

Organizational Learning: Experimental Evidence from Bangladeshi Garment Factories

Job Market Paper 2015/16

Andreas Menzel*[†]
University of Warwick

October 2015

Abstract

Organizational learning, or the sharing of knowledge among co-workers in firms, has long been assumed to be a key driver of productivity growth. However, because knowledge exchange is inherently difficult to observe, identifying the effect of knowledge sharing on productivity has remained problematic. The literature has often measured knowledge exchange in firms through increases in productivity of workers if other workers in the same firm have already produced the same product. However, this approach risks confounding the effect of knowledge exchange with other peer effects, such as competition. This paper provides evidence from a communication intervention in three Bangladeshi garment factories in which randomly selected workers were instructed by their superiors to share production knowledge when one worker started producing a garment that the other had already produced. The intervention increased the productivity of the later workers producing the garment by 0.2 standard deviations during the first one to two days they produce the garment, before their productivity reached its long-run level again. There is some evidence that the effect was stronger if the workers sharing knowledge were socially connected. A back-of-the-envelope calculation indicates that the return on this low-cost intervention is in excess of 600 percent. Furthermore, compliance by the factories in implementing the intervention was higher if the later workers that should receive knowledge were younger. This indicates that workers' status concerns could interfere with the implementation of such a management routine, and could explain why the routine had not been implemented earlier by the factories.

Keywords: Learning, Productivity, Firms

JEL Code: D2, L2, M5, O3

*I am extremely grateful to my thesis supervisors Chris Woodruff and Rocco Macchiavello for their continuous support, guidance and advice throughout this project. I furthermore would like to thank Eric Verhoogen, Emily Breza, Jonas Hjort, Francois Gerard, Suresh Naidu, Cristian Pop-Eleches, Bernard Salanie, Oriana Bandiera, Luigi Guiso, Petra Moser, Florian Englmaier, Robert Akerlof, Clement de Chaisemartin, Sascha O. Becker, Roland Rathelot, Andrew Oswald and participants at the Columbia Development Colloquium, RCEA Workshop Rimini, Oxford Young Economist Workshop 2014, PEDL Workshop Warwick, Italian Development Association Summer School Ascea 2014, Belgrade Young Economist Conference 2014, LMU Workshop Ohlstadt 2014, and Warwick CAGE seminar for many useful comments. Financial help by PEDL is acknowledged.

[†]Address: S2.110 - Department of Economics, University of Warwick, CV4 7AL, United Kingdom. Email: A.Menzel@warwick.ac.uk

1 Introduction

Learning on the job has long been known as a key driver of productivity growth (Arrow (1962); Lucas (1993)). Conceptually, learning on the job can be separated into two aspects, learning by doing something oneself, and learning from co-workers. While the evidence on learning via one’s own experience is extensive and goes well back into time (Wright (1936); Benkard (2000); Hendel and Spiegel (2014); see Thompson (2010) for a review), only limited evidence addresses learning from co-workers. Levitt et al. (2013) in a U.S. car plant, and Thompson and Thornton (2001) in U.S. ship yards show that productivity of workers not only increases with the amount of a certain car or ship model they produced themselves, but also with the amount of the product produced by others in the firm. However, beyond the mere documentation of the existence of learning from others in firms, little is known, including, for example, under which conditions it works particularly well.¹

One reason for the paucity of evidence on learning from co-workers is that knowledge exchange between people is inherently difficult to observe. The above-mentioned studies measure organizational learning through increases in productivity when other workers in the organization have previously produced the same product. However, this leaves open the possibility that the productivity increases are not driven by knowledge transfers, but alternative forms of peer effects.² To clarify the mechanism behind the productivity spill-overs within firms, I collect data and run a randomized experiment in three Bangladeshi garment factories. In the experiment, random pairs of workers are induced by their superiors to brief each other when one of them starts producing a new garment-style that the other one has previously produced (from now on, with “style”, I will refer to differentiated garments on the most disaggregated level possible: the individual garment models ordered by the different buyers). This communication intervention introduces exogenous variation in knowledge exchange across worker pairs, and I show that this intervention increases the productivity of the workers receiving the briefing. This indicates that increases in knowledge exchange among co-workers increase productivity above the levels achieved when workers learn solely from doing a certain task by themselves. Furthermore, using self-reported social ties from surveys I conducted among the workers, I find some evidence that the effect was stronger if the worker who received the briefing was socially connected to the worker who provided the briefing.

¹Kellogg (2011) studies learning within relationships between oil companies and their drilling contractors. He shows that when these relationships break, both the oil company and the drilling contractor become less productive in future drilling operations, indicating that employees accumulate knowledge not only about production processes, but also about working effectively in specific relationships.

²This problem also holds for most studies on social learning outside organizations, for example about new agricultural technologies (Foster and Rosenzweig (1995); Munshi (2004); Bandiera and Rasul (2006); Conley and Udry (2010)), or about microfinance services (Banerjee et al. (2013); Cai et al. (2015)), which use observed technology adoption as proxy for learning. These studies mostly rely on estimating structural models or on placebo tests to separate learning from other possible peer effects. In a recent contribution, Covert (2015) uses mandatory production information, which oil fracking companies have to publish publicly six months after starting a new well, to measure knowledge on optimal input choices which is publicly available in the industry. He finds that firms do use this information, even though they put more weight on the experience gained in their own operations.

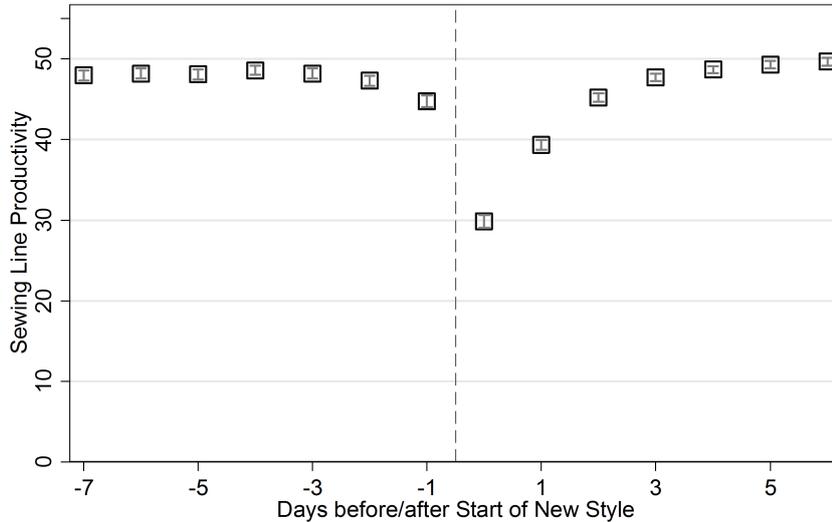


Figure 1: Sewing Line Productivity before and after Start of a New Style. Graph shows average sewing line productivity in the days before and after switching to a new style. The vertical dashed line denotes the switch to the new style, and “Day 0” the first day of production of the new style. Capped bars represent 95% confidence intervals.

Bangladeshi garment factories are an ideal setting for studying organizational learning. The sewing departments of these factories, on which this paper focuses, are organized into parallel sewing lines of 20-80 workers. The lines are designed as independent production units on which the whole process of sewing a style can be completed. The three factories in my sample have more than 200 sewing lines among them. Due to large orders and tight delivery deadlines, most styles are produced on more than one line. In these cases, the different lines typically start producing the style on different days, as they finish previously allocated jobs. Thus, when a given sewing line starts producing a new style, there may or may not be other lines in the factory with production experience on that particular style. When there are such experienced lines, they have gained potentially valuable production information which could be shared with the sewing line starting production at a later time. Sewing lines switch to new styles with high frequency, on average every 10 days. Due to the fast-moving fashion industry and its seasonality, styles are technically differentiated, which is reflected by line productivity dropping by a third on average when lines switch to a new style. Only four to five production days after the start of a new style does line productivity reach its long-run level again (Figure 1). Given these learning processes and fast turnover rates to new styles, the potential gain from knowledge spill-overs is large. Workers - or at least the supervisors of the lines who are held accountable for the productivity of their lines - have a strong incentive to utilize this knowledge.³

³Such incentives to exploit knowledge from co-workers to increase the own line’s productivity would be affected by how the pay of workers and supervisors is determined (Bandiera et al. (2005); Lazear (2000)). In all three factories in my sample, ordinary workers receive hourly pay, while line supervisors receive fixed monthly

The randomized communication intervention was conducted among the line chiefs, the people responsible for the overall management of lines. For a period of four months on randomly selected sewing floors in the factories, line chiefs were instructed by the factory management to brief each other when a style that had previously been produced on one line was being started on another line. The briefings were designed to last about 20 minutes, during which the line chief with experience on that particular style was to share the most important production problems that had to be overcome in its production. I show that these briefings increased productivity; lines supervised by line chiefs who received a briefing experienced productivity increase of 0.2 standard deviations on the first one to two days of production of the new style, before their productivity reached its long-run level again.

The randomized intervention covers only a fourth of the dataset of sewing line productivity that I collected. Using data from as long as two years before the experiment began, I first document that sewing lines are generally more productive on the first days they produce a new style if that style has previously been produced by other sewing lines. This effect is stronger if the lines that previously produced the style are located spatially closer. However, the effects I document in the observational data could also be driven by other forms of peer effects, such as competition. For example, I find that sewing lines are similarly more productive on the first days they make a new style that is completely new to their factory, and thus, no other workers have previous experiences that can be shared, if other lines start to produce the same style on the same day. Again, the effect is stronger if these lines are located spatially closer. In principle, these contemporaneous effects could be driven by the selection of styles produced on multiple lines on the first day - for example, those with especially close delivery dates. Nevertheless, the presence of these productivity spill-over effects among lines starting the style on the same day, which are less likely driven by knowledge transfers, raises the question of whether the productivity increases observed when other lines have prior experience producing the style might also be generated by forces other than learning. Therefore, by exogenously increasing the potential for knowledge exchange between workers, the randomized intervention provides valuable evidence about how much learning from co-workers contributes to these productivity spill-overs.⁴

Previous studies on organizational and social learning typically observed the diffusion process of a single product, or a small number of products; however, the dataset I collected allows

wages (irrespective of the hours of overtime worked). No workers in the factories receive piece rate payments or performance bonuses. However, line supervisors compete for promotions to higher management positions, such as floor managers, and the factories keep track of line productivities, and hold the supervisors accountable for these productivities. Furthermore, the flat pay for supervisors gives them a direct incentive to keep productivity high, as it reduces their overtime, which is not paid separately.

⁴This paper provides experimental evidence on productivity spill-overs by randomly varying knowledge exchange between workers producing the same product. Falk and Ichino (2006) randomly vary the number of workers in a room doing a simple task with little potential for knowledge exchange. They find that workers are more productive when more than one worker is in the room, which provides experimental evidence for the existence of peer effects driven by other forces than learning.

me to observe the production and knowledge diffusion processes of more than 1,000 different styles over a constant set of sewing lines. Previous studies usually relied on multiple observations of the same production unit producing the same product, relating variation in productivity within production units and within products to the amounts of the product already produced by the same and other production units. However, this poses at least two problems. First, prior output of the product on the same and on other units usually evolves in a highly correlated way over time, making the precise disentangling of the separate effects of prior output on the same and on other units econometrically difficult, due to collinearity. Second, capital and labor used in a production unit can evolve over time while the units produce a given product, in ways that most studies cannot properly observe and control for (Thompson and Thornton (2001); Irwin and Klenow (1994)). The dataset used in this study allows to address both problems. First, given the large number of different styles in the data, I can use a single observation of a production line producing a certain style. By using the first daily observation of the line producing the style, I hold previous output of the same style by the same line constant at zero, and can solely focus on variation in output produced on other lines. Second, as I employ only one observation per line and style, I do not need to control for within line-within style changes in labor and capital employed on the line. Finally, a further advantage of the setting is the availability of accurate daily productivity data based on physical output on the sewing line level, which have been standardized across different styles. Therefore, I do not need to compare output measures across heterogeneous products based on assumptions that can be difficult to test (Foster et al. (2008)).

A basic back-of-the-envelope calculation shows that the returns on the intervention are likely to be in excess of 600 percent. The positive effects of the intervention raise the question of why the factories have not previously implemented similar management practices. This question relates to the literature on management practices in large firms (Bloom et al. (2013); Bandiera et al. (2015)), which finds that firms, especially in developing countries, fail to adopt practices that should be universally beneficial. A post-experiment survey was conducted with the three factories' production managers, who supervised the implementation of the experiment. Two of the three managers reported that they had never previously thought of conducting such an experiment; the third had considered the experiment but had not expected it to yield enough benefits to be worthwhile. Thus, lack of awareness about the potential benefits of such an intervention might explain previous non-adoption, mirroring results from Bloom et al. (2013). Additional insights as to why such measures might not have previously been implemented emerge from the analysis of a specific form of non-compliance observed in the implementation of the randomized intervention. Briefings were more likely to be implemented by the factory management the younger the line chiefs were on the receiving end of the briefings. This result holds when controlling for average productivity of the line chiefs' lines, their productivity on the day the briefing should have occurred, their education, and their experience as line chiefs. Thus, the management's reluctance to send line chiefs to brief their older counterparts points towards

the presence of status concerns among the line chiefs. Such concerns can be expected to increase with age; older line chiefs may resist receiving instructions from other line chiefs, especially if they are younger or less-experienced. As a result, factory managements, anticipating such resistance based on status concerns, may generally refrain from routinely having line chiefs exchange information on styles they both produce - despite the potential positive effects of such measures.

Apart from the literature on organizational learning mentioned above, this paper relates to a small but growing literature on experiments within firms. Atkin et al. (2015) vary the pay-schemes for workers in Pakistani soccer ball factories, and show that this variation affects whether employees report truthfully about the benefits of new technologies in the production process. Bloom et al. (2013) provided in-depth management consulting to randomly selected textile factories in India, and showed large effects on productivity of the firms. They cite low competitive pressures as reason for the factories' previous non-adoption of these practices, as well as a lack of trust in managers from outside the owner family, who could introduce better management practices. Bandiera et al. (2013) introduce rank incentives and a tournament for production teams at a soft fruit producer and demonstrate that while the rank incentives decrease productivity, the tournament increases overall productivity. Bandiera et al. (2005) compare worker productivity under piece rate pay and relative payment at the same firm. They show that relative pay reduced worker productivity, but only in cases in which workers can monitor each other's effort. Bandiera et al. (2011) provide an overview over this literature.⁵

This paper proceeds as follows. The next section introduces more background information about the factories and describes in more detail the dataset collected, while section three presents the non-experimental results on productivity spill-overs, using the whole collected dataset. Section four provides more details on the experiment and shows its main effects. Section five presents results on the interplay of the intervention with social ties, while section six provides evidence on how non-compliance in the implementation of the intervention is correlated with line chief characteristics, providing clues why this management routine has not been implemented earlier despite its positive estimated effects. Finally, section seven will conclude.

⁵By using detailed data on worker characteristics and productivity at the sub-firm level, this paper also connects to a broader literature on the interplay between management, worker characteristics and productivity. Amodio and Carrasco (2015) exploit exogenous variation in worker productivity in a setting with quasi-team incentives, and show free-rider effects among co-workers, with the effect being ameliorated by social ties between workers or the introduction of piece rates per worker. Hjort (2014) is a case study of a Kenyan flower packaging factory showing that ethnically diverse work teams have lower productivity than ethnically homogeneous teams. Similarly, within the context of the garment industry, Kato and Shu (2011) use data from a Chinese garment factory to show that the effect of team incentives to increase productivity depends on the composition of work teams out of urban and rural migrant workers. Furthermore, Hamilton et al. (2003) study the introduction of team work in a U.S. garment factory, which has a similar set-up as the factories I study in Bangladesh, and find a significant increase in productivity from team work. Das et al. (2013) investigate the effect of worker training on productivity in an Indian steel mill. Finally, Nagin et al. (2002) randomly vary the supervisory monitoring rate at a telephone solicitation company, to investigate how this affected the shirking behavior of employees.

2 Background and Data

This study was conducted at three large garment factories in Bangladesh, which in recent years has emerged as the third largest garment exporter in the world.⁶ By local industry standards, the three factories involved in this study are very large and modern. Both ownership and management are domestic, and all output is produced for the export market. The factories produce mainly t-shirts, polo shirts, shirts and trousers. The factories vary in size, with 1,200 to 5,000 workers employed in their sewing departments. Table 1 provides key characteristics of the three factories in the sample. The smallest is Factory 2, with 17 sewing lines located on four sewing floors, but more workers per line than the other factories. The largest is Factory 3, with 183 lines located on 14 floors. Though Factory 3 has many more lines, each line in Factory 3 has less than a third of the number of workers than a line in Factory 2. Another distinction: factory 3 uses fewer line chiefs; in most cases, two lines share one line chief, and in some cases even four lines share one. Factory 1 lies in-between the other two factories on most dimensions. It has 59 lines on six different sewing floors; each line has its own line chief, and the size of the lines is closer to those in Factory 3. Note that the bulk of variation in the number of workers per line is cross-factory variation, stemming from the specialization of Factory 1 and 3 into knit garments, and of Factory 2 into woven garments. Within factories, sewing lines are much more homogeneous in terms of the number of workers allocated to lines; this can also be seen in the relatively low standard deviation of workers per line in each factory as shown in Table 1. Finally, workers in Factory 1 are faced with new styles on average every 16 working days, while lines in Factory 2 and Factory 3 take on new styles roughly every 8-10 days.

Sewing lines are organized as assembly lines in which each worker only does one sewing operation, then passes on the garment to the next worker who does another sewing operation. Additionally, each line has one to three quality inspectors, and garments found with quality defects that cannot easily be rectified are sorted out and not counted in the line's output measure. The main tasks of the line chief are to respond to problems from workers, to instruct workers on new tasks associated with a new style's production, and, generally, to keep work discipline and productivity high. The factories' central planning departments allocate workers (including line chiefs) and styles to specific sewing lines. Workers have fixed lines to which they are allocated, and usually change lines only as the result of promotions to new positions - though workers occasionally switch to different lines to replace absent workers. However, workers with production experience on some styles are generally not reallocated to other lines if these lines also start producing the same style. Thus, it is unlikely that such reallocations of workers drive the observed productivity spill-overs across lines producing the same style. Appendix A furthermore presents the results of a placebo test which also indicates that worker movements are not behind the productivity increases of later lines producing the same style.

The main dataset used for the analysis contains line-level production data for all lines in the

⁶Source: WTO, International Trade Statistics 2014: www.wto.org/english/res_e/statis_e/its_e.htm

Table 1: Factory Characteristics

	Factory 1	Factory 2	Factory 3
Nbr. Sewing Floors	6	4	14
Nbr. Sewing Lines	59	17	183
Nbr. Workers in Sewing Section	ca. 2000	ca. 1200	ca. 5000
Nbr. Workers in whole Factory	ca. 5000	ca. 2000	ca. 9000
Nbr. Buyers	28	74	10
Nbr. Styles in Data	866	839	1048
Avg. Nbr. Lines /Style	3.12	1.49	3.94
Avg. Nbr. Days /Style & Line	16.4	9.5	8.5
Avg. Nbr. Workers /Line	30.9	72.2	23.2
S.Dev. Nbr. Workers /Line	8.0	10.8	5.2

Notes: All information from production data collected from factories, except for ‘Nbr. Workers in ...’ which is from surveys of factory management.

factories for 30 consecutive months from Factories 1 and 2, and for eight consecutive months from Factory 3. (This factory was recruited for this project only at a later point in time.) This dataset includes: daily sewing line productivity; the identifier of the style being produced by a line on a given day; its buyer; and the Standard Minute Value (SMV) of the style. The SMV is a style-specific value, calculated prior to the start of production of the style. It is the sum of the time, in seconds, it takes to perform each sewing operation to assemble one piece of the style. Thus, it provides a measure of the required labor input per piece under ideal production conditions. To calculate the daily line-productivity measure, daily piece output is multiplied by the style-specific SMV, and then divided by total labor input on that line and day measured in worker-minutes:

$$\text{Productivity} = \frac{\text{Daily Output} * \text{SMV}}{\#\text{Worker} * \text{Daily Hours} * 60^{\text{mins}}}$$

Because the SMV is also essential in price negotiations with the buyers, the calculation of the SMV is done professionally, and its breakdown into the individual sewing operations is scrutinized by the buyer. Therefore, the productivity measure is of high quality and comparability across different styles.

I also conducted a survey of all factory line chiefs to collect demographic and career information, plus information on the social networks between all line chiefs on six dimensions: kinship, knowing each other prior to working at the factory, having previously worked together at another factory, having visited each other’s home, spending lunch breaks together, and generally being friends. Table 2 presents summary statistics from the surveys, while Appendix E

Table 2: Line Chief Characteristics

	Factory 1	Factory 2	Factory 3
Avg. Age	30.0	30.0	32.2
S.Dev. Age	3.8	5.8	5.6
Avg. Years working in Factory	4.3	5.0	6.2
Avg. Y.s working as LC in Fact.	2.4	1.8	3.5
Avg. Y.s working as LC on current line	1.2	1.5	2.1
Promoted internally	58%	73%	52%
N	53	15	60
N Female LCs	1	1	0

Notes: All information from survey of all line chiefs in factory. Promoted internally is percentage of current line chiefs who worked on lower position in same factory before and were subsequently promoted to line chief.

provides more information on the network data. All line chiefs are around the age of 30, and only two out of 128 line chiefs interviewed were female.⁷ They have worked as line chiefs on average for 1.8-3.5 years in total at the factories, and on average for more than one year with the line they were supervising at the time of the survey. This also fits with the accounts from the factory management and the production data; line chiefs generally have a fixed line and are only rarely reallocated. At all three factories, line chiefs report to have on average ca. 10 years of schooling, which is equivalent with the Bangladeshi Secondary School Certificate (SSC).

Sewing lines are kept homogeneous in terms of size and productivity within the factories by the management, and workers are not sorted to lines according to experience or productivity.⁸ The reason for this lies in the high flexibility required in operations. Buyers place orders with low predictability and close delivery deadlines, and frequent disruptions to the production process (power failures, unrest outside factories, problems in supply and delivery chains, missing inputs) often require reallocations and re-prioritization of orders to lines. Therefore, it is not optimal to have differentiated lines specializing on certain types of garments. This non-specialization of lines on certain types of garments can also be seen in the stacked bar charts for each of the three factories in Figure 2, in which each bar represents a sewing line, and the wider spaces between the bars separate sewing lines located on different floors. The differently colored parts of the bars represent the shares of different garment types (e.g. t-shirts, polo shirts, pants,) among all styles the lines produce. While some variation can be expected, in general, the graphs show few patterns of lines specializing on certain types of garments.

⁷The average share of female workers on the sewing lines is around 80 percent. Workers typically start working in the garment industry at the age of 18 (the result, these days, of child labor regulations enforced through foreign buyers), and stop by the age of 25 to 30, unless promoted to quality control, mechanic or supervisory positions. However, very few women get promoted to these positions.

⁸An exception are the “sample lines” on which the most experienced employees often work to produce samples of new orders for buyers during the negotiations process. Sample lines are not included in my dataset.

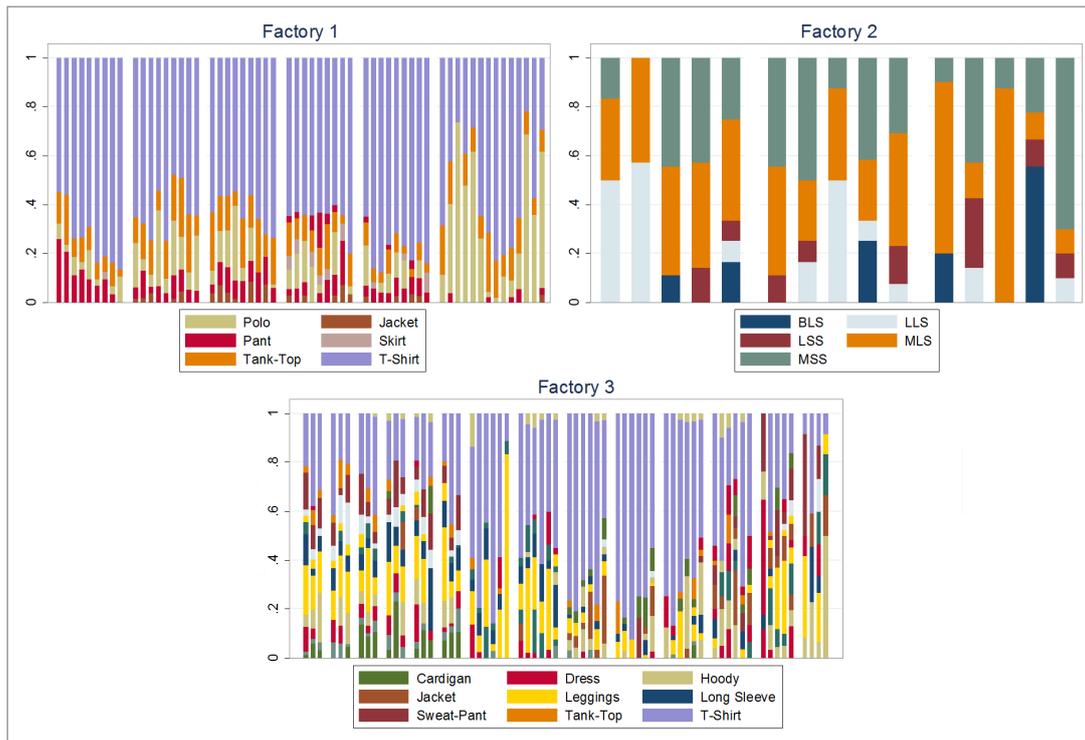


Figure 2: Garment Types produced on different Sewing Lines

Graphs represent the types of garments produced by different sewing lines at the three sample factories. Each bar in the graphs represents a sewing line, and the wider spaces between bars separate sewing lines from different sewing floors. The differently coloured stacked parts of the bars represent different types of garments that the lines produced. Legends show colours for most common garment types only, for illustration. In sub-graph of Factory 3, each bar represents a line chief instead of a line (some line chiefs at this factory look after 2 or 4 lines), to keep the number of bars in graph parsimonious. Graph shows types for only 15 out of 17 lines for Factory 2, as type data is missing for two lines.

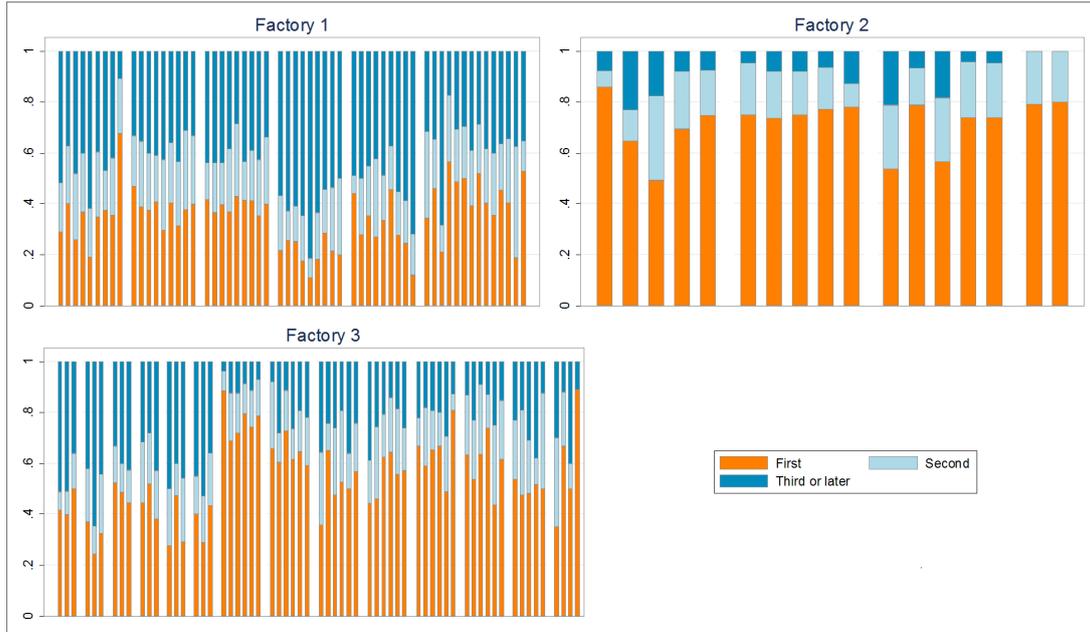


Figure 3: Start Ranks of Styles produced on different Sewing Lines

Graphs show for each sewing line in the three factories for which share of the styles they produce the lines are the first (orange), the second (light blue), or the third or later line (dark blue) to produce that style in the factory. Each bar represents one line, and the wider spaces between bars separate lines located on different sewing floors. In sub-graph of Factory 3, each bar represents a line chief instead of a line (some line chiefs at this factory look after 2 or 4 lines), to keep the number of bars in graph parsimonious.

Lines also do not specialize in whether they are typically the first or a later line in the roll-out of styles across lines. Figure 3 shows a similar stacked bar chart as Figure 2, but this time the differently colored parts of the bars indicate the share of styles the line produced first (orange), second (light blue), or third or later (dark blue) in the factory. Again, few obvious patterns of lines being more often allocated styles early on or later can be seen.⁹ According to the production engineers in the planning departments, incoming orders are prioritized based on the importance of the buyer to the firm, and how close the delivery date is, and are then essentially allocated to the next free line. As a result, higher productivity on lines making previously produced styles is unlikely to be driven by a selection process deploying higher-productivity lines later, rather than first, in the allocation of styles to lines.

3 General Evidence on Productivity Spill-over

Before turning to the results from the randomized intervention, this section explores in the overall production dataset the extent to which line productivity profits from output of the same style

⁹At Factory 3, the six floors to the left of the graph produce for one large buyer, while the other floors produce for other buyers. Because the orders from this buyer are larger, they are produced on average on more lines. As a result, lines on these floors are on average less often the first line to produce a given order, and instead are more often second or later.

previously produced on other lines. Given the evidence from Figure 3, lines do not specialize either in initiating or being a second or later line taking on a new style’s production. Therefore these effects should not be driven by selection of higher-productivity lines into producing styles at a later stage in the roll-out process of styles across lines.

In the overall dataset, I observe 1,257 styles that have been produced on more than one sewing line in the factories. In total, there are 4,964 instances of a sewing line starting to produce one of these styles - a process I refer to from now on as a “style start.” Therefore, these styles are produced on average on 3.95 different lines. Figure 1 shows that on average sewing lines reach their long-run productivity level again five days after starting to produce a new style. Therefore, for the sample of the regressions, I keep the daily line productivity observations from the first five days a line produces a new style. Denote the n^{th} day a sewing line produces a style as the n^{th} “style-day.” Thus, the sample consists of all observations from style-days less or equal to five.¹⁰

The basic econometric model I estimate in this section is of the following form:

$$y_{fisnt} = \sum_n \beta_n^A \ln(A_{isn}) + \sum_n \beta_n^F \ln(F_{isn}) + \alpha_{fin} + \gamma_{ft} + \delta_s + \epsilon_{fisnt} \quad (1)$$

Productivity y_{fisnt} of sewing line i in factory f producing style s in week t on its n^{th} style-day is regressed on output A_{isn} of the same style that has already been produced on all other sewing lines in the factory up to, but excluding, the day on which line i started producing style s . I interact this previous output from other lines with fixed effects for style-day n . Thus, the effect of previous output of the same style is estimated separately for each of the first five style-days included in the sample, to see for how long previous output affects productivity of a new line producing the same style. I use the log of previous output of the same style as I expect each additional produced piece of the style to have a diminishing marginal effect on the stock of knowledge with the style held by workers on other lines.¹¹

I additionally include in the regression the output F_{isn} of the same style produced on all other lines on the same sewing floor, with its effect estimated again separately for each style-day. Sewing lines in the three factories are bundled on sewing floors which contain on average five to 10 lines, and sewing lines located on the same floor are running parallel and only two to three meters apart from each other. Sewing floors, on the other hand, are either located on top of each other in the same building, or in different buildings. Therefore, to get from a sewing

¹⁰At Factory 3, most line chiefs supervise two or four lines. As I am primarily interested in studying learning and productivity spill-overs between line chiefs, I only keep the observations of a line chief producing a certain style from the first of his or her lines the line chief produced the style on. If a line chief produces a certain style on more than one of his or her lines, I drop all observations from those lines that were not the first the line chief produced the style on. Thus, I exclude “within line chief” spill-overs across lines supervised by the same line chief from all analysis in this paper.

¹¹More precisely, I use the log of previous output on other lines plus one, so that when previous output on other lines is equal to zero, it will also be zero after the transformation.

line on one floor to one on another requires at least leaving one’s line out of sight and calling distance. Furthermore, each sewing floor typically has its own floor production manager, who could transfer knowledge he gained on one line to other lines on his floor.¹² For these reasons we could expect a priori the effect of production experience with the same style gained by lines on the same floor to differ from the effect of experience gained by lines on other floors. In total, while at 2,905 of the 4,964 style starts in the sample some other line has previously produced the style, at 1,888 out of these starts was the style already produced on a line on the same floor.

I control for fixed effects α_{fin} on the line chief - style-day level. Thus, I estimate the effect of previous output of the same style (on the same floor) as deviation of line productivity from learning curves estimated separately for each line chief.¹³ Furthermore, I include time fixed effects γ_{ft} on the factory-week level and fixed effects δ_s on the style level. Standard errors are clustered on the line chief level.

Column 1 of Table 3 shows the results from estimating the empirical model from equation 1. The “All Other Lines, Day n ” coefficients represent the effect of previous output of the same style on all other lines in the factory for each of the first five style-days, while the “Lines Same Floor, Day n ” coefficients represent the additional effect of prior output from other lines on the same floor. Previous output on lines on the same floor has a significant positive effect over the first three days the line produces the new style, with the effect becoming gradually smaller and less significant. From the fourth day of production onwards, no effect is visible anymore. On the other hand, overall output on any other line in the factory seems to only have a very weak effect on productivity of later lines, if output from the same floor is included separately. Only on the second day does the coefficient become significant. Thus, indeed, productivity spill-over seem to occur mainly among spatially close lines located on the same floor.

To better gauge the magnitude of the effect, column 2 repeats column 1, but replaces log previous output of the same style (on the same floor) by a dummy indicating whether the style had already been produced on any other line (on the same floor) before. Again, these two dummies are interacted with style-day fixed effects, to separately estimate the effect of previous output for each of the first five style days. The results closely mirror those from column 1; if another line on the same floor has produced the same style before, productivity is increased by almost four efficiency points on the first day later lines produce the same style. Overall average productivity in the production data from the three factories is 47.6 efficiency points, while average productivity on the first day lines produce a new style that has not been produced on another line before is 30.1. Thus, the fact that other lines on the same floor have already produced the style reduces the productivity drop on the first day a new style is produced by 22%, which equals

¹²I use “he” and “his floor” because all floor managers in the three factories are males.

¹³The later results on the experimental intervention use line chief fixed effects because the intervention treats line chiefs. Thus, for consistency, I also use line chief fixed effects in this section. All results are qualitatively similar when using line fixed effects instead of line chief fixed effects.

an increase of 0.17 standard deviations of productivity on the first day lines produce a new style.

Most of the literature on organizational learning used such productivity spill-overs as evidence for learning and knowledge exchange in firms. However, such effects could also be driven by other peer effects. The mere fact that other workers in the factory produce the same product could increase productivity even without learning effects. Workers could compete about who is most productive with a product; or, the productivity of some workers could provide the factory management with a benchmark against which it could compare productivity of other workers, and, therefore, more easily find out if workers slack. To find more evidence on whether knowledge transfers are indeed the drivers of the productivity increases of later lines producing the same style, I run a simple placebo test. If the increases in productivity that we observed are driven by competition or less slack, these effects should arguably be even more prevalent if lines start producing the same style on the same day; this is because, these cases present a more-level playing field for competition or comparing productivity. Furthermore, these cases provide opportunities to study the effect of lines simultaneously starting production of a style when no other lines have previously produced the same style; in these cases, obviously, no other workers have prior experience with the style that could be shared, and learning effects should be absent. Thus, in column 3 of Table 3, I study the productivity of the first lines in the factory to produce a certain style, and check whether it is increased if more than one line starts producing it on the same day, and if this effect is stronger if the other lines starting the style on the same day are located on the same floor.¹⁴ And, indeed, strong and significant such effects can be seen in column 3.

This effect among first lines could be driven by selection of certain styles into instances when more than one line is producing it for the first time - such as when new starts are rushed for a style that needs to be completed quickly, and, as a result, more than one line begins producing it on the same day. Furthermore, because each style is only started once for the first time in the factory, I cannot use style fixed effects. Instead I can only control for observables on the style level, SMV and fixed effects for the type of the garment (t-shirt, polo, pant, etc.). Thus, the effects could be confounded by selection of styles with certain characteristics to multiple first lines which is not controlled for by their type or SMV. Finally, the presence of the effect among the first lines does not automatically imply that the effect of previous output among later lines is not driven by learning effects; the two effects could be explained by different mechanisms.

To gain additional insight on whether the effect among the first and later lines is driven by the same or different mechanisms, I conduct a “horse race” among the two specifications in column 4, which replicates column 1, but adds log output produced by lines starting the style

¹⁴I regress productivity on log output produced on all other lines on the first day of production which also started to produce the style on the same day. There are 2,002 instances in the data of a line starting a style that has not been produced before on another line. At 1,170 of them, some other line in the factory also starts producing the same style on the same day, and at 1,054 instances, at least one of these other lines is located on the same floor.

Table 3: Non-Experimental Results

VARIABLES	(1)		(2)		(3)		(4)	
	Effic.	SE	Dummy Spec. Effic.	SE	First Lines only Effic.	SE	Effic.	SE
PRIOR OUTPUT								
All other Lines, Day 1	0.116	(0.11)	0.265	(0.99)			0.125	(0.10)
All other Lines, Day 2	0.197**	(0.09)	1.303*	(0.71)			0.212***	(0.08)
All other Lines, Day 3	0.089	(0.09)	0.912	(0.74)			0.113	(0.09)
All other Lines, Day 4	0.113	(0.10)	1.465*	(0.86)			0.135	(0.09)
All other Lines, Day 5	-0.008	(0.09)	0.257	(0.75)			0.004	(0.08)
Lines Same Floor, Day 1	0.434***	(0.12)	3.907***	(1.02)			0.537***	(0.11)
Lines Same Floor, Day 2	0.194**	(0.09)	1.854**	(0.78)			0.260***	(0.09)
Lines Same Floor, Day 3	0.147*	(0.09)	1.281*	(0.71)			0.175**	(0.08)
Lines Same Floor, Day 4	-0.040	(0.10)	-0.645	(0.78)			-0.046	(0.09)
Lines Same Floor, Day 5	0.010	(0.10)	-0.223	(0.88)			-0.011	(0.10)
OUTPUT SAME DAY								
All other Lines, Day 1					2.476***	(0.66)	1.158***	(0.28)
All other Lines, Day 2					1.299***	(0.48)	0.464**	(0.19)
All other Lines, Day 3					0.835**	(0.38)	0.100	(0.18)
All other Lines, Day 4					0.376	(0.43)	0.035	(0.18)
All other Lines, Day 5					0.214	(0.37)	-0.158	(0.22)
Lines Same Floor, Day 1					1.315*	(0.70)	0.845***	(0.28)
Lines Same Floor, Day 2					1.164**	(0.51)	0.854***	(0.18)
Lines Same Floor, Day 3					0.698*	(0.42)	0.614***	(0.18)
Lines Same Floor, Day 4					0.484	(0.46)	0.082	(0.19)
Lines Same Floor, Day 5					0.081	(0.37)	0.049	(0.24)
Constant	30.99***	(3.42)	24.95***	(2.80)	23.55***	(5.62)	43.21***	(4.26)
Observations	17,147		17,150		5,684		17,107	
R ²	0.682		0.682		0.604		0.695	
L.Chief-Style.Day FE	YES		YES		YES		YES	
Style FE	YES		YES				YES	
SMV					YES			
Garment Type					YES			
Week-Factory FE	YES		YES		YES		YES	

Column 1 shows the results from regressing daily sewing line productivity from the first five days a line produces a new style (first five style-days under PRIOR OUTPUT) on the log of output of the same style produced on all other lines in the factory up to, but excluding, the day the line starts producing it as well. This previous output is interacted with fixed effects for the n 'th style-day (Coefficients "All other Lines, Day n "). Similarly, log previous output of the style produced on all other lines on the same sewing floor, interacted with style day is included separately (Coefficients "Lines Same Floor, Day n " under PRIOR OUTPUT). Regressions control for fixed effects on the line chief - style-day level, style level, and factory-week level. Column 2 replicates column 1 but with dummy variables for the style having been produced on other lines (on the same or other floors) before, instead of log previous output. Column 3 regresses instead daily line productivity on the first five style days from the first lines which produce a certain style in the factory on log output of the style from lines which also started producing the style on the same day, interacted with style-day (Coefficients: "All other Lines, Day n " under "OUTPUT FIRST DAY"), and separately from lines located on the same floor (Coefficients: "Lines same Floor, Day n " under OUTPUT FIRST DAY). Column 4 regresses output on both previous output and on output from other lines starting to produce the style on the same day. Standard errors are clustered on the line chief level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

on the same day as additional independent variables. Note that this output is not mechanically co-linear with previous output of the same style, because previous output includes output only up to, but excluding, the day the line also starts producing the same style. Thus, it cannot include output from lines starting the style on the same day. In fact, previous output and output from lines starting the style on the same day are negatively correlated, both overall and within units. Because in this specification I can observe again multiple starts per style on different lines, I include style fixed effects again. The results from column 4 show that in this specification, both effects that we saw in column 1 and 3 are still present. Lines are more productive: 1) the more of the style has previously been produced on other lines, and 2) the more of the style is being produced on lines that started producing it on the same day.

To conclude this section, the existence of spill-over effects among first lines that started to produce a certain style on the same day should caution against interpreting the effects of output on other lines as learning - even though both peer and learning effects seem to persevere in a regression that attempts to incorporate both effects. However, even when controlling for output from other lines on the same day, previous output could still capture other possible forms of peer effects. Therefore, the next section presents the results of a randomized intervention that introduced exogenous variation in the likelihood that production knowledge on garment styles is communicated, and shows that this intervention did have an effect on line productivity when lines started new styles.

4 Randomized Help Provision

To identify the effect that knowledge exchange between co-workers has on productivity growth, I carried out a randomized management intervention at the three sample factories. Whenever a line on randomly selected “treatment” sewing floors began producing a new style that had already been produced by any other line chief in the factory, the most senior line chief with previous experience on the style was instructed by the factory’s production management to give a brief to the line chief without the style-specific experience. The 15- to 30-minute brief was intended to explain how the earlier line had overcome initial problems and “bottlenecks” that slowed down the new style’s production.¹⁵

¹⁵Among the 10 to 20 operations needed to assemble one piece, each style typically has a small number (one to five) of bottleneck operations that slow down or, once resolved, increase the productivity of the lines. When lines switch to a new style, the factories attempt to achieve “zero-feeding” which means that one machine after the other on the line is adjusted for the new style, with those machines already switched producing the new style, while those not yet switched still producing the previous style. If zero-feeding works according to plan, each machine can be idle for as little as 15 minutes in the switching process. The adjustment of machines is typically done by the line chiefs together with a production engineer, who also briefs the line chief and workers on the line on the new style. The briefings from the communication intervention were to have been done once the switching process had been completed and all machines on the line had been adjusted for the new style. Its goal was to convey the intricate knowledge exclusively held by a line chief who had already produced a style for several days (e.g. how to hold a garment in the hand when doing a certain stitch); expertise that the production engineers don’t have, who design the sewing line layout for a style from its design template, and who supervise

The intention of this treatment was to exogenously increase the potential for knowledge exchange on the production process of the style between randomly selected pairs of line chiefs, by lowering the costs of helping. In particular, the intervention can be thought to decrease two parts of the cost of seeking and providing help. First, the possible perceived cost of approaching someone else for help, as one exposes a lack of knowledge on how to solve certain production problems (Lee (2002); DePaulo and Fisher (1980)). Because higher ups direct someone to share his or her experience with the style, knowledge is shared without an initial request for help that could reveal a lack of knowledge. Second, especially if the person asked to share his or her knowledge needs to go to the workplace of the other person to provide effective help, help provision can be thought to have a fixed and a variable cost component. The distraction of listening to someone’s request for help, and possibly moving to the other person’s workplace would constitute fixed costs. Once these costs are borne, one would need to decide how much effort to spend analysing the problem and figuring out an effective way to communicate a possible solution; this introduces a variable cost component into help provision. While the randomized help provision eliminates neither the fixed nor the variable costs, it does render the fixed costs (engaging with the other worker and moving to his or her workplace) as sunk, because the worker cannot decide anymore whether or not to bear this cost. Thus, we can think of this fixed cost as being taken out of the cost-benefit analysis undertaken by the line chief asked to provide help, when deciding whether to do so.

The experiment ran on the treatment floors for four months, from June through September 2014. The production data show 377 “non-first” style starts on the treatment floors during this time; these are style starts at which some other line chief had already produced the same style, and therefore, the line chief now starting the style could and should have been briefed in the randomized experiment. The treatment protocol was implemented by the industrial engineers from the factories. The engineers were provided with experimental logbooks to record each instance of a treatment of a style start. According to these logbooks, 154 briefings among line chiefs were actually conducted, of which 98 could be matched with a style start in the production data.¹⁶ However, it is likely that compliance was much higher than indicated by these numbers. The implementing engineers admitted underreporting of treatments in the logbooks. Among the actually treated style starts, there is likely to be selection into treatment of starts for which the treatment was expected to have a stronger effect.¹⁷ For these reasons, the analysis

the switching process of lines to new styles, but do not stay on the line after the switch.

¹⁶26 further briefings out of the 154 recorded ones could also be matched to style starts in the production data; however, according to the production data, in these instances no other line chief had previously produced the style. It seems that in these instances, line chiefs who had already produced similar styles were sent to give instructions.

¹⁷All three factories reported that prior to the intervention, they occasionally sent line chiefs with experience producing a new style to other lines to help co-workers with new styles. However, this behavior was not institutionalized in any of the three factories. To the extent that the factories already had asked line chiefs to help each other, the factories were instructed to not change their behavior on the control floors, while on the treatment floors, the factories were instructed - without exception - to send line chiefs to brief others who were starting production of the same styles.

will focus on the “intention to treat” effect, assuming that any start of a style that should have been treated was actually treated. Appendix B shows results when regressing productivity on recorded instances of the treatment.¹⁸

The sample from which floors were randomly selected consisted of 17 floors across the three factories (Factory 3 requested to include only six out of its 14 floors in the sample).¹⁹ Randomization across floors was chosen to make compliance with the randomized implementation as simple as possible for the factory management. The original intention was to randomize treatment across sewing lines. However, the factory managements were worried that it would be too difficult for their staff to remember which lines should be treated and which not. Furthermore, there was the concern that if the intervention were implemented at some lines on a given floor, and if it were to prove effective, its implementation would quickly be copied by other lines in the same floor, which usually operate just a few meters away. Table 4 shows tests of balanced outcomes over observable average line and line chief characteristics between treatment and control floors from April and May 2014, just before the start of the intervention, when the random selection of units was done. No observable line or line chief characteristics differ significantly on conventional levels, except for average line productivity, which was lower on treated floors (p-val. 0.082).²⁰

Because the intervention was conducted at the end of the time covered by the collected production data, a substantial amount of pre-intervention data is available. Figure 4 plots the average productivity over the first four days a line produces a new style for four different cases: treated lines before and during the intervention, and control lines before and during the intervention. I use data from the beginning of 2014 until end of September 2014, when the initially defined treatment time ended, and only include data from the non-first style starts, which are in principle “treatable” because another line chief has previously produced the same style. Prior to the start of the intervention, and compared to control lines, treated lines had on average lower productivity values in the first days a new style was produced, which fits with the overall lower productivity among lines in treatment floors shown in Table 4. This difference is not accounted for by observable characteristics of lines or line chiefs, and is driven by two out of the three factories. However, as shown in Figure 4, while productivity remains constant across pre-treatment and treatment time on control floors, treated lines experience an upward shift in their learning curves during the time of the treatment. Interpreting the results in a difference-in-differences framework, the intervention indeed had an effect in raising productivity, especially on the first day a new style was produced.

¹⁸Note that there is no indication of style starts on control floors being treated. The logbooks do not show any such treatment; in addition, the production managers who implemented the intervention showed no sign of confusion about which style starts should and should not be treated.

¹⁹In fact, the sample of floors over which the randomization occurred consisted of only 15 sewing floors. However, at factory 1 and 2, one floor was randomly chosen at each factory, and one (physical) half of the floor randomly selected into treatment. Therefore, the randomization occurred effectively across 17 units, 13 full floors, and 4 half floors.

²⁰All differences, except for average productivity, remain insignificant when controlling for factory fixed effects.

Table 4: Balancing of Randomization across Sewing Lines

	Control Lines	Diff. Treated Lines	Diff. p-val.	N
<u>Line Chief Characteristics</u>				
Age	30.06	0.16	0.89	73
Seniority Factory (months)	61.79	2.48	0.80	74
Seniority as Line Chief (months)	34.71	0.12	0.99	74
Seniority as LC current line (months)	9.45	6.54	0.10	71
Promoted Internally	64.7%	-0.02%	0.85	74
Nbr. Social Connections	2.66	0.47	0.46	74
Education	15.27	-0.22	0.66	72
<u>Line Characteristics</u>				
Avg. Number Worker	28.92	1.59	0.35	140
Avg. Daily Hours	9.57	0.16	0.36	140
Avg. Productivity	53.53	-3.09	0.08	140
Avg. SMV	10.76	-0.81	0.40	140

Notes: Line Chief characteristics from line chief surveys. Line characteristics from production data. Values in “Diff. Treated Lines” columns show deviation of average values from treated lines from those from control lines, with “Diff. p-val” showing the p-value of the difference.

Using a difference-in-differences framework to identify treatment effects, one should check whether pre-trends differ between treated and non-treated units. Figure 5 provides such a check. For the first day a line produces a new style that has already been produced by another line chief before, it plots monthly average productivity from January through September 2014, separately for lines selected for treatment (square symbols) and as control lines (triangle symbols). The solid vertical line indicates the start of the treatment with June 2014.²¹ Indeed, first-day productivity was systematically lower on floors selected for treatment in the months before the start of the treatment. This difference is then greatly reduced with the onset of the intervention due to an upward shift of first day productivity on treatment floors, especially when compared to the three months directly preceding the start of the intervention. The data, however, are noisy, leading to fluctuations in first-day productivity across months. These fluctuations are greatly reduced when introducing style fixed effects, as shown in Figure 6. In this specification, the effect of the intervention is derived from styles that are either started on different lines both before and after the start of the intervention, or on both treatment and control floors. This is the case for 35 percent of all styles in the sample of the graphs, which, however, make up 64 percent of all style starts in that sample. Using this specification, especially in the three months before the start of the intervention, first-day productivity on treated and control lines

²¹Data from Factory 3 becomes available from April 2014 onwards, resulting in smaller confidence intervals from that month onwards. To disentangle pre-trends from composition effects, Appendix C1 shows graphs for Factory 1+2 and Factory 3 separately, which look qualitatively similar.

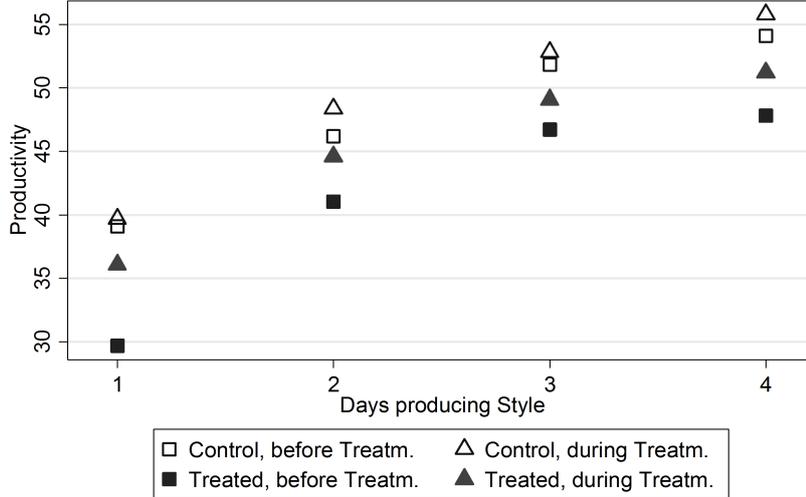


Figure 4: Pre-and Post- Treatment Learning Curves of Treated and Control Lines. Figure plots average productivity over first four days a line produces a new style previously produced on other lines for four different cases: From treatment floors prior to start of treatment (solid square symbols), treatment floors during experiment (solid triangle symbols), control floors prior to start of treatment (hollow square symbols), and control floors during treatment (hollow triangle symbols). Productivity standardized on factory-level.

follows parallel downward trends. And while this trend more or less continues on control lines, it clearly reverses on treatment lines with the onset of the intervention, with gains in productivity especially notable from the second month of the intervention onwards. That the effects are only visible from the second month of the intervention onwards could indicate that the intervention was initially not very effective in increasing productivity, possibly because the line chiefs needed to get used to the communication intervention.²²

To estimate the intention-to-treat effect of the intervention in a difference-in-differences approach, I keep, similar as in the previous section, the observations from the first three style-days from each style start. I restrict the sample to observations from January through September 2014, the same as used for Figures 4 - 6. Using this sample, I run the following baseline regression:

$$y_{fisnt} = \sum_n \beta_n^T Treat_{fisn} + \alpha_{fin} + \gamma_{ft} + \delta_{isn} + \epsilon_{fisnt} \quad (2)$$

Productivity y_{fisnt} of line i producing style s on one of the first three style-days n is regressed on a dummy $Treat_{fisn}$ for the start of the style being randomly selected for treatment, interacted with fixed effects for the three different style-days included in the sample. I control,

²²The logbooks do not indicate that in the first month of the intervention, a lower share of style starts was treated than in the other months of the intervention. The share was roughly equal across the months, except for July 2014, where the share was roughly twice as large

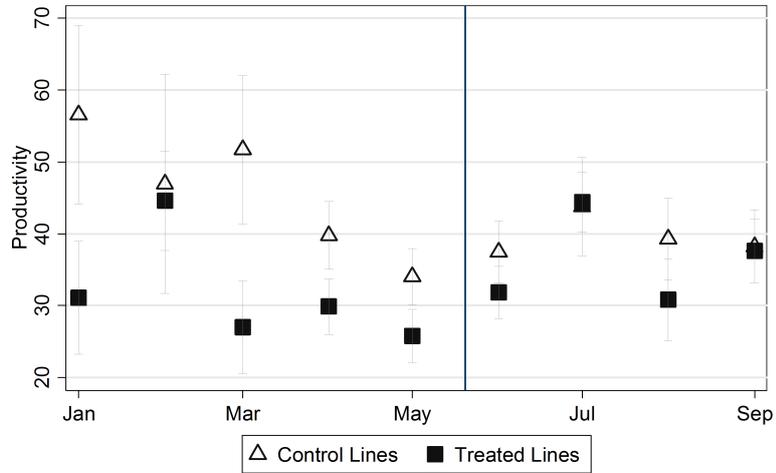


Figure 5: Pre-Post Intervention Start Trends for First Day Productivity. Graph shows average monthly productivity of lines on the first day they start producing a new style, separately for lines selected for treatment (solid squares) and lines not selected (hollow triangles). The solid vertical line indicates start of treatment from June 2014 on. Capped bars represent 95% confidence intervals.

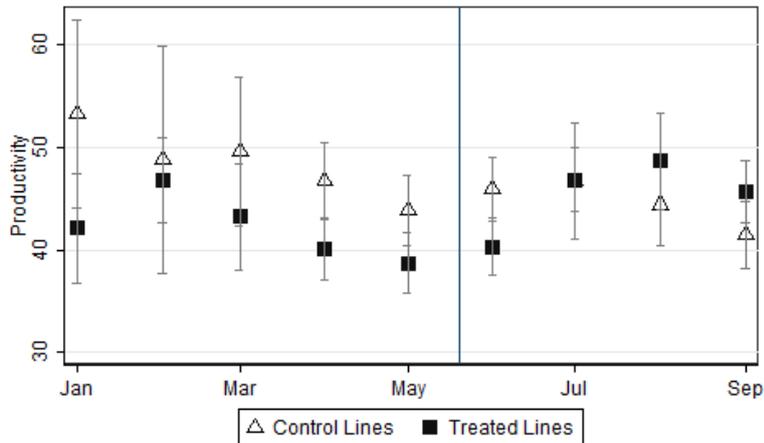


Figure 6: Pre-Post Intervention Start Trends with Style Fixed Effects. Graph replicates Figure 5, after taking style fixed effects out of data. It shows average monthly productivity of lines on the first day they start producing a new style, separately for lines selected for treatment (solid squares) and lines not selected (hollow triangles). The vertical line indicates start of treatment from June 2014 on. Capped bars represent 95% confidence intervals.

as in the previous section, for fixed effects a_{fin} on the line-chief - style-day level, and γ_{ft} on the factory-week level. Furthermore, I include fixed effects δ_{isn} for the “rank” of the line in the roll-out of the style, which indicates how many other line chiefs in the factory have already produced the style, interacted with style-day fixed effects.²³

The reduction in the difference of starting productivity with the onset of the intervention could imply that the results are caused by some other form of catch up of productivity on treatment floors relative to control floors, which coincided with the start of the intervention. To address this concern, I apply the reweighting approach by DiNardo et al. (1996). It reweights observations from the treatment floors such that in the pre-treatment time the average learning curves no longer differ between treatment and control floors. I use the approach in a similar way as Dufflo et al. (2013), who adapted it to control for possible endogenous selection into treatment. Their basic idea is to reweight observations from a controlled experiment such that independent variables that were not balanced pre-treatment between treated and control units become balanced after the reweighting. In this paper, I apply this approach to correct for the fact that the dependent variable of productivity on the first days a line starts a new style is not balanced between treatment and control groups prior to the start of the randomized experiment. Identification using the reweighting approach relies on the assumption that the treatment effect does not depend on the distribution of the independent and dependent variables, as the approach creates artificial counterfactual distributions in the sample used to estimate the treatment effect. More details on the implementation of the approach are shown in Appendix C2.

Figure 7 demonstrates the reweighting approach, replicating Figure 6 after reweighting the data such that productivity on the first day a line produces a new style which has already been produced on another line is balanced in the two months before the start of the intervention, the same time frame used to create the results for the balancing tests from Table 4. Indeed, the graph now shows that in the reweighted sample, productivity on the first day a line starts producing a new style that has already been produced on another line no longer significantly differs between treatment and control floors in the two months before the start of the intervention. In general, productivity on treated and control floors follows a close trend in the four months before the start of the intervention. After the start of the intervention, as the graphs using the unweighted data show, no treatment effect is visible during the first month of the intervention, while a strong positive effect can be seen from the second months onwards.

Due to the small number (only 17) of clusters over which the randomization was conducted, special attention needs to be given to inference, because even standard errors clustered on the 17 floors can be biased downwards, as shown by Cameron and Miller (2008). I follow their suggestion to use wild cluster bootstrap to obtain standard errors at all regressions estimating

²³Instead of rank fixed effects, I could have also controlled for cumulative output of the same style on previous lines, the central variable of interest in the regression from the previous section. However, to more flexibly control for how many lines have previously produced the style, I instead use rank fixed effects (interacted with style-day).

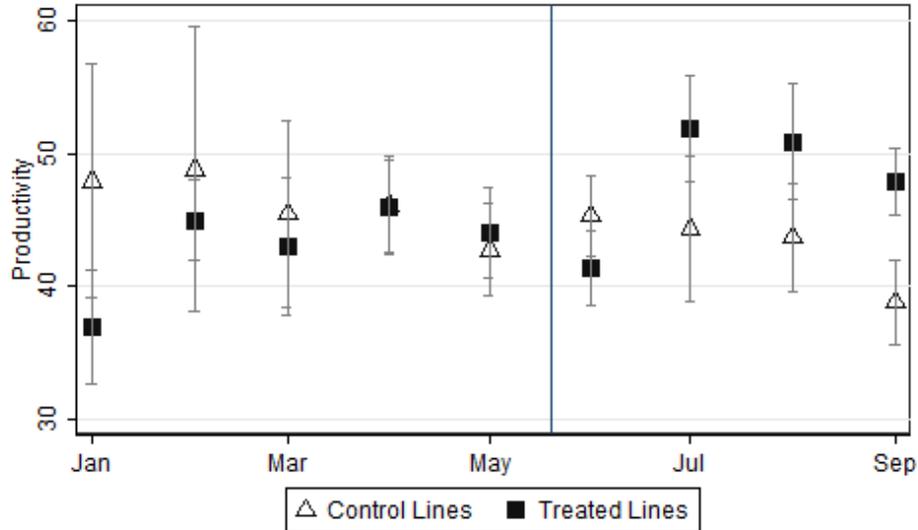


Figure 7: Pre-Post Intervention Start Trends for First Day Productivity, Reweighted Data. Graph replicates Figure 6, with productivity data reweighted following the approach from DiNardo et al. (1996). It shows average monthly productivity of lines on the first day they start producing a new style, separately for lines selected for treatment (solid squares) and lines selected as controls (hollow triangles). The vertical line indicates start of treatment from June 2014 on. Graph controls for style fixed effects. The lines through the symbols represent 95% confidence intervals.

the effects of the randomized intervention.

Table 5 shows the main results from the communication intervention. The regressions shown in each column in the upper and lower panel of the table use the same specifications; however, the lower panel shows the results with the data reweighted using the approach from DiNardo et al. (1996) outlined above, while the upper panel shows the results using the unweighted data. Column 1 shows the results when estimating the model from equation 2 with and without reweighing the data. In both cases, a significant positive effect on productivity on the first day of production of a new style can be seen. Using the reweighted data, first-day productivity is increased by 4.09 efficiency points, which resembles 19.2 percent of the standard deviation of first day productivity when other lines have already produced the same style. Average first day productivity of lines if other lines have already produced the style is 36.5 efficiency points, and overall long-run productivity is 47.6. Thus, the intervention reduces the average gap of first-day-to long-run- productivity by roughly 37 percent. The effect becomes successively smaller and insignificant on the second and third days of production in both cases when using reweighted or non-reweighted data.

Column 2 adds style fixed effects. In principle, the characteristics of the styles produced on treatment and control lines should be balanced, due to their random selection, and Table 4 shows no significant difference in the SMV of styles across these two types of lines. Also Figure 6

showed that the lower pre-treatment productivity of treatment lines is not captured by style fixed effects. However, it did show that including style fixed effects greatly reduces months-to-months fluctuations in first-day productivity, in turn making the treatment effect more visible in the raw data. Thus, column 2 includes style fixed effects, which indeed makes the estimate of the effect on first-day productivity somewhat larger, and more significant when using the reweighted data.

So far, I included all instances of a line starting a new style in the sample. However, only those style starts undertaken when at least one other line has already produced the same style can be treated; the design of the intervention requires the presence of one line chief already experienced with the style, who can administer the briefing. Thus, a more direct way of estimating the intention-to-treat effect is to restrict the sample to the non-first style starts only, both before and after the start of the intervention at the factory, and on treatment and control floors; this is also the sample used for Figures 4 - 7, and column 3 in Table 5 shows the results when repeating the regression from column 2 on this restricted sample. The results remain qualitatively the same, using either the original or the reweighted data.

While style fixed effects control for all unobservable characteristics of the style, they do not capture possible interaction effects of styles and lines. And while lines are, in principle, not specialized on certain types of garments, they could nevertheless be differentially productive for different types of garments, for example, if they happened to have produced more of a certain garment type in the past than another. And while I cannot control for style - line chief fixed effects, because each line chief produces a style for the first time only once, I can control for garment type - line chief fixed effects, to capture interaction effects between lines and classes of garments, such as t-shirts, polo shirts, or pants. Thus, column 4 uses line chief - garment type fixed effects instead of style fixed effects, which leads to a large increase in the estimated effect of the intervention. Given that also the inclusion of style fixed effects increased the estimate of the treatment effect, this result again argues against the possibility that a shift towards easier-to-produce styles on treatment lines spuriously induces a treatment effect. If anything, it points toward a shift in the direction of more difficult-to-produce styles on treatment lines. Furthermore, when using line chief - garment type fixed effects, the effect of the intervention now also becomes marginally significant on the second day of production of the new style.

Figure 8 shows distributions of line productivities on the first day they produce new styles that were already produced on other lines, for four different cases: treatment lines before (Jan-May 2014) and during (Jun-Sep 2014) the implementation of the experiment, and control lines during the same time frames. The increase in first-day productivity of treatment lines during the implementation of the intervention seems to be driven by a strong reduction of the left tail of the productivity distribution. The number of starts with very low productivity is greatly reduced, which is indicative of the individual treatments being enacted specifically when very

Table 5: Main Experimental Results

VARIABLES	(1)		(2)		(3)		(4)	
	Efficiency	se	Efficiency	se	Efficiency	se	Efficiency	se
PANEL 1: ITT REGRESSION								
Treatment, Day 1	5.584***	(0.00)	6.162**	(2.66)	6.596**	(2.83)	12.380**	(6.02)
Treatment, Day 2	2.373	(1.93)	3.325	(2.21)	2.814	(2.55)	6.534	(5.10)
Treatment, Day 3	1.169	(1.82)	1.564	(1.59)	1.347	(2.37)	2.404	(3.74)
Constant	70.271***	(0.00)	28.170	(18.73)	24.067*	(12.79)	39.094	(33.97)
Observations	4,384		4,384		2,942		2,441	
R ²	0.425		0.707		0.730		0.695	
PANEL 2: ITT REGRESSION, REWEIGHTED								
Treatment, Day 1	4.096**	(1.76)	6.320***	(0.00)	5.605**	(2.58)	11.049**	(5.37)
Treatment, Day 2	2.483	(3.15)	3.249	(2.59)	3.659	(2.98)	7.850*	(4.77)
Treatment, Day 3	1.252	(2.45)	1.097	(1.98)	2.151	(2.85)	1.681	(2.08)
Constant	42.354***	(0.00)	48.193***	(0.00)	72.414***	(0.00)	37.688	(24.88)
Observations	4,094		4,094		2,743		2,338	
R ²	0.436		0.729		0.748		0.709	
Start.Rank-Style.Day FE	YES		YES		YES		YES	
LC-Style.Day FE	YES		YES		YES		YES	
Style FE			YES		YES		YES	
LC-Style.Day-Grmt.Type FE							YES	
Week-Fact. FE	YES		YES		YES		YES	

Notes: Table shows the results from regressing daily line productivity from the first three days a sewing line produces a new style (first three “style-days”) on a dummy that the style start should have been treated, interacted with style-day (“Treatment, Day 1”), and on fixed effects for the number of lines that already produced the style before, interacted with style-day, line chief fixed effects interacted with style day, and fixed effects on the week-factory level. Column 2 and 3 add style fixed effects. Column 4 interacts line-chief - style-day fixed effects with fixed effects for the type of the garment (t-shirt, polo, pant.). Panel 2 shows results when estimated using reweighted data according to DiNardo et al. (1996) to balance first day productivity in pre-treatment period (1st April-20th May 2014). Standard errors obtained by wild-bootstrap, clustered on 17 units of randomization. Standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1.

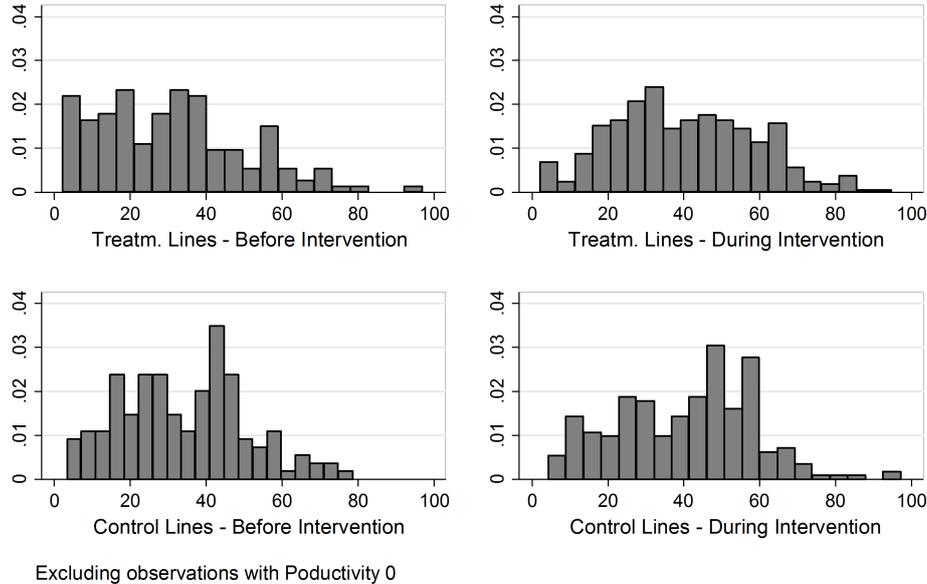


Figure 8: First Day Productivity Distribution before and during Intervention. Graph shows distribution of productivity on first day sewing lines produce new styles already produced on some other line in the factory before, on treatment floors, before and during implementation of intervention (top row), and on control floors, before and during the implementation (bottom row).

low productivity could have been expected. This fits with the fact that fewer treatments were reported in the logbooks than should have been according to the production data. And while the production engineers said that the logbooks underreport the number of actually conducted treatments, they also explained that in cases in which a line chief could be expected to start the new style without any problems, because he or she had prior experience producing very similar styles, no treatment was done because no effect of the treatment was expected.

As an additional check on whether this treatment effect could be caused spuriously by a change in the characteristics of styles on treatment lines with the onset of the intervention, Figure 9 replicates the distribution graphs of Figure 8, however using the SMV of the styles produced on treatment and control lines before and after the start of the intervention instead of first-day productivity. Given that the SMV of a style captures the required labor input to produce one piece of it, the SMV is a proxy for the complexity of the style, and is highly negatively correlated with productivity in the overall data. However, there seems to be no shift in the distribution of SMVs of the styles on treatment lines with the onset of the intervention. Average SMV actually slightly increases, which fits with the estimated treatment effect being smaller when not controlling for style or garment-type fixed effects, even though the shift toward styles with slightly higher SMVs is also visible on control lines.²⁴ Nevertheless, this suggest that

²⁴The increase in average SMVs of the styles could be due to the factories gradually shifting production towards warmer (and more complex) winter garments during the time of the intervention, which are produced in Bangladesh ahead of the cold season in Europe and North America.

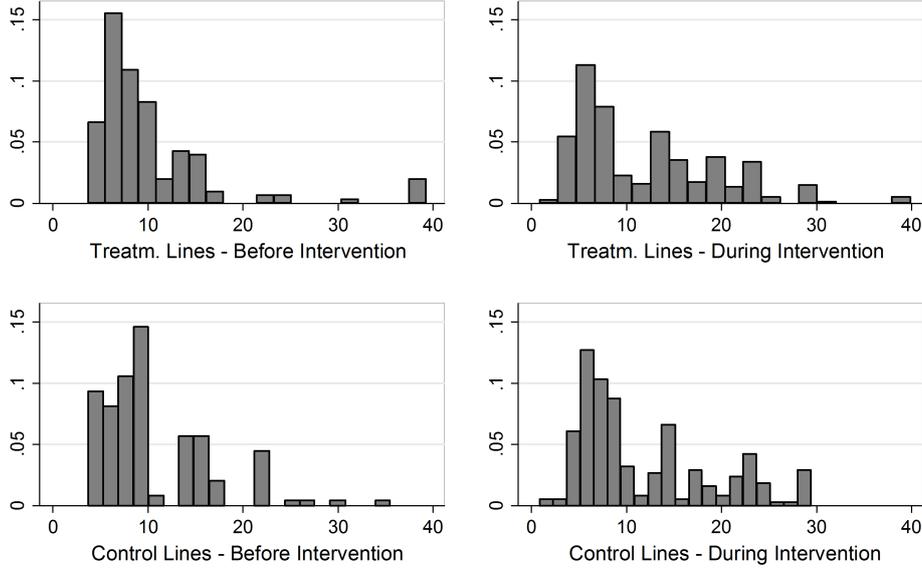


Figure 9: SMV Distribution of Styles before and during Intervention. Distribution of SMV of the styles already produced on some other line in the factory before, on treatment floors, before and during implementation of intervention (top row), and on control floors, before and during the implementation (bottom row).

the treatment effects are unlikely to be caused spuriously by treatment or control lines shifting to a different composition of the styles they produce with the onset of the treatment.

To obtain an idea of the overall effect of this intervention on factory productivity and profits, I use the conservative baseline estimate from column 1, Table 5, which shows a significant increase of productivity of about 4.1 efficiency points on the first day a line produces a new style. Sewing lines on average switch to a new style every 10 days, and at roughly every second of these starts, another line has already produced the style before. Average daily productivity in the three factories is 47.4 efficiency points; thus, a very basic back-of-the-envelope calculation shows that output was increased by $4.1/(10 \times 2 \times 47.4) = 0.43$ percent. Anecdotal evidence shows that labor costs make up around 12 percent of revenue on average in these factories, while the profit margin is about 6 percent. If we assume that the intervention would save 0.43 percent of labor costs, this would translate into an increase in profits of 0.86 percent. On the other hand, the pure monetary costs of the intervention are very low. The hourly wage of a line chief in the factories is about U.S. \$0.50; therefore, the wage cost of a half-hour briefing is about \$0.25. In the largest factory, with more than 180 sewing lines, for example, roughly 3,000 briefings per year would be required to treat lines making a style that is novel to them but has previously been produced by another line. Thus, the yearly monetary cost of the intervention would be \$750. I do not have information on the revenues of the firms, but local business newspapers report that factories of this size generate revenue in excess of \$10 million per year. Using the commonly referenced margin of 6 percent yields estimated profits of \$600,000. A 0.85 percent increase would thus imply an increase in profits of around \$5,000, or a return on the intervention

of more than 600 percent. These high estimated returns are also in line with anecdotal reports from the production managers in the post-intervention survey, that the monetary costs of the intervention were not deemed as a hindrance for its implementation.

5 Randomized Help Provision and Social Ties

The surveys that I conducted among all line chiefs in the three factories in the sample contained questions on social ties they had to other line chiefs in the same factory. This allows for an examination of whether the effect of the treatment is influenced by the presence of social ties with the line chief who provided the briefing. Social ties have been shown to play important roles within firms, such as matching firms with workers (Granovetter (1973, 1995), or Heath (2015) in the case of Bangladeshi garment factories), or to affect effort choice of workers (Bandiera et al. (2010)). Furthermore, outside of organizations, social ties have also been shown to play important roles in learning about new technologies - such as in small-scale agricultural settings (Foster and Rosenzweig (1995); Munshi (2004); Bandiera and Rasul (2006); Conley and Udry (2010)), or about new microfinance services (Banerjee et al. (2013); Cai et al. (2015)). However, due to the lack of documentation of many instances in which help should have been provided, it is often not clear which line chief was sent for the briefing. Furthermore, to the extent that the factory management is aware of social ties among line chiefs, management could have directed socially connected line chiefs to provide the briefings in select instances in which managers expected the line chief to profit either particularly much or little from the briefing, which would bias the estimate of the interaction effect. To address this problem, I exploit the factories initial instructions specifying that the most “senior” line chief who has already produced the style should be the one sent to provide the briefings.²⁵ Therefore, I interact treatment with whether the line chief receiving the treatment reported social ties to the most senior line chief who already had produced the style; this was the case in 59 out of the 377 style starts which should have been treated.²⁶ I measure seniority by the time a line chief already worked as line chief in the factory. To the extent that it was not the most senior line chief according to this measure who was sent to provide the treatment, or that no treatment occurred, the estimated effect can be interpreted as an intention to treat (ITT) effect of the interaction.

Column 1 of Table 6 shows the results of this interaction, replicating column 1 from Table 5, but adding two further dummy variables, each interacted with style-day. The first (“Connected,

²⁵This specific instruction was given to minimize possible resistance among line chiefs against the intervention, in cases in which less senior line chiefs were sent to more senior line chiefs. In these cases, help provision might have not been accepted by the line chief who was supposed to be briefed. Given that only randomly selected line chiefs receive briefings by these senior line chiefs, the intervention essentially estimates the treatment effect of receiving help from the most senior line chief who already produced the same style. This estimate might be the most useful for policy implications, as other factories implementing such a management intervention could likely adapt the same policy, of always sending the most senior worker who has experience with certain processes to instruct co-workers on them.

²⁶The network data I use are directed, in the sense that line chief A is considered socially connected to line chief B if and only if line chief A reported a link on one of the six dimensions asked to line chief B, regardless of whether line chief B reported a connection to line chief A.

Day n) indicates that the line chief starting the style shares social ties with the most senior line chief who has so far produced the style, and the second is an interaction of this variable with whether the style start should have been treated. All results shown in this section use directly the reweighted data (results are qualitatively the same when using the unweighted data). In this specification, a positive but insignificant effect can be seen on the first day a line produces a style, while the effect turns negative, and still insignificant on the second day.

In column 2 the usual line-chief - style-day fixed effects are interacted with garment type fixed effects, as was already done in column 4 of Table 5. Now, the effect of the interaction becomes very large, 17 efficiency points, or 84 percent of the standard deviations of productivity on the first day a line produces a new style that has previously been produced on another line. However, this large effect comes against the backdrop of a large and negative (but insignificant) effect of being connected to the most senior line chief in general. This could point towards the effect being driven by a few influential observations, and it indeed seems that two observations with unusually high daily productivity values of over 150 efficiency points have an over proportional effect on the results. I thus drop all three observations with productivity of more than 150 efficiency points.²⁷ The size of the interaction term now drops to 11 efficiency points (and its p-value to 0.051), while the general effect of being connected to the most senior person who produced the style so far drops to 6 efficiency points, remaining insignificant.

Thus, there is some evidence that the effect of the treatment was stronger when the line chiefs providing and receiving the briefing shared social ties - if we condition the regression on line-chief - garment type fixed effects. This is in line with findings from the previous section that also the overall treatment effect was stronger when conditioned on these fixed effects. In fact, once an interaction term with social ties is included in the regression, the size of the general effect of the treatment remains close to the ones estimated without line-chief - garment type fixed effects. The increase in the overall effect we saw in Table 5 when including these fixed effects seems to be driven by those instances in which the briefing was done by a socially connected supervisor.

6 Who is being treated?

Given the positive estimated treatment effects and high returns of the intervention, the question remains as to why such a management routine has not been implemented earlier at the factories. Possible reasons fall into two broad categories: managers did not realize the benefits of the measure (Bloom et al. (2013); compare Hanna et al. (2014) on evidence of failure of micro-entrepreneurs to notice possible improvements in operations by focusing attention on a too-limited set of dimensions for possible improvements); and resistance by key workers, such as the line chiefs, to such a measure. Resistance by line chiefs could stem from at least two sources;

²⁷This is done after the data had already be cleaned of clear outliers with productivity of more than 200 efficiency points.

Table 6: Communication Intervention and Social Ties

VARIABLES	(1)		(2)		(3)	
	Efficiency	se	Efficiency	se	Efficiency	se
Treatment, Day 1	3.278**	(1.67)	4.051**	(1.82)	4.510**	(1.91)
Treatment, Day 2	3.452*	(2.04)	3.880*	(1.93)	4.140*	(2.09)
Treatment, Day 3	1.009	(1.46)	-0.139	(2.34)	0.180	(2.39)
Connected, Day 1	-1.972	(3.98)	-9.169	(5.33)	-6.658	(3.89)
Connected, Day 2	-1.765	(1.67)	-2.161	(2.93)	-2.122	(2.93)
Connected, Day 3	-3.507	(3.91)	-2.457	(3.24)	-2.355	(3.27)
Treatm. x Conn., Day 1	4.170	(4.85)	17.799**	(7.24)	11.855*	(5.51)
Treatm. x Conn., Day 2	-4.245	(7.47)	-3.132	(10.01)	-3.198	(9.93)
Treatm. x Conn., Day 3	2.457	(10.22)	1.569	(8.86)	1.469	(8.92)
Constant	42.169***	(0.00)	45.652***	(4.44)	39.500***	(5.05)
Observations	4,094		3,381		3,379	
R^2	0.438		0.657		0.667	
Start.Rank-Style.Day FE	YES		YES		YES	
LC-Style.Day FE	YES					
LC-Style.Day-Grmt.Type FE			YES		YES	
Week-Fact. FE	YES		YES		YES	

Notes: Table shows the results from regressing daily line productivity from the first three days a sewing line produces a new style it has not produced before (first three “style-days”) on a dummy that the style start should have been treated, interacted with style-day (“Treatment, Day n”), on a dummy that the line chief of the line reports social ties to the most senior line chief who already produced the style, interacted with style-day (“Connected, Day n”), and on an interaction term of the two for each style-day (“Treatm. x Conn., Day n”). It controls for fixed effects for the number of lines that already produced the style before, interacted with style-day, line chief fixed effects interacted with style day, and fixed effects on the week-factory level. Column 2 and 3 interact line-chief - style-day fixed effects with fixed effects for the type of the garment (t-shirt, polo, pant.). All results estimated using reweighted data according to DiNardo et al. (1996) to balance first day productivity in pre-treatment period (1st April - 20th May 2014). Standard errors obtained by wild-bootstrap, clustered on 17 units of randomization. Standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1.

first, an unwillingness to provide help to co-workers whom they might regard as competitors; and, second, status concerns - that is, the perception that accepting help might negatively affect how they are regarded by peers and superiors (Lee (1997); Bunderson and Reagans (2011)). If status concerns play a role, they could be amplified by the age and seniority of line chiefs; older and more senior line chiefs could especially dislike being briefed, especially by younger, less-experienced line chiefs.

I test the extent to which status concerns related to the age and seniority of the line chiefs play a role in shaping resistance to the intervention. I do this by exploiting the imperfect compliance in the implementation of the intervention, to see if certain line chief characteristics predict the likelihood that a line chief who should have received a briefing actually received one. I control the regression for the average productivity of the line chief, and for his or her productivity on the day the treatment should have occurred; this way, I test whether any effects of a line chief's age or seniority on the likelihood that he or she receives a briefing is confounded by the line chief being more productive (on average, or on the day of the briefing), and, thus, less in need of a briefing.²⁸ A remaining negative effect of age or seniority would indicate a reluctance on the part of factory managements to instruct line chiefs to brief their counterparts who are older or more senior. This would indicate status concerns among older and more-senior line chiefs. Anticipating resistance by line chiefs and poor outcomes of the intervention, factory management might therefore not attempt to send line chiefs to brief each other.

Table 7 below shows the results from a test of the extent to which line chief characteristics affect the probability of actually receiving a briefing if randomly selected to receive one. Among the sample of the 377 style starts which should have been treated, I regress a dummy variable for the style start being matched to one of the treatments reported in the log-books, on age and seniority of the line chief who should have received the briefing, using probit regressions. Given that age and seniority (years line chief already worked as line chief in factory) are highly correlated (corr. = 0.61), I include these terms separately in the regression. Column 1 of Table 7 shows the results when only using age and factory fixed effects. Age shows a significantly negative effect on the likelihood of receiving a briefing. Column 2 includes three controls for the productivity of the line chief: 1) productivity of his or her line on the day the new style is begun and, thus, the day the briefing should have occurred ("Productivity Day"); 2) overall average productivity of the line chief's line in all observations in the data ("Productivity Overall"); and 3) average productivity on the first days the line chief started a new style that had been previously produced on other lines during the two pre-intervention months, which should have not been affected by the intervention ("Productivity Pre-Treatm."). None of these three controls show any effect on the likelihood of receiving a briefing, while the effect of age remains unchanged with a significant negative effect.

²⁸Empirically, age or seniority of line chiefs is not correlated with the line chief's daily productivity, even when controlling for style fixed effects. The point estimate for the correlation is often negative, though mostly not statistically significant when clustering standard errors on the line chief level.

Column 3 replaces the factory fixed effect with fixed effects on the level of individual sewing floors, to test whether the effects are driven by the intervention being implemented more widely on floors with younger or less-senior line chiefs. The inclusion of these fixed effects makes the effect of age twice as large. This stronger effect within these groups of line chiefs who work in close proximity on the same floors supports the hypothesis that the age effects represent status concerns; any briefing is more likely to be noticed by other line chiefs from the same floor, and most briefings are conducted among line chiefs from the same floor. Thus, to some extent, floors can be considered self-contained laboratories in which the intervention was done, and it is the relative age within these floors that predicts the likelihood of receiving a briefing. Furthermore, as one would expect, the overall productivity of a line chief's line now also has a negative effect on the likelihood of receiving briefings; better-performing line chiefs are less likely to receive a briefing. However, this effect seems to operate independently of the effect of age on the likelihood of receiving a briefing.

Column 4 uses seniority (months already spent working as line chief in factory) instead of age; seniority similarly yields a negative effect. However, when including both age and seniority in one regression (Column 5), the effect is driven by age, even when including the line chief's education level and number of reported social ties as additional controls, which themselves have no significant effect.²⁹ Age dominates the effect, even though, a priori, age might have been expected to be less related to a line chief's productivity than seniority or education (even though none are empirically correlated with productivity in the data). This result again supports the hypothesis that the effect represents status concerns; the result is not confounded by these less-likely-to-be-treated line chiefs somehow being more able or productive, and, thus, less in need of the briefings.

Thus, the non-compliance in the implementation of the treatment provides clues as to why such an intervention might not have been implemented before at these factories. The briefings were less likely to be conducted for older and more-senior line chiefs. The data show that this proved to be the case even though these line chiefs are not more productive; controlling for a line chief's average productivity did not affect the results. Thus, factories seem to be more willing to implement the briefing intervention when younger line chiefs are on the receiving end of the briefings. This points towards the existence of status concerns, which are likely to be more pronounced among older line chiefs, and which interfere with the implementation of the intervention, as they might provoke resistance against receiving help which was not requested by the line chiefs themselves. As a result, factory management may refrain from implementing such interventions more broadly, despite their positive effects as shown in the randomized trial.

²⁹Education obtained from line chief surveys, using an IPA Bangladesh specific ordinal code capturing educational attainment.

Table 7: Who is being Treated?

	(1)	(2)	(3)	(4)	(5)
VARIABLES	Actual Treatment Recorded				
Age	-0.027** (0.01)	-0.031** (0.01)	-0.060*** (0.01)		-0.041** (0.02)
Seniority				-0.018** (0.01)	-0.008 (0.01)
Education					0.060 (0.10)
Nbr. Soc. Connections					0.003 (0.03)
Productivity Day		0.003 (0.01)	0.003 (0.01)	0.003 (0.01)	0.003 (0.01)
Productivity Overall		0.014 (0.03)	-0.102*** (0.04)	-0.083** (0.03)	-0.086*** (0.03)
Productivity Pre-Treatm.		-0.004 (0.01)	-0.006 (0.01)	0.005 (0.01)	-0.004 (0.01)
Constant	-0.966** (0.41)	-1.542 (1.37)	6.150*** (2.14)	3.477* (2.02)	3.958* (2.31)
Observations	361	299	265	266	265
Factory FE	YES	YES			
Factory-Floor FE			YES	YES	YES

Notes: Among sample of 377 style starts that should have been treated, a dummy for the treatment having been actually conducted according to experimental logbooks is regressed on characteristics of the line chief who should have received briefing, using probit regression. “Productivity Day” is productivity on day treatment should have occurred, “Productivity Overall” average productivity of line chief who should have received briefing, and “Productivity Pre-Treatm.” average first day productivity of line chief during two months before start of treatment. Standard errors clustered on line chief level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

7 Conclusion

This paper presents evidence from a randomized communication intervention in three garment factories in Bangladesh; the intervention aimed to reduce the costs of sharing knowledge about the production processes of garment styles produced on different sewing lines in the same factory. In the intervention, supervisors of randomly selected sewing lines receive a briefing when they started producing a style that was new to them but had already been produced by another line in the factory. The briefing was conducted by the most senior supervisor of a line that had previously produced the style. I show that productivity of lines increased whenever they were selected to receive a treatment. The increase in productivity was mainly visible on the first day a line produced a new style, while in some specifications also on the second day. Thus, the increase was visible during the steepest part of the average learning curve when lines start producing a new style. The effect was driven by a reduction of the number of style starts with extremely low productivity. This points towards the treatment having been most effective at, and possibly targeted towards, those style starts at which productivity would have otherwise been very low.

Given these estimated returns from the intervention, the puzzle remains as to why it was not previously implemented by the factories. Prior to the intervention, factory managers told us that they occasionally told line chiefs to ask other line chiefs for help on specific production processes; however, this routine was not institutionalized at any of the three factories. And while two factories did not report problems with the implementation of the intervention, the third factory reported that line chiefs complained about the need to give other line chiefs briefings. This could point towards a conflict between the beneficiaries of the interaction (the factories through increased productivity), and the bearers of the bulk of its costs (the line chiefs in terms of time spent providing the briefings). Furthermore, I find direct evidence that the factories were more likely to comply with the instructions to implement the briefings with all line chiefs starting new styles previously produced on other lines, if the line chief receiving the briefing was younger. This effect is not accounted for by lower performance of younger line chiefs in terms of productivity of their lines, nor by their actual experience as line chiefs or education levels. It thus supports the hypothesis that status concerns among older line chiefs invoke resistance against receiving the briefings. For this reason, factories might generally be less inclined to implement such knowledge exchange interventions among its workers, despite their positive estimated effects.

References

Amodio, F. and M. M. Carrasco (2015). *Input Allocation, Workforce Management and Productivity Spillovers: Evidence from Personnel Data*. McGill University.

- Arrow, K. (1962). The Economic Implications of Learning by Doing. *The Review of Economic Studies* 29(3), 155–173.
- Atkin, D., A. Chaudhry, S. Chaudry, A. Khandelwal, and E. Verhoogen (2015). Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan. NBER Working Paper 21417.
- Bandiera, O., I. Barankay, and I. Rasul (2005). Social Preferences and the Response to Incentives: Evidence from Personnel Data. *Quarterly Journal of Economics* 120(3), 917–962.
- Bandiera, O., I. Barankay, and I. Rasul (2010). Social Incentives in the Workplace. *Review of Economic Studies* 77(2), 417–458.
- Bandiera, O., I. Barankay, and I. Rasul (2011). Field Experiments with Firms. *Journal of Economic Perspectives* 25(3), 63–82.
- Bandiera, O., I. Barankay, and I. Rasul (2013). Team Incentives: Evidence From A Firm Level Experiment. *Journal of the European Economic Association* 11(5), 1079–1114.
- Bandiera, O., A. Prat, and R. Sadun (2015). Managing the Family Firm: Evidence from CEOs at Work. NBER Working Paper 19722.
- Bandiera, O. and I. Rasul (2006). Social Networks and Technology Adoption in Northern Mozambique. *Economic Journal* 116(514), 869–902.
- Banerjee, A., A. Chandrasekhar, E. Duflo, and M. Jackson (2013). The Diffusion of Microfinance. *Science* 341(6144).
- Benkard, L. (2000). Learning and Forgetting: The Dynamics of Aircraft Production. *American Economic Review* 90(4), 1034–1054.
- Bloom, Nicholas, B. E., D. McKenzie, A. Mahajan, and J. Roberts (2013). Does Management Matter: Evidence from India. *Quarterly Journal of Economics* 128(1), 1–51.
- Bunderson, J. S. and R. E. Reagans (2011). Power, Status, and Learning in Organizations. *Organization Science* 22(5), 1182–1194.
- Cai, J., A. D. Janvry, and E. Sadoulet (2015). Social Networks and the Decision to Insure. *American Economic Journal: Applied Economics* 7(2), 81–108.
- Cameron, A. Colin, J. G. and D. Miller (2008). Bootstrap based Improvements for Inference with Clustered Errors. *Review of Economics and Statistics* 90(3), 414–427.
- Conley, T. and C. Udry (2010). Learning about a new Technology: Pineapple in Ghana. *American Economic Review* 100(1), 35–69.
- Covert, T. (2015). Experiential and social learning in firms: the case of hydraulic fracturing in the Bakken Shale. Working Paper, University of Chicago.

- Das, S., K. Krishna, R. Somanathan, and S. Lychagin (2013). Back on the Rails: Competition and Productivity in State-owned Industry. *American Economic Journal: Applied Economics* 5(1), 136–62.
- DePaulo, B. and J. Fisher (1980). The Costs of Asking for Help. *Basic and Applied Social Psychology* 1(1), 23–35.
- DiNardo, J., N. Fortin, and T. Lemieux (1996). Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semi-Parametric Approach. *Econometrica* 64(5), 1001–1044.
- Dufo, E., M. Greenstone, R. Pande, and N. Ryan (2013). Truth-telling by Third-party Auditors and the Response of Polluting Firms: Experimental Evidence from India. *Quarterly Journal of Economics* 128(4), 1499 – 1545.
- Falk, A. and A. Ichino (2006). Clean Evidence on Peer Effects. *Journal of Labor Economics* 24(1), 39–57.
- Foster, A. and M. Rosenzweig (1995). Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture. *Journal of Political Economy* 103(6), 1176–1209.
- Foster, L., J. Haltiwanger, and C. Syverson (2008). Reallocation, Firm Turnover and Efficiency: Selection on Productivity or Profitability. *American Economic Review* 1(98).
- Granovetter, M. (1973). The Strength of Weak Ties. *American Journal of Sociology* 78, 1360–1380.
- Granovetter, M. (1995). *Getting a Job: A Study of Contacts and Careers*. Chicago, 2.Ed: University of Chicago Press.
- Hamilton, B., J. Nickerson, and H. Owan (2003). Team Incentives and Worker Heterogeneity: An Empirical Analysis of the Impact of Teams in Productivity and Participation. *Journal of Political Economy* 111(3), 465–497.
- Hanna, R., S. Mullainathan, and J. Schwartzstein (2014). Learning Through Noticing: Theory and Experimental Evidence in Farming. *The Quarterly Journal of Economics*.
- Heath, R. (2015). Why do Firms Hire using Referrals? Evidence from Bangladeshi Garment Factories. University of Washington.
- Hendel, I. and Y. Spiegel (2014). Small Steps for Worker, a Giant Leap for Productivity. *American Economic Journal: Applied Economics* 6(1), 73–90.
- Hjort, J. (2014). Ethnic Divisions and Production in Firms. *Quarterly Journal of Economics* 129(4), 1899–1946.
- Irwin, D. and P. Klenow (1994). Learning-by-Doing Spillovers in the Semiconductor Industry. *Journal of Political Economy* 102(6), 1200–1227.

- Kato, T. and P. Shu (2011). Competition, Group Identity, and Social Networks in the Workplace: Evidence from a Chinese Textile Firm. IZA Discussion Paper 6219.
- Kellogg, R. (2011). Learning by Drilling: Interfirm Learning and Relationship Persistence in the Texas Oilpatch. *The Quarterly Journal of Economics* 126(4), 1961–2004.
- Lazear, E. (2000). Performance Pay and Productivity. *American Economic Review* 90(5), 1346–1361.
- Lee, F. (1997). When the Going Gets Tough, Do the Tough Ask for Help? Help Seeking and Power Motivation in Organizations. *Organizational Behavior and Human Decision Processes* 72(3), 336 – 363.
- Lee, F. (2002). The Social Costs of Seeking Help. *The Journal of Applied Behavioral Science* 38(1), 17–35.
- Levitt, S., J. List, and C. Syverson (2013). Towards and Understanding of Learning by Doing: Evidence from an Automobile Plant. *Journal of Political Economy* 121(4), 643–681.
- Lucas, R. (1993). Making a Miracle. *Econometrica* 61(2), 251–272.
- Munshi, K. (2004). Social Learning in a heterogeneous population: Technology Diffusion in the Indian Green Revolution. *Journal of Development Economics* 73, 185–213.
- Nagin, D., J. Rebitzer, S. Sanders, and L. Taylor (2002). Monitoring, Motivation, and Management: The Determinants of Opportunistic Behavior in a Field Experiment. *American Economic Review* 92(4), 850–873.
- Thompson, P. (2010). Learning by Doing. In B. Hall and N. Rosenberg (Eds.), *Handbook of Economics of Technical Change*. North-Holland: Elsevier.
- Thompson, P. and R. Thornton (2001). Learning from Experience and Learning from Others: An Exploration of Learning and Spill-overs in Wartime Shipbuilding. *American Economic Review* 91(5), 1350–1368.
- Wright, T. (1936). Factors affecting the Cost of Airplanes. *Journal of Aeronautical Sciences* 3(4), 122–128.

Appendix A: Worker Movement

While sewing worker are allocated to fixed lines, they do at times switch lines on a day-to-day basis to replace absent workers. With average daily absenteeism rates in the sample factories non-negligible at 3 percent to 5 percent, there is scope that within these reallocations, enough workers with relevant production knowledge on specific styles are moved across lines to drive the observed increases in productivity when other lines have already produced the same style

Table 8: Worker Movement

VARIABLES	Efficiency
Additional Line	3.249*** (0.838)
Constant	43.353*** (3.640)
Observations	18,749
R^2	0.513
Factory-week FE	YES
L.Chief-Style.Day FE	YES

Notes: Table shows result of regressing in sample of first lines in factory to produce a style daily line productivity on a dummy indicating that by that day further lines have started producing same style as well (“Additional Line”). Regressions control for factory-week fixed effects and line chief fixed effects interacted with style-day. Standard Error clustered on line chief level in parantheses: *** p<0.01

before. This would indicate that productivity spill-overs to later lines producing the same style are driven by within-worker transfers of production knowledge across sewing lines, and not by knowledge exchange across workers.

To test for the likelihood that short-term movements of workers indeed account for the higher productivity if other lines have previously produced the style, a tentative test is conducted. If workers with experience on the styles are reallocated across lines, the lines from which the workers are taken should experience a negative effect on their productivity if they still produce the style when further lines start producing the style as well, and workers are reallocated away to these lines. Table 8 below shows the results when regressing the daily productivity of first lines that produced a style in the factory on line chief - style-day fixed effects, factory-week fixed effects, and on a dummy indicating that by that day, additional lines (“Additional Line”) have also started producing the same style. The results show that instead of a drop in productivity, if anything, the first line experiences an increase in productivity when other lines also start producing the style. The source of these positive effects are not immediately clear. These effects could be due to reverse knowledge spill-overs from the additional lines back to the first line, or due to other forms of peer effects, such as competition. However, these results are not in line with what could be expected if systematic movement of workers with production experience on certain styles would cause the observed productivity spill-over.

Appendix B: Results using Recorded Treatments

The main analysis focuses on the Intention-to-Treat effect, assuming that any style start which was randomly selected for treatment was treated. This appendix presents results when regressing productivity on the first three days a line produces a new style on a dummy variable (interacted with style-day fixed effects) indicating that according to the logbooks the line chief received a briefing for the style by another line chief who already produced it. As already mentioned in the main part of the paper, the production data records 377 starts of new styles on treatment lines which should have been treated, while 98 of these starts could be matched with recorded treatments in the log-books. However, given that the factory managements admitted that not all treatments were recorded in the logbooks, the number of actual treatments is likely to be much larger. On the other hand, 26 further style starts at which no other line chief had produced the same style before were recorded as treated, indicating that a certain share of actual treatments has occurred on style starts not intended to be treated. Note, however, that there is no recorded treatment on control lines, or on dates outside the time the intervention was implemented at the factories. Thus, when dropping style starts from the sample in which no other line has yet produced the style (the same sample as used in column 3 and 4 of Table 5), the only type of non-compliers left in the sample are style starts which should have been treated but were not.

Column 1 of Table 9 shows the results using the usual sample of daily sewing line productivity observations from the first three days a line produces a new style from January through September 2014, excluding observations from style starts at which no other line has produced the style before. It regresses these productivity observations on an indicator that the style start had been treated according to the logbooks, controlling for the usual start-rank and line chief fixed effects (interacted with style-day) and factory-week fixed effects. In this specification, no effect of recorded treatment can be observed. When interacting the line chief fixed effects with garment type fixed effects in column 2, a positive but insignificant effect (at least on style-day 1 and 3) of actual treatment can be seen. This is in line with also the stronger effects of the intention to treat effect, and its interaction with social ties, that we see when interacting line chief and garment-type fixed effects. Finally, when controlling for past output of the style on other lines not by fixed effects for the number of lines that already produced the style before, but by the log cumulative output of the same style on other lines before, the effect become even larger, but remaining insignificant (column 3).

The insignificant and small (when not using line chief - garment type fixed effect) estimated effects of recorded treatments are puzzling in light of the significant intention-to-treat effects. The following points could provide explanations for this pattern. First, as indicated by the factory managements, many more style starts were likely treated than indicated as “treated” in the regressions from this section. This would dilute the estimated effect of the treatment, especially if the likelihood of the treatment of not being recorded was uncorrelated with the effect of the particular treatment on productivity. Second, actually administered treatments

Table 9: Results using Recorded Treatments

VARIABLES	(1)		(2)		(3)	
	Effic.	se	Effic.	se	Effic.	se
Treatment, Day 1	1.766	(2.50)	3.944	(5.98)	5.215	(3.98)
Treatment, Day 2	-0.477	(0.67)	3.937***	(0.00)	4.376**	(1.71)
Treatment, Day 3	2.599*	(1.33)	5.570	(3.96)	4.675**	(2.18)
Constant	88.38**	(37.96)	39.82	(31.75)	51.92**	(25.55)
Observations	2,942		2,441		2,441	
R^2	0.485		0.691		0.680	
Start.Rank-Style.Day FE	YES		YES			
Ln Output Otherl. * Style.Day					YES	
LC-Grmt.Day FE	YES					
LC-Grmt.Day-Grmt.Type FE			YES		YES	
Week-Fact. FE	YES		YES		YES	

Notes: Table shows results when regressing line productivity on the first three days a line produces a new style already produced on another line before on a dummy variable for a style start having been matched to a reported treatment in the logbooks, interacted with style-day. Column 1 controls for line-chief fixed effects interacted with style-day, while column 2 and 3 interact these fixed effects with garment type fixed effects. Columns 1 and 2 control for fixed effects for number of lines that previously produced style, while column 3 controls for log previous output on other lines, each time interacted with style-day. All regressions control for start-rank fixed effects interacted with style-day and factory-week fixed effects. Standard errors clustered on floor level in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

were probably directed towards style starts at which the factory management expected a poor initial productivity. Therefore, the estimated effect of the recorded treatments is small, as the right counterfactual is not the productivity at other untreated style starts on the same line or on control lines, but the (unobserved) productivity if these treatments would not have been conducted. This also fits with the reduction in the left tail of the productivity distribution on treated lines with the onset of the treatment as shown in 8.

Appendix C1: First Day Productivity by Factory

Figures 5 - 6 combined first day productivity from Factories 1-3 from January through September 2014, though data from Factory 3 becomes available only from April 2014 onwards, resulting in smaller standard errors from that month onwards in the graph. To disentangle pre-trends in first day productivity from possible composition effects due to the inclusion of data from Factory 3 from April 2014 onwards, Figure 10 replicates Figure 6 separately for Factory 1 and 2 on the left-hand side, for which data throughout January to September is available, and for Factory 3 on the right-hand side, with data starting in April 2014 for this factory. The graphs show that a difference-in-differences treatment effect is visible in both of the two separate samples from

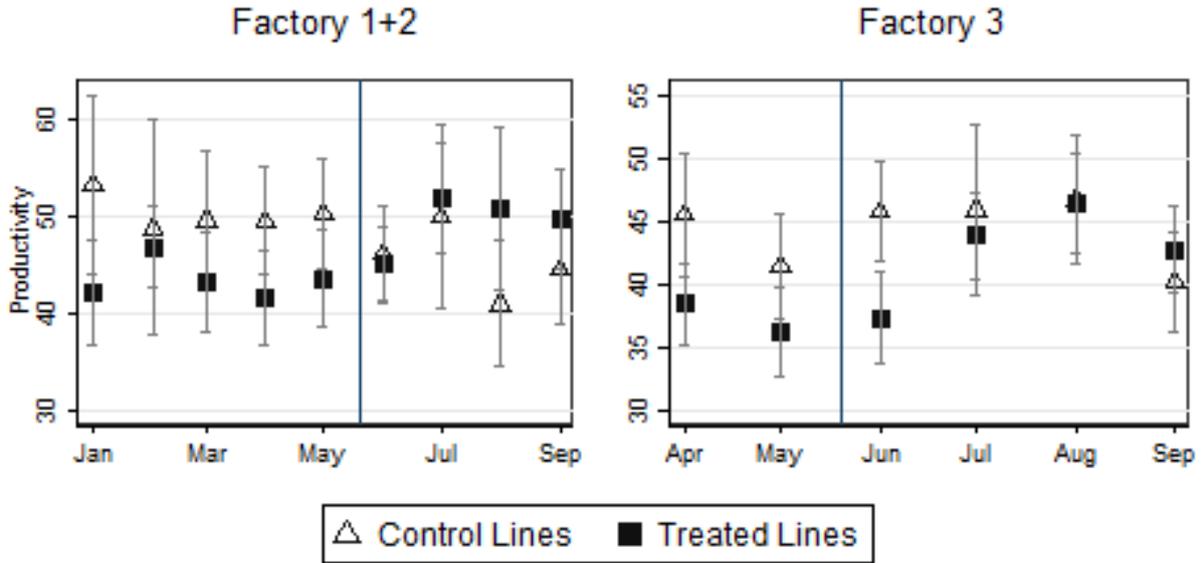


Figure 10: Pre-Post Intervention Start Trends for First Day Productivity, Factory 1+2 and Factory 3 separately. Graph shows average monthly productivity of lines on the first day they start producing a new style, separately for lines selected for treatment (solid squares) and lines not selected (hollow triangles). Left hand graph shows productivities for Factory 1 and 2 only, for which data throughout January to September 2014 is available, while right hand side graph shows productivities for Factory 3, for which data is available from April 2014 onwards. The solid vertical lines indicate start of treatment from June 2014 on. Capped bars represent 95% confidence intervals.

Factory 1 and 2, and from Factory 3.

Appendix C2: Application of Reweighting Approach from DiNardo et al. (1996)

The implementation of the approach of DiNardo et al. (1996) requires the estimation of two probit models; first, of a dummy indicating whether a unit i in the sample is selected for treatment ($T_i = 1$) on the unbalanced variable z_i , and, second, of a dummy indicating whether the unit is selected as control ($T_i = 0$) on z_i . The predicted probabilities $P(T = 1|z_i)$ for each unit i for being in the treatment group, and $P(T = 0|z_i)$ for being in the control group conditional on the unbalanced variable z , and the unconditional probabilities $P(T = 1)$ and $P(T = 0)$ of being selected into treatment or control sample, respectively, are then used to calculate weights w_i for each unit i according to:

$$w_i = \frac{P(T = 0|z_i)P(T = 1)}{P(T = 1|z_i)P(T = 0)} \quad (3)$$

To implement the approach, I first regress on a sample of all sewing lines a dummy indicating

that a sewing line is located on a treatment floor, on the line’s average productivity on the first days it produced new styles that have already been produced on other lines before, during the pre-intervention time April and May 2014. I control for factory fixed effects. The predicted values of this regression for each sewing line yield $P(T = 1|z_i)$ for calculating the weights w_i , according to equation 3. Similarly, I also regress a dummy indicating that a line is located on a control floor on its average first-day productivity during April and May 2014, to obtain $P(T = 0|z_i)$. I then reweight all observations from the treatment lines with weight w_i for the respective line (control units are not reweighted in this approach, therefore weights w_i for lines from control floors are set to 1).

Appendix D: Further Results

Table 10 in this Appendix presents further results on the communication intervention. Columns 1-3 show the results from a placebo test. The intervention was designed to target style starts at which another line had already produced the same style before. Only in these cases is another line chief available who had already experience with the style and could provide the briefings. Thus, the average productivity of lines starting to produce a new style which no other line chief has previously produced should not be affected by the intervention. This allows the implementation of a placebo test, in which I replicate the specification of column 1, Table 5, but adding an additional “*Treatm.Date&Line*” dummy, interacted with style-day, which takes value one at any style start on treatment lines during the time in which the treatment was implemented, regardless of whether another line has already produced the style before or not. If this variable captures the increase in productivity instead of the usual “*Treatment*” variable, which only takes value one at style starts at which another line did produce the style before, this could be indicative that the increase in productivity is due to a general increase in (first day) productivity on treatment lines with the onset of the intervention, as opposed to an actual treatment effect.

Column 1 of Table 10 uses the non-reweighted data. And, indeed, the term capturing all style starts shows a slight positive, but completely insignificant effect. However, this inclusion does reduce the size and the statistical significant of the main treatment effect. Note, however, that style starts at which no other line chief has produced the style before might ultimately not be the cleanest comparison group, as the logbook data reveals 26 instances in which also those style starts have been treated, presumably by line chiefs who produced very similar styles before. Thus, the intervention might have also have had some effect on first style starts, even if this was not intended in the initial design. Column 2 uses the reweighted data according to DiNardo et al. (1996). Now the treatment effect on non-first style starts is of similar size as in the main results from Table 5, even though still remaining insignificant. The additional control term for all style starts on treatment lines during the time of the treatment shows as relatively precisely estimated zero effect. The results remain unaffected when including style fixed effects, as in column 3. This indicates that the treatment effect in the main results is indeed driven by non-first style starts which were targeted for treatment, and not all style starts on treatment

lines during the treatment time.

Finally, column 4 of Table 10 shows the standard experimental results when including as additional controls log previous output of the same style on all other lines and on lines from the same floor, and log output of the same style from lines starting producing it on the same day. These are the main regressors included in Column 5 of Table 3, demonstrating the ambiguity of whether the effect of previous output on productivity is driven by learning effects or other forms or peer effects. Column 4 shows that the treatment effects are not affected by these additional controls.

Appendix E: Social Networks in the Factories

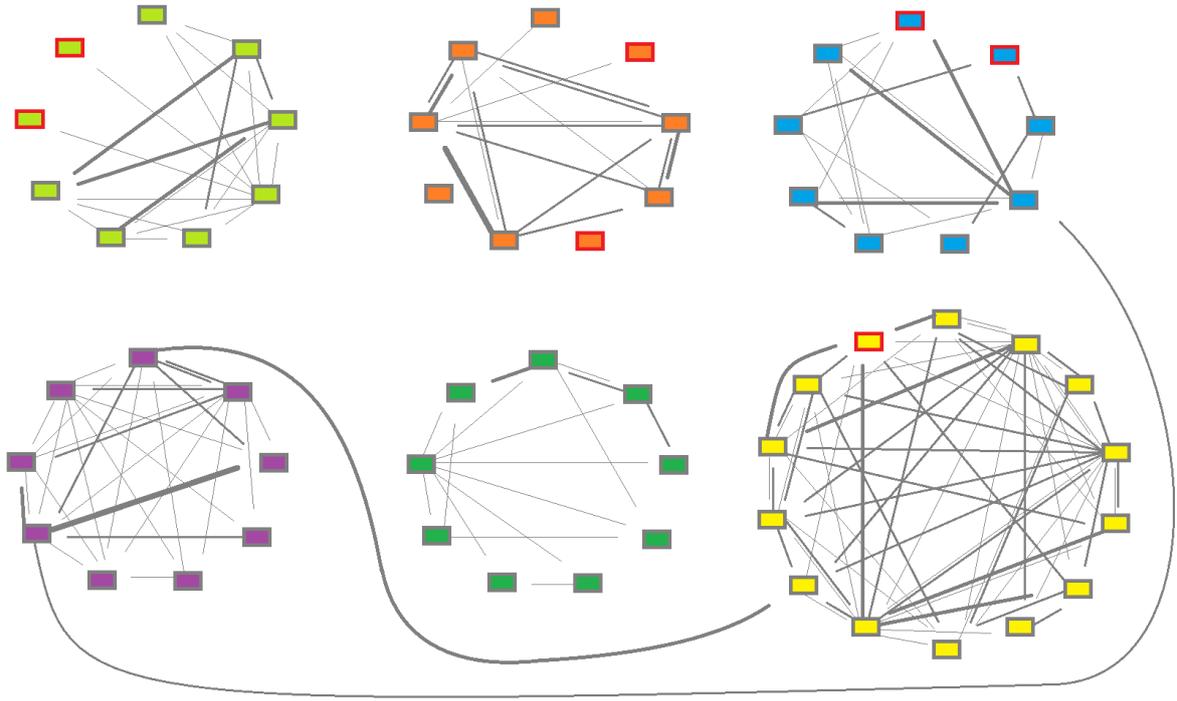
All line chiefs were asked to which other line chiefs in the factory they share social ties on the following six dimensions: being related, knowing each other from before working at the factory, having worked together at another factory before, visiting each other at home, spending lunch breaks together, and being friends. The three subsequent graphs depict the network information among line chiefs collected at the three factories in the sample. Nodes with same filling colour represent line chiefs on the same sewing floor. Thickness of links represents number of dimensions connection was reported on. Red frames of nodes represent line chiefs not surveyed due to being absent on the survey day. All graphs show directed links; the link is mentioned by the node which ‘touches’ link, and to the node which does not touch the link.³⁰

³⁰The network data is described in more detail, and its formation process analysed, in Chapter four of my PhD Thesis, “Social Network Formation in the Workplace: Evidence from Four Bangladeshi Garment Factories” (University of Warwick)

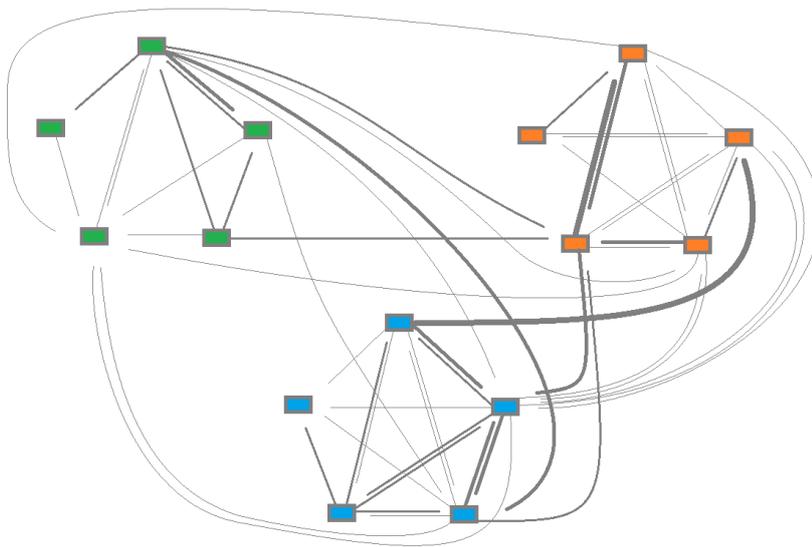
Table 10: Further Experimental Results

VARIABLES	(1)		(2)		(3)		(4)	
	Triple Difference		Reweighted Difference-in-Differences		Reweighted Difference-in-Differences		Reweighted Table 3 Contr.	
	Effic.	se	Effic.	se	Effic.	se	Effic.	se
Treatment, Day 1	3.563	(3.51)	5.038	(5.17)	7.226	(6.69)	5.973**	(2.32)
Treatment, Day 2	2.356	(2.57)	1.826	(5.50)	2.415	(5.32)	2.911	(2.32)
Treatment, Day 3	1.534	(3.01)	0.974	(3.33)	-0.362	(5.78)	0.600	(1.36)
Treatm. Date&Line, Day 1	2.920	(3.54)	-1.241	(5.46)	-0.988	(4.11)		
Treatm. Date&Line, Day 2	0.449	(8.95)	0.800	(3.01)	1.207	(3.25)		
Treatm. Date&Line, Day 3	-0.006	(0.03)	0.321	(1.59)	2.005	(3.62)		
<u>PREVIOUS OUTPUT:</u>								
All Lines, Day 1							-0.533	(0.94)
All Lines, Day 2							-1.119**	(0.54)
All Lines, Day 3							-0.814	(0.62)
Same Floor, Day 1							0.049	(0.12)
Same Floor, Day 2							-0.057	(0.22)
Same Floor, Day 3							-0.212	(0.26)
<u>OUTPUT FIRST DAY:</u>								
All Lines, Day 1							0.670	(0.43)
All Lines, Day 2							0.123	(0.36)
All Lines, Day 3							-0.372	(0.30)
Same Floor, Day 1							0.558	(0.43)
Same Floor, Day 2							0.927**	(0.40)
Same Floor, Day 3							0.910***	(0.00)
Constant	70.48***	(0.00)	42.89***	(0.00)	49.18***	(0.00)		
Observations	4,384		4,094		4,094		4091	
R ²	0.426		0.436		0.729		0.736	
Start.Rank-Style.Day FE	YES		YES		YES		YES	
L.Chief-Style.Day FE	YES		YES		YES		YES	
Style FE					YES		YES	
Week-Fact. FE	YES		YES		YES		YES	

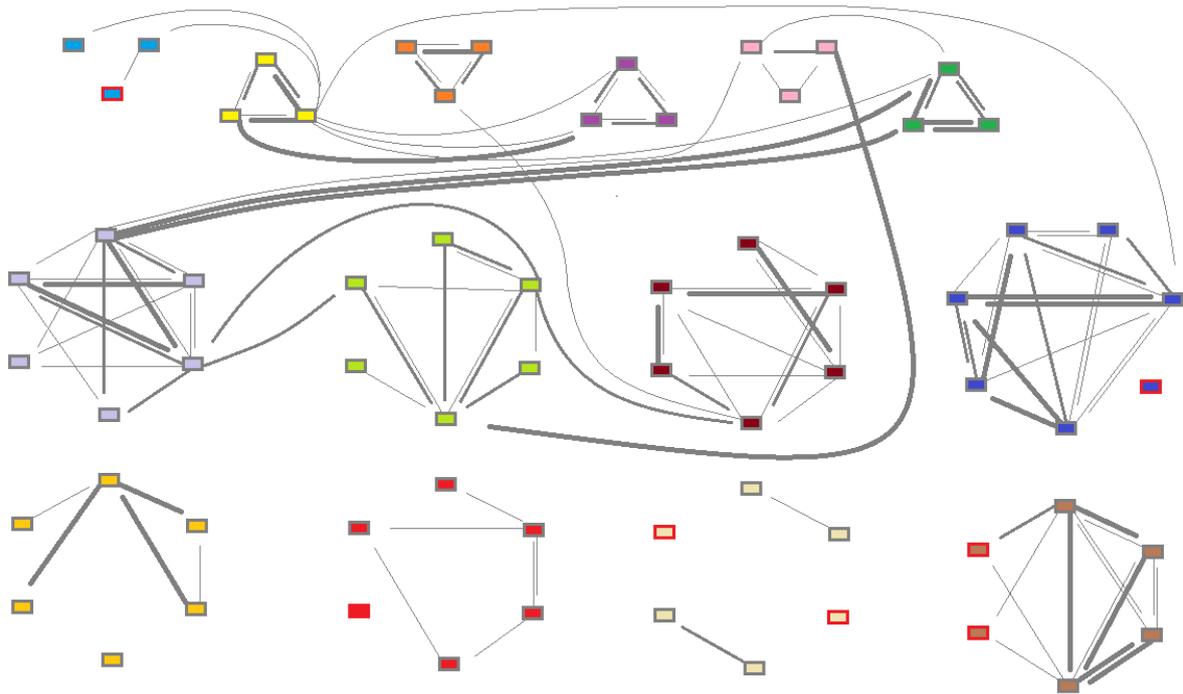
Notes: Table shows results from regressing line productivity on first three days line produces a new style on dummy indicating that style start should have been treated, interacted with style-day (“Treatment, Day n”), and on dummy indicating line starting style is on a treatment floor and style start occurred during time of treatment (“Treatm.Date&Line, Day n”). Column 4 controls for log previous output of style on other lines and other lines on same floor, and log output from lines starting style on same day, and those lines from the same floor, all four interacted with style day. All Regressions control for line chief fixed effects and fixed effect for number of lines that already produced style, both interacted with style-day. Regressions also include factory-week fixed effects, and columns 3 and 4 style fixed effects. Standard errors clustered on floor level in parentheses: *** p<0.01, ** p<0.05, * p<0.1



Factory 1



Factory 2



Factory 3