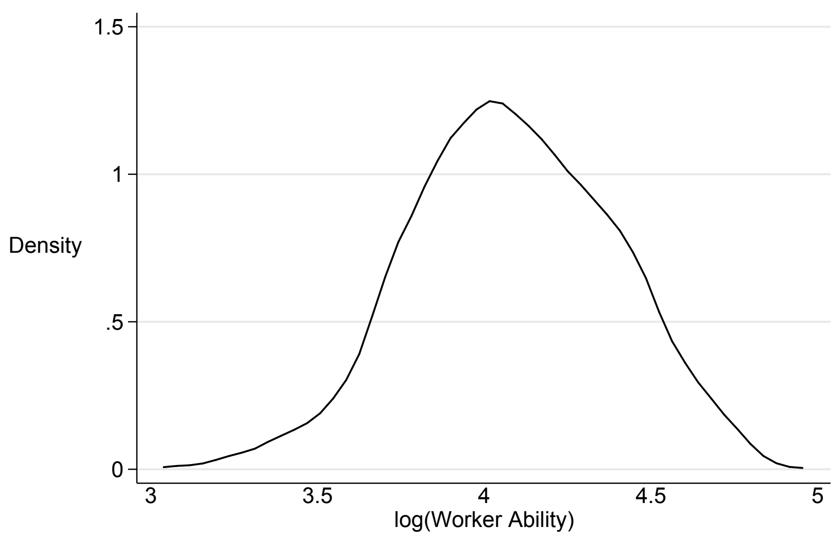

Notes: Figures present the density of estimated peer ability for a sample of tea pluckers at the Lujeri Tea Estates in Malawi. See Section 5 for details on how we construct these estimates.

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

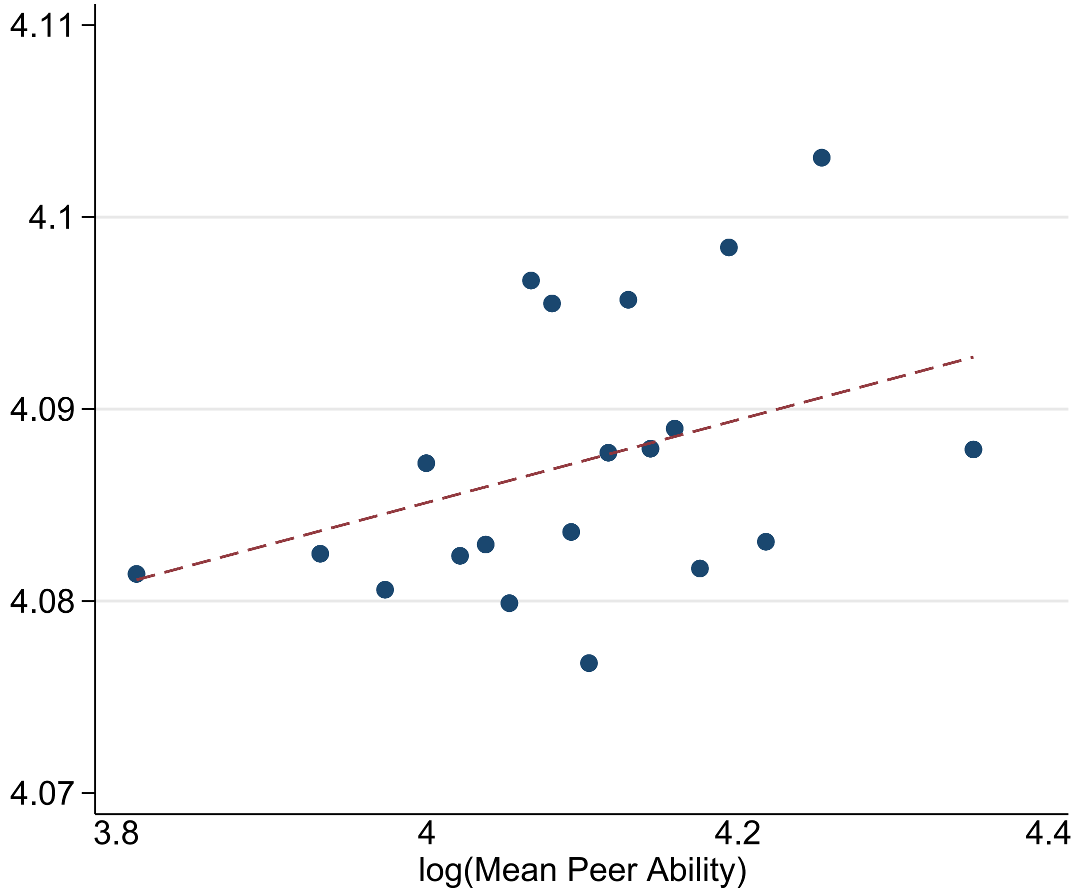

Notes: This figure plots (log) worker output (y-axis) after controlling for worker and date-by-location fixed effects by bins of mean (log) peer ability. There are 20 bins based on the ventiles of the mean (log) peer ability.

Figure 4: Distribution of Worker Ability by Gender

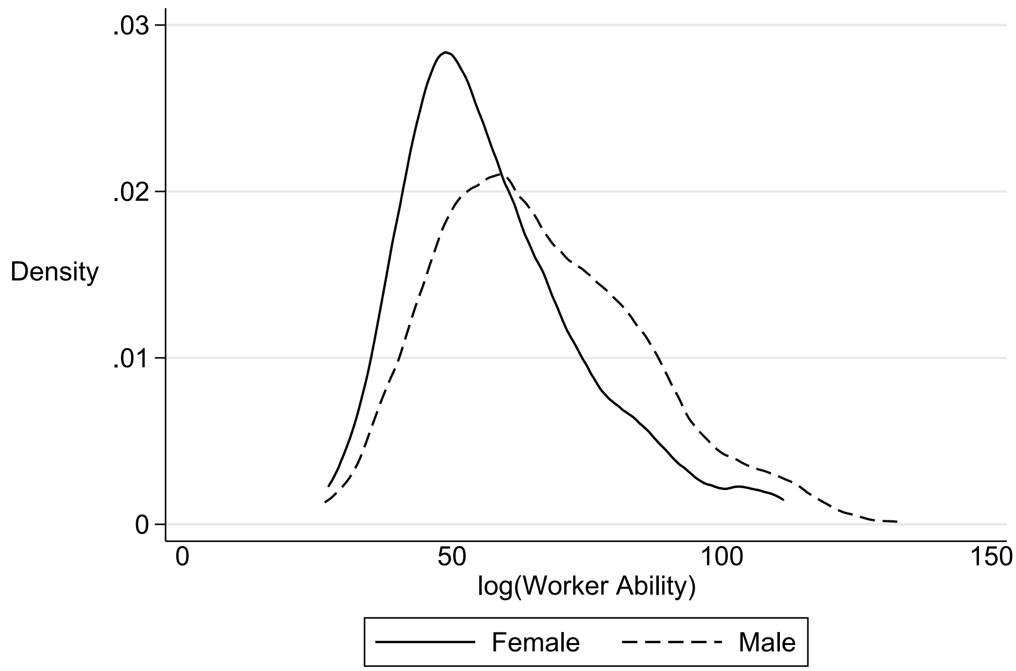

Notes: This figure presents density plots of estimated peer ability for male and female workers at Lujeri Tea Estates in Malawi. See Section 5 for details on how we construct these estimates.

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

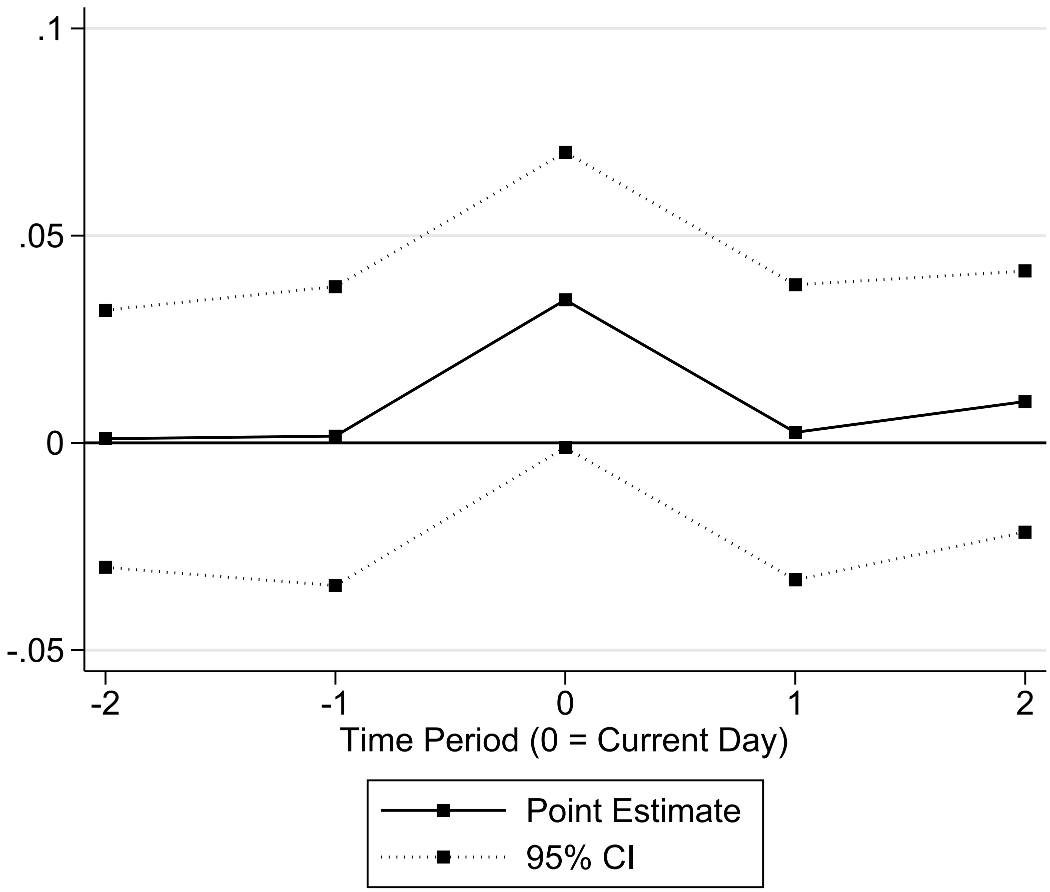

Notes: This figure plots the estimated elasticity of own output (y-axis) with respect to mean peer ability from a model which includes contemporaneous peer ability (t=0) as well as two leads (t=1, t=2) and lags (t=-1, t=-2).

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 $N = 55$ individuals.

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 Fixed Effects	No	No	No	Yes
Observations	9,313	44,858	35,545	35,449
Adjusted R-squared	0.246	0.233	0.230	0.397

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

Table 3: Effects of Workplace Peers, Linear Model

	Dependent Variable: Log of Daily Output	
	(1)	(2)
Log(Mean Peer Ability)	0.028**	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 daily output (kilograms of tea plucked) on the mean ability of physically nearby co-workers. The underlying data is a panel at the worker and day level. The results in Columns (1) and (2) use two different approaches to control for date and location effects. The measure of daily output comes from administrative data obtained from Lujeri Tea Estates; information on neighbors was recorded and collected staff for this project Standard errors are two-way clustered at the level of a worker-by-cycle-day and date-by-field and adjusted using the Bayesian parametric bootstrap of Mas and Moretti (2009).

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

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

	Dependent Variable: Log of Daily Output				
	(1)	(2)	(3)	(4)	(5)
Log(Mean Peer Ability), Non-Friends	0.028**		0.028**		
	(0.014)		(0.014)		
Any Non-friends (=1)	-0.068		-0.072		
	(0.064)		(0.064)		
Log(Mean Peer Ability), Friends		0.006	0.006		
		(0.013)	(0.014)		
Any 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 daily output (kilograms of tea plucked) on measures of the mean ability of nearby co-workers who are friends and non-friends. The underlying data is a panel at the worker and day level. The measure of daily output comes from administrative data obtained from Lujeri Tea Estates; information on neighbors was recorded and collected by staff for this project. Standard errors are two-way clustered at the level of a worker-by-cycle-day and date-by-field and adjusted using the Bayesian parametric bootstrap of Mas and Moretti (2009).

Table 6: Peer Effects by Experience Level

	Dependent Variable: Log of Daily Output		
	(1)	(2)	(3)
Log(Mean Peer Ability)	0.030**	0.032	
	(0.014)	(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 daily output (kilograms of tea plucked) on the mean ability of physically nearby co-workers, broken down by workers' experience at the firm. The underlying data is a panel at the worker and day level. The results in Column (2) and (3) are from specifications that include additional interactions based on the worker's self-reported experience at Lujeri. The measure of daily output comes from administrative data obtained from Lujeri Tea Estates; information on neighbors was recorded and collected staff for this project. Standard errors are two-way clustered at the level of a worker-by-cycle-day and date-by-field and adjusted using the Bayesian parametric bootstrap of Mas and Moretti (2009).

Table 7: Effects of Previous Days' Peers

	Dependent Variable: Log of Daily Output					
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Mean Peer Ability)	0.030**	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 daily output (kilograms of tea plucked) on the mean ability of physically nearby co-workers, including lagged as well as contemporaneous peer ability. The underlying data is a panel at the worker and day level. The measure of daily output comes from administrative data obtained from Lujeri Tea Estates; information on neighbors was recorded and collected staff for this project. Standard errors are two-way clustered at the level of a worker-by-cycle-day and date-by-field and adjusted using the Bayesian parametric bootstrap of Mas and Moretti (2009).

Table 8: Preferences for Fast Peers

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

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

Supplemental Appendix for Online Publication

A1 Appendix Figures and Tables

Figure A1: Photograph of a Tea Field at Lujeri Tea Estates

Notes: This photograph shows a tea field at Lujeri Tea Estates. The photograph was taken by the authors in November 2014.

Figure A2: Illustration of Peer Pressure Function Cases

(a) High-Ability Peers Increase Utility

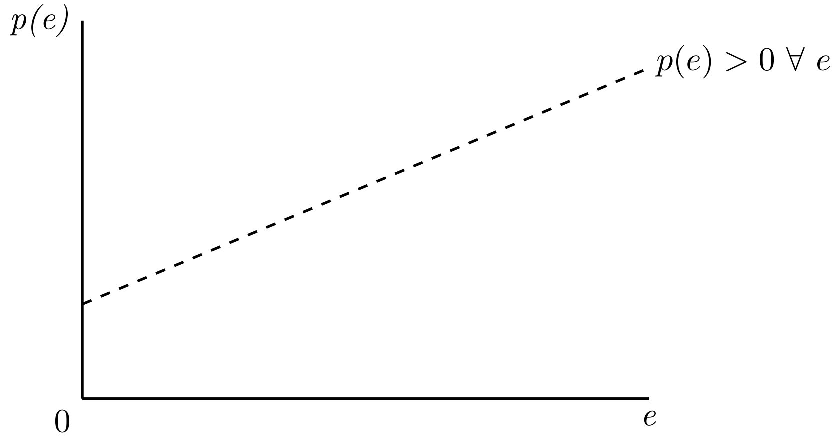

(b) High-Ability Peers Reduce Utility

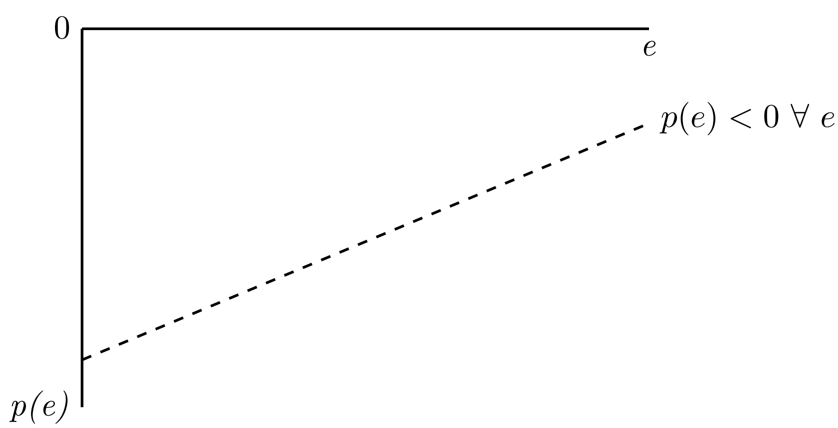

Notes: This figure illustrates two possible cases for the peer pressure function $p(.)$ for the stylized model of utility presented in Section 7. The x-axis shows effort (e). Panel A shows shows the case where having high ability peers increases total utility. This is consistent with the idea that high-ability peers provide motivation. Panel B shows the case where having high ability peers reduces utility. This is consistent with the idea that high-ability peers reduce utility due to shame or rank preferences.

Table A1: Regression of Worker Ability on Worker Attributes

	Dependent Variable: Worker Ability	
	(1)	(2)
Female	-5.676***	-6.088***
	(2.062)	(2.080)
Married	2.127	2.214
	(2.096)	(2.116)
Household Size	0.852	0.701
	(0.519)	(0.524)
Household Spending per Capita	0.000	0.000
	(0.000)	(0.000)
Age	-0.015	
	(0.068)	
Quartiles of Age		
Age 30 to 35		2.113
		(1.647)
Age 36 to 44		2.049
		(1.801)
Age 44 to 72		1.766
		(1.877)
Experience	0.272***	
	(0.090)	
Quartiles of Experience		
2.1 to 5 Years		5.694***
		(1.671)
5.1 to 10.7 Years		6.258***
		(1.645)
10.8 to 49.5 Years		7.742***
		(1.862)
Worker Fixed Effects	Yes	Yes
Date by Location Fixed Effects	Yes	Yes
Observations	909	909
Adjusted R-squared	0.062	0.075

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

Table A2: Additional Balance Tests: Comparing Worker Characteristics to Peer Ability

	Dependent Variable						
	Age	*Female*	*Married*	*Experience*	*Household Size*	*New Worker*	*HH Spending per Capita*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log(Mean Peer Ability)	-0.257	0.091	-0.097	0.495	-0.044	-0.017	62.794
	(1.33)	(0.062)	(0.063)	(1.07)	(0.164)	(0.031)	(351)
Log(Leave-One-Out Gang Mean Ability)	-3.226	5.742***	-4.939***	-55.027***	-6.880***	0.865***	1,665.492
	(13.1)	(0.787)	(0.773)	(11.7)	(1.60)	(0.277)	(3521)
Cycle Day 1	No	No	No	No	No	No	No
Remaining Cycle Days	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Worker Fixed Effects	No	No	No	No	No	No	No
Date by Location Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32,914	32,914	32,914	32,914	32,914	32,914	32,914
Adjusted R-squared	0.035	0.026	0.010	0.060	-0.006	0.067	0.003

Notes: This table presents results from regressions of a worker's characteristics on the mean ability of physically nearby co-workers. The underlying data is a panel at the worker and day level. 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 A3: Peer Effect Estimates Controlling for Other Peer Characteristics

	Dependent Variable: Log of Daily Output								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Log(Mean Peer Ability)	0.030**	0.030**	0.037**	0.043***	0.040**	0.037**	0.034**	0.030**	0.040**
	(0.014)	(0.014)	(0.015)	(0.017)	(0.017)	(0.016)	(0.016)	(0.014)	(0.017)
Number of Neighbors		0.000							-0.001
		(0.001)							(0.001)
Mean Peer Age			0.000						0.000
			(0.000)						(0.000)
Share of Peers who are Female				0.009					0.016
				(0.007)					(0.013)
Share of Peers who are Married					-0.004				0.006
					(0.008)				(0.013)
Mean Peer Household Size						0.001			0.001
						(0.003)			(0.003)
Mean Peer Experience							0.001		0.001
							(0.000)		(0.000)
Share of Same-Gender Peers								-0.002	-0.001
								(0.007)	(0.007)
Worker Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Date by Location Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	35,545	35,545	34,301	34,301	34,301	34,301	34,301	35,517	34,301
Adjusted R-squared	0.715	0.715	0.721	0.721	0.721	0.721	0.721	0.715	0.721

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

Table A4: Effects of Direct and Second-Order Peers

	Dependent Variable: Log of Daily Output		
	(1)	(2)	(3)
Log(Mean Peer Ability)	0.030**	0.036***	0.037***
	(0.014)	(0.014)	(0.014)
Log(Mean Ability of Strictly 2nd-Order Peers)			0.000
			(0.019)
Worker Fixed Effects	Yes	Yes	Yes
Date by Location Fixed Effects	Yes	Yes	Yes
Observations	35,545	35,291	35,291
Adjusted R-squared	0.715	0.716	0.715

Notes: This table presents results from a regression of daily output (kilograms of tea plucked) on the mean ability of physically nearby co-workers. The underlying data is a panel at the worker and day level. The results in Columns (1) replicate our preferred specification from Table 3. The results in Column (2) are based on the restricted sample of observations where information on second-order neighbors is available. Second-order neighbors are defined as the co-workers adjacent to a focal worker's neighbors who are *not* directly adjacent to the focal worker. Columns (3) and (4) estimate versions of Equation 1 that include measures of ability for second-order neighbors. Standard errors are two-way clustered at the level of a worker-by-cycle-day and date-by-field.

Table A5: Effects of Assigned Peers on Attendance and Tea Plucking

	Dependent Variable: Attendance	*Dependent Variable: Tea Plucking*
	(1)	(2)
Log(Mean Peer Ability)	0.003	0.004
	(0.012)	(0.013)
Worker Fixed Effects	Yes	Yes
Date-by-Location Fixed Effects	Yes	Yes
Observations	47,959	47,959
Adjusted R-squared	0.133	0.234

Notes: This table presents results from a regressions of an indicator for the worker being present at work (column 1) or being engaged in tea plucking (column 2) on the mean ability of the physically nearby co-workers for their assigned field for the day. The underlying data is a panel at the worker and day level. The measure of daily attendance and plucking comes from administrative data obtained from Lujeri Tea Estates; information on neighbors was recorded and collected by staff for this project. Standard errors are two-way clustered at the level of a worker-by-cycle-day and date-by-field and adjusted using the Bayesian parametric bootstrap of Mas and Moretti (2009).

Table A6: Effects of Workplace Peers without Double Leave-One-Out Correction

	Dependent Variable: Log of Daily Output	
	(1)	(2)
Log(Mean Peer Ability)	0.053***	0.043***
	(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 daily output (kilograms of tea plucked) on the mean ability of physically nearby co-workers. The underlying data is a panel at the worker and day level. The results in Columns (1) and (2) use two different approaches to control for date and location effects. The measure of daily output comes from administrative data obtained from Lujeri Tea Estates; information on neighbors was recorded and collected by staff for this project Standard errors are two-way clustered at the level of a worker-by-cycle-day and date-by-field and adjusted using the Bayesian parametric bootstrap of Mas and Moretti (2009).

Table A7: Male and Female Peer Effects, Robustness

	Dependent Variable: Log of Daily Output	
	(1)	(2)
Male X Log(Mean Peer Ability)	0.006	0.020
	(0.020)	(0.041)
Female X Log(Mean Peer Ability)	0.060***	0.083*
	(0.023)	(0.046)
Quartiles of Age		
[Age 20 to 29] X [Log(Mean Peer Ability)]		0.031
		(0.045)
[Age 30 to 35] X [Log(Mean Peer Ability)]		-0.014
		(0.045)
[Age 36 to 44] X [Log(Mean Peer Ability)]		-0.015
		(0.042)
Quartiles of Own Ability		
[Own Ability Quartile 1] X [Log(Mean Peer Ability)]		-0.063
		(0.045)
[Own Ability Quartile 2] X [Log(Mean Peer Ability)]		-0.015
		(0.041)
[Own Ability Quartile 3] X [Log(Mean Peer Ability)]		0.004
		(0.041)
Worker Fixed Effects	Yes	Yes
Date by Location Fixed Effects	Yes	Yes
Observations	33,010	33,010
Adjusted R-squared	0.725	0.725
Male-Female Treatment Effect Difference	0.054*	0.063*
	(0.030)	(0.037)

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

Table A8: Male and Female Peer Effects, Interaction Model Results

	Dependent Variable: Log of Daily Output	
	(1)	(2)
Log(Mean Ability of Female Peers)	0.024***	
	(0.008)	
Log(Mean Ability of Male Peers)	0.023***	
	(0.008)	
Male X Log(Mean Ability of Female Peers)		0.013
		(0.010)
Female X Log(Mean Ability of Female Peers)		0.037***
		(0.012)
Male X Log(Mean Ability of Male Peers)		0.012
		(0.011)
Female X Log(Mean Ability of Male Peers)		0.036***
		(0.012)
Worker Fixed Effects	Yes	Yes
Date by Location Fixed Effects	Yes	Yes
Observations	26,911	26,394
Adjusted R-squared	0.728	0.728

Notes: This table presents results from a regression of daily output (kilograms of tea plucked) on the mean ability of physically nearby co-workers, interacted with own and peer worker characteristics. The underlying data is a panel at the worker and day level. The measure of daily output comes from administrative data obtained from Lujeri Tea Estates; information on neighbors was recorded and collected by staff for this project. Sample is restricted to observations where gender is non-missing for all neighbors. Standard errors are two-way clustered at the level of a worker-by-cycle-day and date-by-field.

Table A9: Summary Statistics for Worker Ability by Gender

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

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

Table A10: Test for Asymmetry in Peer Effects

	Dependent Variable: Log of Daily Output	
	(1)	(2)
Log(Mean Peer Ability)	0.030**	
	(0.014)	
Log(|Peer Ability Difference|)X(Below Peer Ability)		0.017***
		(0.004)
Log(|Peer Ability Difference|)X(Above Peer Ability)		-0.018***
		(0.005)
Below Peer Ability		0.008
		(0.016)
Worker Fixed Effects	Yes	Yes
Date Fixed Effects	No	No
Location (Field) Fixed Effects	No	No
Date by Location Fixed Effects	Yes	Yes
Observations	35,545	35,449
Adjusted R-squared	0.715	0.717

Notes: This table presents results from a regression of daily output (kilograms of tea plucked) on the mean ability of physically nearby co-workers. The underlying data is a panel at the worker and day level. The results in Columns (1) replicate our preferred specification from Table 3. Column (2) shows a test for asymmetry in the peer effects, following the analysis of the impact of friends in Bandiera et al. (2010). We compute the absolute value of the difference between the worker's own ability and the mean of their co-workers. We interact the log of this measure with an indicator for whether the worker's ability level is below the mean of their co-workers. The measure of daily output comes from administrative data obtained from Lujeri Tea Estates; information on neighbors was recorded and collected by staff for this project. Standard errors are two-way clustered at the level of a worker-by-cycle-day and date-by-field and adjusted using the Bayesian parametric bootstrap of Mas and Moretti (2009).

Table A11: Summary Statistics for Mean Ability of Peers who are Friends and Non-Friends

	(1) Average	(2) Std. Deviation	(3) 10th Percentile	(4) 25th Percentile	(4) 50th Percentile	(4) 75th Percentile	(4) 90th Percentile	(5) Obs (N)
Non-Friends	61.28	13.03	46.89	51.75	58.87	69.38	79.23	35,456
Friends	64.23	18.85	43.18	50.18	61.59	75.47	90.25	7,409

Notes: This table presents descriptive statistics for mean peer ability among workers' friends and non-friends. The ability measure is estimated using Equation 2.

B Bayesian Parametric Bootstrap

To account for the fact that neighbor (ability) types are estimated, we construct all standard errors in the paper using the Bayesian parametric bootstrap of Mas and Moretti (2009). This procedure consists of four steps. First, we draw simulated ability types (i.e., fixed effects for productivity) for each worker from a joint normal distribution that has a vector of means and a variance-covariance matrix that are equal to the results from our type estimation procedure. Second, for each draw of the worker types, we re-run the regressions in our analysis using the draws from the simulation to construct our peer ability measure. Third, we estimate σ_{sd} for a given regression as the standard deviation of the point estimates across draws. Fourth, we combine the across-simulation standard error with the typical clustered standard error. Let SE_{clust} be the usual cluster-adjusted standard error. The Bayseian standard is then equal to $SE_{bayes} = \sqrt{\widehat{\sigma_{sd}}^2 + SE_{clust}^2}$. We follow Mas and Moretti in using 10 draws for our analysis.

C Double Leave-One-Out Approach

This section provides further details on our approach to addressing spatial spillovers, which may bias estimates of peer effects. Consider worker i who is on a high-productivity part of a field f, and let k index her K plot neighbors on that day. Suppose we estimate the following model of peer effects:

$$y_{ift} = \mu_i + \beta \overline{Ability}_{-ift} + \delta_{tf} + \epsilon_{ift}. \tag{C1}$$

A potential problem occurs because the higher-productivity part of field f increases the error term ϵ_{ift}. We can control for field fixed effects in the model through the field-by-date fixed effects δ_{tf}. However, this does *not* control for variations in plot quality *within* a field.

In this case, the issue is that being on a high-quality part of a field f will increase the average value of output for worker i and all of worker k's neighbors (i.e., both y_{ift} and y_{kft} go up). This becomes a problem if we attempted to estimate ability for each worker using the specification:

$$y_{ift} = \mu_i + \mathbf{M}_{ift}\gamma' + \delta_{tf} + \tau_{ift} \tag{C2}$$

where the term $\mathbf{M}_{ift}$ is a vector of dummy variables which indicate whether worker j in a gang is working next to worker i in field f on date t. The issue in Equation C2 is that spatial spillovers will be absorbed in all of the neighbor k fixed effects (μ_k). This is problematic if one constructs the mean ability of the worker i's neighbors as $\overline{Ability}_{-ift} = \overline{\widehat{\mu}}_{-ift} = \frac{1}{K}\sum_{k=1}^{K} \widehat{\mu}_k$. With this definition, spatial spillovers would generate a correlation between $\overline{Ability}_{-ift}$ and the error term (ϵ_{ift}) in Equation C1.

Our approach eliminates this correlation by estimating each of the $\widehat{\mu}_k$ for field f using a dataset that excludes field f. Specifically, we estimate Equation 2, reproduced below:

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

where g indexes all fields except field f. The term μ_{i-f} is the ability measure for worker i using all fields except field f. This allows us to construct a double leave-one-out estimate of a worker's mean peer ability, $\overline{Ability}_{-i-ft} = \overline{\widehat{\mu}}_{-i-ft} = \frac{1}{K}\sum_{k=1}^{K}\widehat{\mu}_{k-f}$. This version of mean peer ability ($\overline{Ability}_{-i-ft}$) is preferable because it is constructed without using any data from field f. Hence, spatially-correlated variation in plot quality within a field f does not affect our estimates of permanent productivity (ability type) for worker's peers.

One drawback of using this approach is that it will tend to induce additional classical measurement error in our ability estimates — and hence attenuate our estimated coefficients — because the ability measure is estimated using a smaller sample. In a situation with no spatially-correlated shocks, then, the approach that uses all the data to estimate ability is less biased.

To estimate the amount of additional measurement error induced by our approach, we first note that the average worker in our analytic sample has 5.61 cycle days of data. The double leave-one-out estimator drops one cycle day, and so uses 4.61 cycle days instead. This has the effect of using, on average, 1/5.61 fewer total observations to estimate the coefficients in the double leave-one-out approach. Since the standard error is proportional to $1/\sqrt{N}$, this inflates the standard error of the ability estimate by $\sqrt{5.61/4.61} = 1.103$. Thus, the variance of the ability estimate rises by 1.216. Under classical measurement error, the attenuation bias factor λ is $Var_x/(Var_x + Var_u)$, where x is the regressor of interest and the noise term u is additional variation in the measured value of x that arises as a result of using less data. The prior calculation shows that the double leave-one-out estimate has variance $Var_{x'} = 1.216 \cdot Var_x = Var_x + 0.216 \cdot Var_x$. In other words, the additional measurement error induced by the double leave-one-out approach implies that $Var_u = 0.216 \cdot Var_x$. Hence, the attenuation factor is $\lambda = Var_x/(Var_x + 0.216 \cdot Var_x) = 1/1.216$. This means that our estimates are attenuated by a factor of 0.822, and we should see results that are attenuated by an additional 18 percent.

The estimated attenuation factor is a lower bound, and thus the additional attenuation will be less than 18 percent, for two reasons. First, the ability estimates are correlated across days for a given worker, and so losing any given day is less costly in terms of mismeasurement of ability. Second, our preferred specification averages multiple independent worker ability estimates, which means that the calculations above overestimate the attenuation bias problem.[58] We study the attenuation bias issue numerically in Appendix D and find that our double leave-one-out estimates are attenuated by about 10 percent relative to the true parameter value.

[58] The above calculation also ignores the fact that the number of cycle days varies somewhat by worker, but this should not substantively affect our results.

D Monte Carlo Simulations

This section uses Monte Carlo simulations to assess the importance of two features of the empirical strategy described in Section 5. First, we show the sensitivity of the estimated of impact of mean peer ability using different levels of fixed effects. The results show that our preferred estimates, based on including worker and field-by-date fixed effects, perform well relative to other choices of fixed effects. Second, we assess how the estimated impact of mean peer ability depends on the use of the double leave-one-out approach to estimating worker own ability. For this exercise, we use simulations that assume the existence of spatially-correlated (time-invariant) shocks to plot quality. The results show that peer effect estimates based on the double leave-one-out estimator are much closer to true parameter value relative to estimates based on an approach that estimates worker ability using all daily data.

D.1 Field-by-Date Fixed Effects

Our first set of simulations is based on the following data generating process for the log of daily output (measured in kilograms) y_{ift}:

$$y_{ift} = \mu_i + 0.1\overline{Ability}_{-ift} + \delta_{tf} + \epsilon_{ift}. \tag{D1}$$

We build the simulated data to mirror the following features of the observed sample of workers:

1. The simulated data has the same number of observations as our experimental data.

2. The date t and field f fixed effect δ_{tf} is estimated from the real data on productivity for a given field and date via our main regression specification (Equation 1).

3. Worker ability levels μ_i are drawn randomly, with replacement, from the estimated (log) ability levels for the entire gang that works on a given field.

4. Peers are determined by randomly assigning the simulated workers to positions in the real adjacency matrix for the field, which then determines log mean peer ability $\overline{Ability}_{-ift}$.

5. The error term ϵ_{ift} is generated by randomly drawing a residual, with replacement, from the estimated distribution of residuals obtained from the main regression specification (i.e., Equation 1).

Under these conditions and the model in Equation D1, we create 150 simulated samples by re-generating all variables in each simulation iteration.

Using the simulated data, we use the double leave-one-out approach (as in the main text) to estimate the impact of mean peer ability using different sets of fixed effects. Appendix Table D1 (on page 63) shows that estimates are much closer to the true parameter value when the peer effect regressions use worker and field-by-date fixed effects. The first row reports the true value of the peer effect parameter (which is constant across simulations). As expected, the second row shows that peer effect estimates are severely upward-biased when no fixed effects are included. The estimated coefficient is 0.55 — five times the true value. This upward bias occurs because our experiment randomized workers to co-workers only within a given field. Fields vary systematically in terms of productivity, which will increase both daily output y_{ift} and estimated co-worker ability $\overline{\hat{Ability}}_{-ift}$, leading to positive omitted variable bias. The third row reports the estimates from a model which only includes field-by-date fixed effects. There is evidence of exclusion bias due to the lack of worker fixed effects: workers cannot be their own co-workers, creating a mechanical negative bias in our point estimates (Guryan et al., 2009; Caeyers and Fafchamps, 2016). Finally, the fourth row shows that our preferred specification, which includes both worker and field-by-date fixed effects, produces estimates of the impact of mean peer ability that are close to the true parameter value. As discussed in Appendix C, the estimates are slightly attenuated due to classical measurement error.

D.2 The Double Leave-one-out Approach

Our second set of simulations assume the existence of spatially correlated shocks to plot quality. Specifically, this set of simulations is based on the following data generating process:

$$y_{ift} = \mu_i + 0.1\overline{Ability}_{-ift} + \delta_{tf} + \sigma Shock_{if} + \sigma\rho Peers\text{–}Shocked_{if} + \epsilon_{ift}. \tag{D2}$$

The term $Shock_{if}$ is an indicator for worker i being shocked on field f; these shocks are permanent attributes of a specific plot, so each time a worker returns to a given field her shock status is the same. The term $Peers\text{–}Shocked_{if}$ is the number of peers near worker i that receive a shock. In this model, the idiosyncratic shocks to plot quality boost output by σ log points. The shocks can spill over onto neighboring plots as governed by the spatial correlation factor ρ.

Based on this model, we create simulated data where we always assume that five percent of all plots receive a shock.[59] Across simulations, we vary two key parameters: the shock intensity, σ, which we allow to be 0, 0.25, or 0.5; and the spatial correlation factor, ρ, which we allow to be 0, 0.5, or 1. We simulate each combination of parameter values 50 times, randomly re-generating all variables for each simulation (holding the peer effect coefficient and shock rate constant).

Using the simulated data sets, we compare estimates the impact of mean peer ability using the double leave-one-out approach to estimates based on an approach that uses all the data to estimate peer ability. All estimates of peer effects in this section are based on a regression which includes both worker and field-by-date fixed effects. Panel A of Appendix Figure D1 (on page 64) shows estimates of the impact of mean peer ability while the shock intensity varies and the spatial correlation factor is held constant at 0.5. The bias of the estimator that uses all the data to estimate ability is an increasing function of the shock intensity. In

[59] We build the simulated data using all of the conditions listed in the prior section (e.g., the simulated sample has the same number of observations as the real data).

contrast, the double leave-one-out estimator performs well even when the shocks are larger. Next, Panel B of Appendix Figure D1 shows estimates that vary the spatial correlation factor and hold the shock intensity fixed at 0.25. Again, the double leave-one-out estimator has low bias for all values of the spatial correlation factor. In contrast, as the spatial correlation factor increases, the bias increases when using estimates of ability based on all the data.

Table D1: Monte Carlo Comparison of Peer Effects Estimates with Different Levels of FEs

	(1) Average	(2) Std. Deviation	(3) Min	(4) Max	(5) Obs (N)
True Peer Effect Coefficient	0.10	0.00	0.10	0.10	150
No Field-by-Date FEs, no Worker FEs	0.55	0.03	0.47	0.62	150
Field-by-Date FEs, no Worker FEs	0.01	0.04	-0.09	0.13	150
Preferred Estimates (Field-by-Date FEs and Worker FEs)	0.09	0.01	0.06	0.13	150

Notes: This table presents the distribution of the estimated coefficient on (log) mean peer ability across 150 Monte Carlo simulations of the data generating process described in Section D, with no spatially-correlated shocks. The first row shows the true parameter value within the simulation. Row 2 shows the estimates without any fixed effects in the model. Row 3 shows the estimates with field-by-date FEs but no worker FEs, and row 4 shows our preferred specification, which includes both field-by-date FEs and worker FEs. All specifications use the double leave-one-out estimator, which drops the cycle day in question from the sample when estimating peer ability measure for a given day.

Figure D1: Monte Carlo Simulation of the Performance of Different Peer Effects Estimators

(a) Varying Degree of Shock Intensity

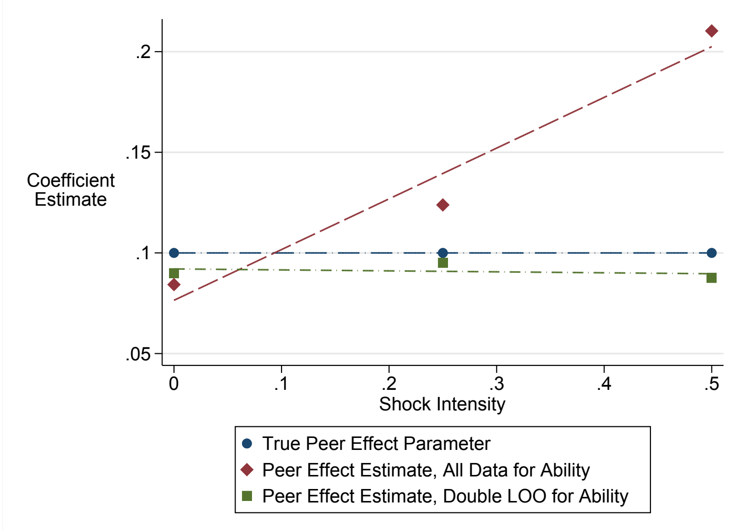

(b) Varying Degree of Spatial Correlation

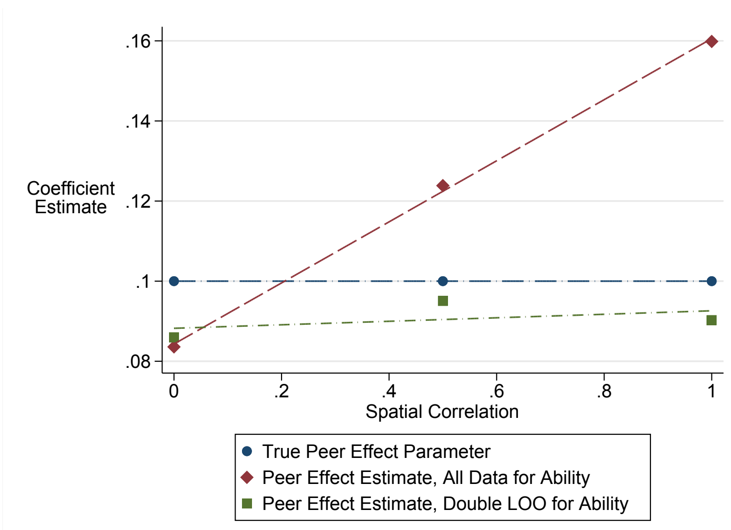

Notes: This figure presents estimates the impact of mean peer ability from different simulations of data. The "All Data" estimates (diamond symbol) are based on the simulation data and estimating worker ability using the entire sample. The "Double LOO" estimates (square symbol) are based on the simulation data and estimates of worker ability are based on a double leave-one-out approach where one cycle day is dropped. Panels A and B show points and regression lines where the underlying model for simulations varies by the shock intensity and spatial correlation parameters. See the text in Appendix D for further details.