Working Paper Series 607
(ISSN 1211-3298)

Knowledge Exchange and Productivity Spill-overs in Bangladeshi Garment Factories

Andreas Menzel

CERGE-EI
Prague, November 2017

ISBN 978-80-7343-414-4 (Univerzita Karlova, Centrum pro ekonomický výzkum a doktorské studium)
ISBN 978-80-7344-450-1 (Národohospodářský ústav AV ČR, v. v. i.)

Knowledge Exchange and Productivity Spill-overs in Bangladeshi Garment Factories

Andreas Menzel*†
CERGE-EI, Prague

November 2017

Abstract

Productivity spill-overs within firms have commonly been used as a proxy measure for organizational learning. Using novel data from more than 200 production lines in three garment factories in Bangladesh, this paper extends the evidence on such productivity spill-over in two directions. First, I find that spatial distance within firms matters greatly for the strengths of productivity spill-overs, while product complexity matters little. This has important implications for firms in rapidly developing countries such as Bangladesh, as spill-over strength seems less affected when firms upgrade to more complex products, but seems more affected if firms grow larger. Second, I provide evidence from a randomized communication intervention in the three factories to determine the extent to which productivity spill-overs are indeed a measure of knowledge exchange within firms, and not of other types of peer effects, such as competition. In the intervention, randomly selected line supervisors were instructed by their superiors to share production knowledge when their lines were allocated the same garment for production. The intervention increased the strength of the productivity spill-overs between the targeted production lines. It thus supports the view that productivity spill-overs can be used as a measure of knowledge exchange within firms.

Keywords: Learning, Productivity, Firms

JEL Code: D2, L2, M5, O3

*I am grateful to Chris Woodruff and Rocco Macchiavello for their continuous support and advice throughout this project. I furthermore would like to thank Eric Verhoogen, Emily Breza, Jonas Hjort, Steven Levitt, Francois Gerard, Suresh Naidu, Cristian Pop-Eleches, Bernard Salanie, Oriana Bandiera, Florian Englmaier, Robert Akerlof, Clement de Chaisemartin, Sascha O. Becker, Roland Rathelot, Andrew Oswald, Stepan Jurajda, Michal Bauer, Nikolas Mittag, and seminar participants at Columbia University, CERGE-EI, University of Warwick, Urbana Champaign LER, CUNEF, as well as at the RCEA Workshop Rimini 2015, Oxford Young Economist Workshop 2014, PEDL Workshop London 2016, Italian Development Association Summer School 2016, Belgrade Young Economist Conference 2014, and the LMU Workshop Ohlstadt 2014 and 2016 for many useful comments. Financial help by PEDL is acknowledged.

†Address: R.330 - CERGE-EI, a joint workplace of Charles University and the Economics Institute of the Czech Academy of Sciences, Politickych veznu 7, P.O. Box 882, 111 21 Prague 1, Czech Republic. Email: Andreas.Menzel@cerge-ei.cz

1 Introduction

Organizational learning, "the creation, retention and transfer of knowledge within organizations" (Argote (2013)), has long been viewed as a major driver of firm productivity growth, and thereby also of overall economic growth (Arrow (1962); Lucas (1993)). Given the inherent difficulty to observe these processes in firms, the literature has usually approximated the study of organizational learning with the study of productivity spill-over within firms. Such spill-overs, defined as instances of increased productivity of production units (e.g. workers, teams, shifts) if other such units in the same firm have previously produced the same product, were interpreted as evidence for knowledge exchange between workers in these firms. A substantial literature has documented the existence of such spill-over in different sectors (Levitt et al. (2013); Darr et al. (1995)), studied the extent to which they occur between products that are technologically differentiated to various degrees (Thompson and Thornton (2001); Egelman et al. (2016)), and to what extent their strength depends on how long ago the other units had previously produced the product (Benkard (2000); Thompson (2007); David and Brachet (2011)).

However, our understanding of productivity spill-over remains incomplete, particularly along dimensions which are important as economies grow. Can the strength of internal spill-overs be maintained as firms grow in size as well as in organizational complexity? Can it be maintained, or does it even strengthen as firms upgrade to technologically more complex products when economies grow (Hidalgo et al. (2007); Kaldor (1967))? A new dataset collected from three large factories in a rapidly growing export oriented manufacturing sector, the Bangladeshi garment sector, allows me to address these questions. The three factories are organized into more than 260 parallel production lines and produce garments ordered by a large number of different international buyers. I find that the spatial proximity of production units within firms matters greatly for the strength of spill-overs; productivity spill-overs between production lines producing the same garment are much larger between lines located on the same production floors as compared to between lines located on different floors. This suggests that as firms or establishments grow in size, they loose their ability to transfer knowledge effectively between all employees, providing an explanation as to why young small firms often exhibit higher productivity growth (Haltiwanger et al. (2016)). On the other hand, I find no effect of product complexity, measured in terms of the labor input requirement to produce a specific garment (holding capital constant), on spill-over strength. This suggests that as factories in countries like Bangladesh upgrade to more advanced goods, they would not lose their (or gain) capacity to profit from organizational learning.

Furthermore, the data also allows me to shed light on a more fundamental question the literature has had difficulties to address, namely the extent to which observed productivity spill-over can indeed be considered a proxy measure for knowledge exchange. Increased productivity in response to previous output of the same product by others in the firm could also be caused by other types of peer effects that do not involve knowledge sharing, such as competition or benchmark setting effects.[1] The data allows me to exploit random variation in the (potential for) knowledge exchange between workers on different production lines induced by a management intervention implemented on randomly selected floors at the three factories, to test the extent to which knowledge exchange is indeed the driver of productivity spill-overs in this setting. During the four-month intervention period, whenever a production line located on a selected floor began producing a new garment style that had previously been produced by another line in the factory, the supervisor of the line that had already gained experience with the garment was sent by his superiors to the supervisor of the line now starting it, to brief him for 15-30 minutes on the most important production problems that had to be overcome on the earlier line. This intervention roughly doubled the strength of the production spill-over on the selected floors, showing that increased knowledge exchange leads to increased productivity spill-over. It thereby supports the long held view in organizational economics that productivity spill-over can be used as a proxy measure for organizational learning.[2]

Figure 1 below describes the nature of the productivity spill-overs we observe in these factories. As mentioned above, the factories are organized into a total of 260 parallel production lines, located on 23 floors. Production lines are self-contained production units, designed such that the total sewing process for a garment can be completed on one line. However, due to large order sizes, most orders (or "styles") are produced on more than one line, and in these cases, these lines typically start producing the same style on different days, after finishing previously allocated styles ("style" is the industry term for the different garment designs ordered by buyers). Thus, there is considerable (within-line) variation in the amount of a style already produced by other lines in the factories when a line starts producing a new style for the first time. Due to the fast-moving fashion industry and its seasonality, styles are technologically differentiated, which is reflected by line productivity dropping by a third on average on the first day lines produce a new style, as shown by the solid squares in Figure 1. Only on the fourth production day of the new style does line productivity reach its previous levels. The hollow circles in Figure 1, on the other hand, show how productivity evolves if another line on the same floor has previously produced the same style. While line productivity does not differ in the days before a line switches to producing the new style, the initial drop in productivity is reduced by almost 40 percent in these cases, and productivity remains higher for the first three production

[1] This 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 technology adoption as a proxy for learning. These studies mostly rely on estimating structural models or on placebo tests to separate learning from other peer effects. Covert (2015) uses mandatory production information, which oil fracking companies have to publish, to measure knowledge publicly available in the industry. He finds that firms use this information even though they put more weight on the experience gained in their own operations.

[2] In general, incentives to share knowledge or to compete with co-workers would be affected by workers' pay structure (Bandiera et al. (2005); Lazear (2000)). In all three factories in my sample, ordinary workers receive hourly pay, while line supervisors receive fixed monthly wages (irrespective of the hours of overtime worked). No workers in the factories receive piece rate payments or performance bonuses. However, workers and 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.

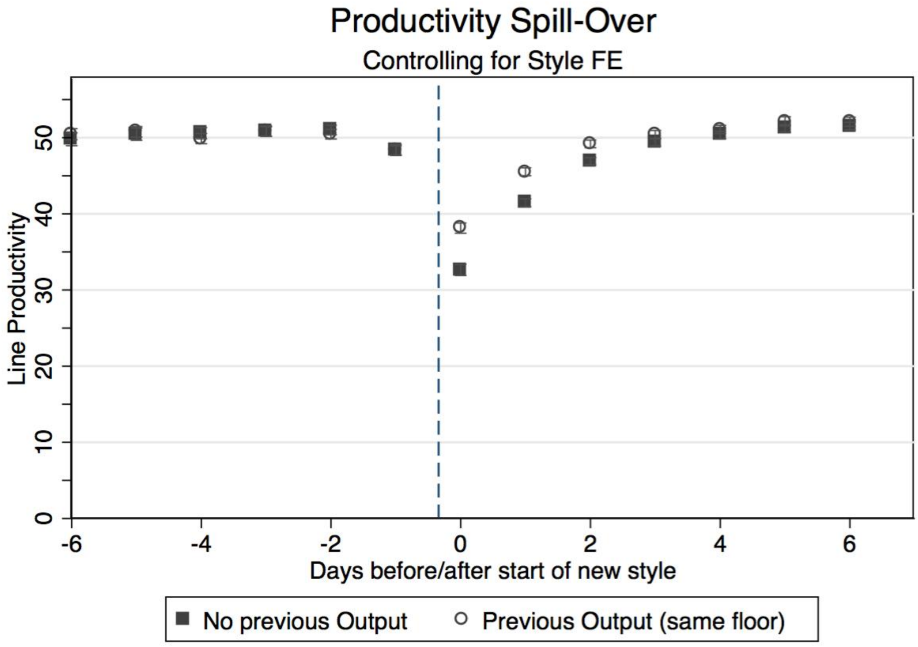

Figure 1: Sewing Line Productivity before and after the Start of a New Style. The graph shows average sewing line productivity in the days before and after switching to a new style. The vertical dashed line denotes the switch to the new style, and "Day 0" the first day of production of the new style.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

days, before it reaches its previous levels. It is this increased productivity during the first three to four days a line produces a new style if other lines have previously produced the same style that I refer to as the "productivity spill-over" and which I analyse in the remainder of this paper.

The rich nature of the dataset covering the production of more than 1,500 different garment styles allows me to disentangle more cleanly productivity spill-over from learning-by-doing effects than previous studies in this literature could, which mostly had data covering at most the production of a few dozen different products. Therefore, they usually related daily (or piece-wise) productivity of a production unit to the cumulative output of the same product already produced by other production units in the factory, controlling for cumulative previous output of the product on the same production unit. However, these two measures, cumulative output on the same and on other production units, are typically highly correlated over time, making it difficult to reliably measure their two distinct effects. The large dataset I have, on the other hand, allows me to use only a single observation per production line and garment style when the line starts producing the particular garment for the first time. Thus, I can hold previous output on the same line constant at zero, and focus on variation in cumulative output on other lines only.[3]

[3] What is more, (unobserved aspects of) capital and labor used in a production unit can evolve over time while the units produce a given product, in ways that most studies that rely on within-production-unit – within-product-variation in productivity have difficulties to control for (Thompson and Thornton (2001)). Again, by

Apart from the large size of the dataset and the availability of exogenous variation in the potential for knowledge exchange, the three factories also provide an ideal laboratory to study organizational learning for a number of different reasons. First, the factories work with a basic, labor-intensive, homogeneous technology. Thus, the results might be less dependent on idiosyncratic features of the production process, supporting their external validity. This study can therefore be considered complementary to most previous studies on productivity spill-over that focused either on high technology or capital intensive industries, such as electronics, car, ship, or airplane plants (Egelman et al. (2016); Benkard (2000); Thompson and Thornton (2001); Levitt et al. (2013)) Second, the production lines switch to a new style on average every 12 working days, each time going through the several-day learning process with reduced productivity. Thus, the learning effects are observed in a setting where the efficiency of the knowledge exchange process has a direct bearing on the overall productivity of the whole factory. Finally, a further advantage is the availability of accurate daily productivity data based on physical output on the sewing line level, which have been standardized across different styles using detailed garment complexity measures. 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)).

Apart from the literature on organizational learning mentioned above, the experimental section of this paper also contributes to a growing literature on experiments within firms to inform models of organizational behavior, often run in developing countries. Atkin et al. (2017) 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 the productivity of the firms. 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.[4]

using just one observation per line producing a given garment, this study circumvents this problem.

[4] By 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, sometimes referred to as insider econometrics. This term was introduced after the early study from Ichniowski and Shaw (1997), who study the effect of different human resource management practices on line productivity in US steel mills. Amodio and Carrasco (2017) exploit exogenous variation in worker productivity in a setting with quasi-team incentives, identifying free-rider effects among co-workers. Hjort (2014) is a case study of a Kenyan flower packaging factory showing that ethnically more diverse work teams have lower productivity. 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.

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. Section four provides more details on the randomized management intervention and shows its effects. Section five discusses possible reasons why, despite the positive effects found, the factories have not previously implemented such an intervention, while section six concludes.

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.[5] By local industry standards, the three factories involved in this study are 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, dress shirts and trousers, with 1,200 to 5,000 workers employed in their respective sewing departments. Table 1 provides key characteristics of the three factories. 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, but with each line having less than a third of the number of workers than a line in Factory 2. Another distinction is that factory 3 uses fewer supervisors (or "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 in 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, as can be seen in the relatively low standard deviation of workers per line within each factory in Table 1. Finally, workers in Factory 1 encounter 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 are sorted out and not counted in the line's output. The factories' central planning departments allocate workers, supervisors, and styles to specific sewing lines. Workers have fixed lines to which they are allocated, and usually change lines only as a 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

[5]Source: WTO, International Trade Statistics 2015: www.wto.org/english/res_e/statis_e/its2015_e/its15_toc_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	25	67	10
Nbr. Styles in Data	866	839	1048
Avg. Nbr. Lines /Style	3.11	1.49	3.94
Avg. Nbr. Days /Style & Line	16.4	9.5	8.5
Avg. Nbr. Workers /Line	30.9	72.8	23.2
S.Dev. Nbr. Workers /Line	8.0	11.1	6.9

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

reallocations of workers drive the observed productivity spill-overs across lines producing the same style. Furthermore, Appendix A presents the results of a placebo test which also suggests 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 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; daily production hours and number of workers for each line, 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 required labor input per piece under ideal production conditions, and is a proxy for the production complexity of the garment. Because the SMV is also essential in price negotiations with the buyers, its breakdown into the individual sewing operations is scrutinized by the buyer. Therefore, the SMV can be considered a reliable measure of product complexity, comparable across different styles.

To calculate the daily line-productivity measure, daily piecewise output is multiplied by the style-specific SMV, and then divided by total labor input on that line and day measured in worker-minutes:

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

I also conducted a survey of all line chiefs at the three factories, all of whom are around the age of 30, and only two of 128 line chiefs interviewed were female. They have worked as line chiefs on average for 1.8-3.5 years at the factories, and on average for more than one year on 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).[6]

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.[7] The reason for this lies in the high flexibility required in operations. Buyers place orders with low predictability, particularly regarding the order's specifications, and with close delivery deadlines. Furthermore, 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 in 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 in certain types of garments.

Lines also do not specialize in whether they are typically an earlier or 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.[8] According to

[6]Among ordinary workers, the 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 are promoted to these positions, with reasons discussed for this gender inequality in Macchiavello et al. (2015)

[7]An exception is 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.

[8]At 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.

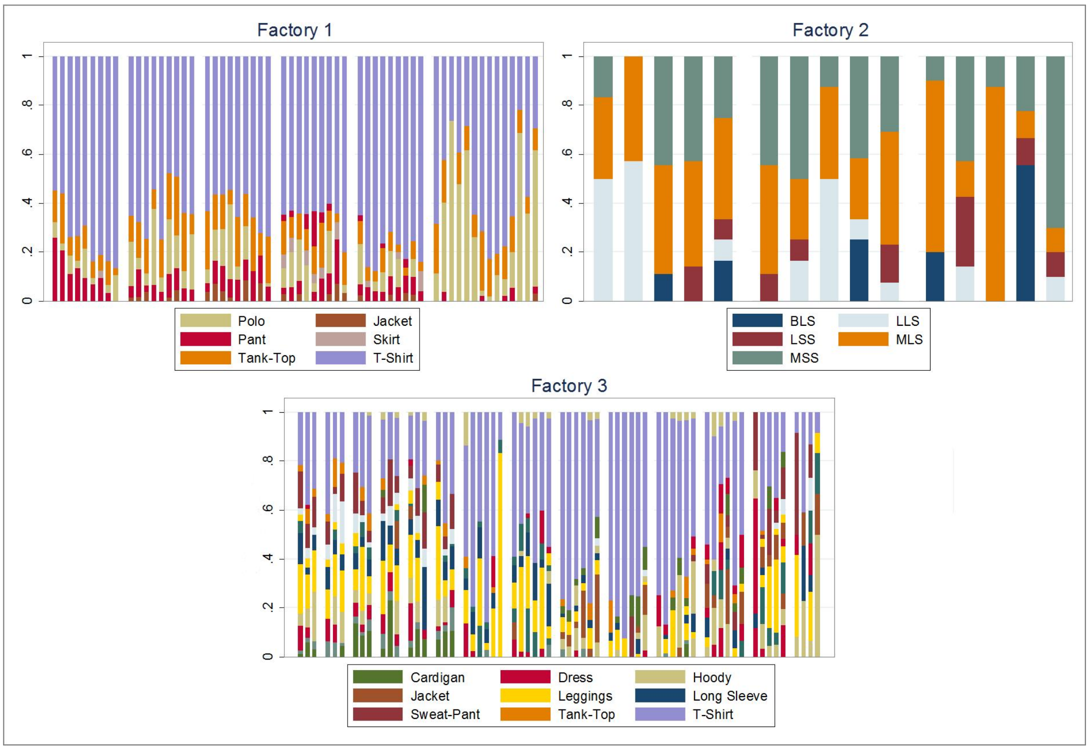

Figure 2: Garment Types Produced on Different Sewing Lines
The graphs represent the types of garments produced by different sewing lines at the three sample factories. Each bar in the graphs represents a sewing line, and the wider spaces between bars separate sewing lines from different sewing floors. The differently colored stacked parts of the bars represent different types of garments that the lines produced. Legends show colors for the 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 the graph parsimonious. The graph shows types for only 15 out of 17 lines for Factory 2.

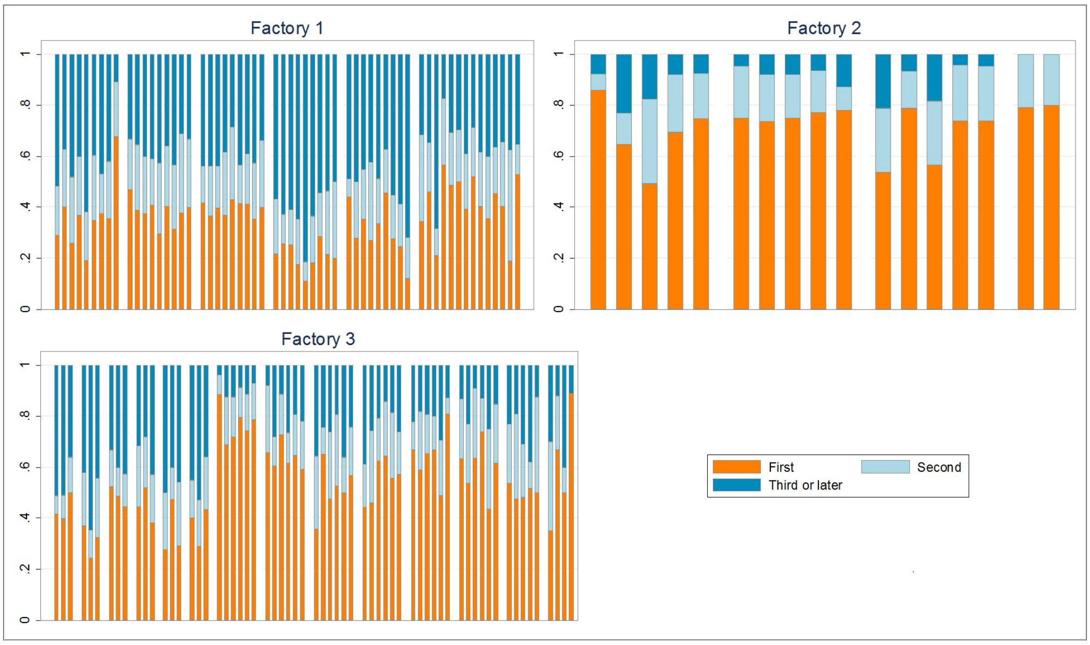

Figure 3: Start Ranks of Styles produced on different Sewing Lines
Graphs show for each sewing line in the three factories for which share of the styles they produce the lines are the first (orange), the second (light blue), or the third or later line (dark blue) to produce that style in the factory. 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 the graph parsimonious.

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. In any case, such selection concerns would only threaten our results if garment characteristics or the position in the roll-out of garments across lines would be correlated with time-variant characteristics of lines, as all results only exploit within line variation in productivity.

3 General Evidence on Productivity Spill-over

This section explores the extent to which line productivity profits from output of the same style previously produced on other lines. In the overall dataset, I observe 1,523 styles that have been produced on more than one sewing line in the factories. In total, there are 5,939 instances of a sewing line starting to produce one of these styles - a process I refer to as a “style start.” Therefore, these styles are produced on average on 3.9 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 following regressions, I keep the daily line productivity observations from the first ten days a line produces a new style. Denote from now on the n^{th} day a sewing line produces a style as the n^{th} "style-day." Thus, the sample that I use consists of all observations from style-days less or equal to ten, which amount to 58% of all observations on the day-line level from the original dataset.[9]

The basic econometric model I estimate in this section is as follows:

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

Productivity y_{fisnt} of sewing line i in factory f producing style s in month 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 first 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 ten style-days included in the sample to see for how long previous output affects the 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. Since Wright (1936), the literature on learning within firms has frequently used the log function to approximate knowledge with cumulative output, though alternative functional forms have been discussed, e.g. in Thompson (2007).[10]

In some regressions, I also include the output F_{isn} of the same style produced on all other lines on the same 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 ten 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 move from a sewing 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 manager, who could transfer knowledge from one line to other lines on his floor. 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, for 3,470 of the 5,939 style starts in the sample some other line has previously produced the style, while for 2,346 of these starts the style was already produced on a line on the same floor.

[9] At Factory 3, most line chiefs supervise two or four lines. As I am mainly interested in cross-worker spill-over and less in within-line chief spill-overs, and given that the management intervention studied in the next section is an intervention between line chiefs, I only keep the observations of a line chief producing a certain style from the first of his or her lines on which the line chief produced the style. 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 on which the line chief produced the style.

[10] More 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.

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 a deviation of line productivity from learning curves estimated separately for each line chief.[11] Furthermore, I include time fixed effects γ_{ft} on the factory - month - 'style-day' level. Standard errors in this section are clustered at the line chief level.[12]

General Results

Column 1 of Table 2 shows the results from estimating the empirical model from equation 1. While Panel 1 in that table always shows the effects with only total previous cumulative output in the factory, Panel 2 always adds the output from the same floor as an additional regressor. We can see that for each of the first 10 days a line produces a garment, previous output has a positive effect on line productivity, with the size of the effect falling from somewhat more than 0.6 to around half that size at the 10th day. The effect eventually becomes insignificant around the 12th day, which is not shown here due to space constraints. In Panel 2 below, log cumulative output from the same floor is added. Particularly during the first two to three production days, the effect of previous output is much stronger when having been produced on the same floor. After about seven days, it no longer makes a difference.

These effects, however, could be due to selection; e.g. styles that can be produced more efficiently are rolled out across more lines. Therefore, column 2 repeats column 1, but with style fixed effects interacted with dummies for each of the first ten 'style-days'. The effect of log overall previous output becomes smaller but remains significant at least until the 8th production day (Panel 1). When splitting up the effect in overall effect and from output on the same floor only, we can see that on the first production day, productivity is enhanced only when the style has previously been produced on the same floor, but not on others (Panel 2). On the second day the effect is statistically somewhat difficult to disentangle (both effects have a p-value of 0.1-0.12), while for the third and fourth production day, the effect is driven by overall previous output, without any difference if it was previously produced on the same floor. From the fifth day onwards, this regression no longer shows any significant separate effects for previous output, whether overall or from the same floor.

To better gauge the economic significance of these effects, Column 3 uses dummies for positive previous output of the same garment (on the same floor) instead of log previous cumulative output. The estimated effects are very similar to those in column 2. During the first two days of production, line productivity particularly increased if output has been previously produced

[11] The 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 do not differ qualitatively when using line fixed effects instead of line chief fixed effects.

[12] Results are qualitatively the same when clustering at the style level, or two-way on style and line chief level.

on the same floor, while relative to output produced on any floor in the factory, this no longer makes a difference by the fourth day of production. If another line has previously produced the same style, productivity is increased on average by 2.5 productivity units, or about a sixth of the average productivity drop of about 15 productivity units on the first day a line produces a new style. This effect is increased to about 3.5 productivity units if the garment has previously been produced on the same floor, or almost a quarter of the productivity drop of 15 units.

Knowledge Exchange vs other Spill-over Drivers

Most of the literature on organizational learning has 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 style; or, the productivity of some workers could provide the factory management with a benchmark against which it could compare the productivity of other workers and therefore more easily determine if workers slack. The fact that spill-over are mainly observed as long as line productivity is still in the learning curve is supportive of the knowledge exchange hypothesis. As long as the workers on the line have not yet gained the necessary production knowledge themselves, obtaining it from other workers could have direct effects on productivity. However, if the spill-overs are caused by mere differences in worker effort, e.g. enticed by competition, we could expect such effects to appear also outside the learning curve, as the scope for effort to affect productivity should not be constrained to the learning curve phase.

I also tested the hypothesis whether the spill-overs are stronger among more complex styles, as indicated by a higher SMV. A higher SMV usually indicates that a larger number of sewing processes have to be done to finish one piece of the garment, which could make it more likely that there are processes with which the workers are unfamiliar and for which they could profit from knowledge already gained by other workers. However, as Appendix B.1 shows, I do not find any systematic effect of style complexity on the strength of the spill-over. The interaction effects of SMV and previous output are mostly insignificant, and if not they do not have consistent signs.

To find more evidence of 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

Table 2: General Productivity Spill-Over & Peer Effects

	(1) Log Outp.		(2) Log Outp.		(3) Outp. >0		(4) Outp. >0		(5) Outp. >0	
Panel 1: All Factory										
Cumul. Previous Output x ...										
Day 1	0.613***	(0.08)	0.333***	(0.12)	2.554**	(1.00)			2.533**	(1.00)
Day 2	0.473***	(0.06)	0.259***	(0.08	2.146***	(0.67)			2.263***	(0.67)
Day 3	0.392***	(0.06)	0.293***	(0.07)	2.342***	(0.62)			2.357***	(0.59)
Day 4	0.436***	(0.06)	0.341***	(0.08)	2.632***	(0.67)			2.552***	(0.67)
Day 5	0.384***	(0.05)	0.186**	(0.08)	0.920	(0.73)			0.910	(0.74)
Day 6	0.332***	(0.06)	0.204***	(0.07)	1.342**	0.65)			1.338**	(0.66)
Day 7	0.284***	(0.06)	0.185**	(0.09)	1.423*	(0.74)			1.407*	(0.76)
Day 8	0.391***	(0.07)	0.266**	(0.10)	2.051**	(0.85)			2.027**	(0.86)
Day 9	0.194***	(0.07)	0.097	(0.08)	-0.059	(0.75)			-0.199	(0.76)
Day 10	0.371***	(0.08)	0.264**	(0.10)	1.466	(0.90)			1.418	(0.92)
Output Lines Starting Same Day x ...										
Day 1							11.971***	(1.88)	5.158***	(1.27)
Day 2							10.086***	(1.75)	3.904***	(1.08)
Day 3							6.698***	(1.55)	2.331**	(0.93)
Day 4							5.746***	(1.51)	0.802	(0.92)
Day 5							2.448	(1.51)	0.132	(0.92)
Day 6							1.485	(2.02)	0.403	(0.98)
Day 7							1.657	(1.87)	-0.372	(0.98)
Day 8							1.972	(2.18)	-1.179	(1.02)
Day 9							2.920	(2.15)	-0.111	(1.04)
Day 10							1.851	(2.51)	-0.105	(1.23)
Panel 2: All Factory & Same Floor										
Cumul. Previous Output x ...										
Day 1	0.272**	(0.11)	0.112	(0.13)	0.716	(1.10)			0.438	(1.07)
Day 2	0.271***	(0.09)	0.155	(0.10)	1.168	(0.84)			1.015	(0.83)
Day 3	0.248***	(0.08)	0.212**	(0.09)	1.444**	(0.70)			1.317*	(0.69)
Day 4	0.375***	(0.08)	0.335***	(0.09)	2.531***	(0.76)			2.410***	(0.77)
Day 5	0.263***	(0.06)	0.091	(0.10)	0.057	(0.86)			0.053	(0.86)
Day 6	0.206**	(0.08)	0.109	(0.09)	0.874	(0.85)			0.873	(0.87)
Day 7	0.181*	(0.09	0.148	(0.11)	1.182	(0.84)			1.220	(0.84)
Cumul. Previous Output Same Floor x ...										
Day 1	0.543***	(0.13)	0.339***	(0.12)	2.922***	(1.05)			3.319***	(1.01)
Day 2	0.327***	(0.10)	0.162	(0.10)	1.565*	(0.83)			2.024**	(0.82)
Day 3	0.238***	(0.09)	0.130	(0.09)	1.498**	(0.76)			1.765**	(0.76)
Day 4	0.101	(0.08)	0.009	(0.09)	0.174	(0.78)			0.237	(0.78)
Day 5	0.199***	(0.08)	0.154*	(0.09)	1.424*	(0.79)			1.406*	(0.80)
Day 6	0.208**	(0.10)	0.154	(0.11)	0.778	(0.89)			0.834	(0.89)
Day 7	0.170*	(0.10)	0.060	(0.10)	0.391	(0.84)			0.345	(0.86)
Output Lines Starting Same Day x ...										
Day 1							4.979*	(3.00)	1.683	(1.62)
Day 2							3.256	(2.63)	0.858	(1.26)
Day 3							2.739	(2.38)	1.574	(1.17)
Day 4							5.285*	(2.68)	1.738	(1.23)
Day 5							2.152	(2.17)	0.609	(1.35)
Day 6							-0.376	(2.65)	-1.320	(1.66)
Day 7							2.811	(3.00)	-2.389	(1.67)
Output Lines Starting Same Day, Same Floor x ...										
Day 1							7.967***	(3.01)	4.892***	(1.54)
Day 2							7.811***	(2.71)	4.278***	(1.48)
Day 3							4.481*	(2.41)	1.274	(1.11)
Day 4							0.524	(2.78)	-1.195	(1.02)
Day 5							0.336	(2.20)	-0.424	(1.37)
Day 6							2.109	(2.30)	2.445	(1.84)
Day 7							-1.330	(2.90)	2.771*	(1.65)
N	33,153		33,153		33,162		13,077		32,890	
Controls	YES		YES		YES		YES		YES	
Line Chief FE	YES		YES		YES		YES		YES	
Month FE	YES		YES		YES		YES		YES	
Style FE			YES		YES				YES	
Type FE							YES			

Notes: Column 1 reports effects of regressing daily line productivity on the first 10 days of producing a new style on log cumulative previous output of the same style on other lines, interacted with dummies for each of the first 10 production days of a style on a line (Panel 1). Panel 2 also includes log previous output on other lines on the same floor only, interacted in same way. Column 2 adds style fixed effects interacted with dummies for the first ten production days. Column 3 transforms these variables into dummies indicating positive amount of previous output. Col. 4 instead regresses line productivity on a dummy for other lines starting the same style on the same day among first lines in the factory to produce that style. Col. 5 regresses productivity on dummies for positive previous cumul. output and for lines starting same style on the same day, as usual interacted with dummies for the first 10 production days. Controls are SMV, daily runtime, and number of workers on line. All Controls and FE in turn interacted with n'th day of production of style on the line. Panel 2 shows only coefficients for 1st-7th day of production due to space constraints. Due to computational constraints, regressions were estimated separately for each of the first ten days a line produces a new style. 'N' refers to summed 'N' for all ten of these regressions. Standard errors clustered at the line chief level in brackets; *: $p < 0.1$, **: $p < 0.05$, ***: $p < 0.01$

style. In these cases, no other workers have prior experience with the style that could be shared, and learning effects should be absent. Thus, in column 4 of Table 2, 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.[13] Indeed, strong and significant effects can be seen in column 4, Panel 1, again over the first five days a line produces the new garment. Again, this effect is mainly driven by instances where the other lines starting the new style on the same day are located on the same floor (Panel 2).

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 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” between the two specifications in column 5, which replicates column 3, but adds dummy variables for at least one other line (on the same floor) starting to produce the same garment on the same day. Note that this dummy variable is not mechanically correlated 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 only 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 again include style fixed effects. The results from column 5 show that in this specification, both effects that we saw in columns 1 to 4 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. The next section will exploit randomized variation in (the potential for) knowledge exchange across floors to find more conclusive evidence of whether knowledge exchange is indeed a major driver of these productivity spill-overs.

[13] I regress productivity on the first ten style-days on a dummy indicating that some other line (on the same floor) starts producing the same style on the same day as well, interacted as usual with dummy variables for each of the first ten style-days included in the sample. There are 2,398 instances in the data of a line starting a style that has not been produced before on another line. For 1,329 of them, some other line in the factory also starts producing the same style on the same day, and for 1,198 instances, at least one of these other lines is located on the same floor.

Knowledge Depreciation

To conclude this section, I investigate whether there is evidence for "organizational forgetting", or knowledge depreciation in this setting. There is an established subsection in the literature on organizational learning that investigates whether knowledge about production processes depreciates over time, which has found conflicting evidence so far. While Benkard (2000) and David and Brachet (2011) found strong evidence for depreciation, Thompson (2007) and Egelman et al. (2016) find little. I follow the approach from, among others, Argote et al. (1990) and Thompson (2007), with a grid search of daily knowledge depreciation parameters over one percent intervals, which I use to construct the "depreciated" past cumulative output of the same style (from the same unit). With this augmented variable, I re-run columns 1 and 2 from Table 2 using maximum likelihood estimation, and look for the depreciation parameter that yields the highest likelihood function value in these estimations. When not using style fixed effects, as in column 1 of Table 2, then the highest likelihood value is obtained with zero depreciation. When using style fixed effects as in column 2, then a daily depreciation factor of 5% yields the highest likelihood value when restricting the sample to the first style-day only, while a depreciation factor of 2% maximizes the value when using the sample of the first three or ten days. On the other hand, the variation in likelihood values is not strong, and the estimation results of the spill-over remain more or less unchanged when using different depreciation factors between 0 and 20% per day (see Appendix B.2 for more results). Thus, there is some evidence that daily knowledge depreciation rates could be around 2-5%, though including depreciation or not in the estimation does not affect the qualitative results on productivity spill-overs in this study.

4 Randomized Help Provision

To better identify the effect that knowledge exchange between co-workers has on productivity growth, I exploit a randomized management intervention that was carried out at the three sample factories under the supervision of my research team. Whenever a line on randomly selected "treatment" floors began producing a new style that had already been produced by any other line in the factory, the most senior line chief with previous experience of the style was instructed by the factory's production management to brief the line chief who is now also starting the same style. The 15- to 30-minute brief was intended to explain how the earlier line had overcome initial problems, particularly at "bottlenecks" operations that slowed down the new style's production.[14]

[14] 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 on the new style. The briefings from the communication intervention were to have been done once all machines on the line had been adjusted for the new style. Its goal was to convey the knowledge 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 do not have, as they design the sewing line layout for a style from its

The intention of this treatment was to exogenously increase the potential for knowledge exchange in the production process of the style between randomly selected pairs of line chiefs.[15] The experiment ran on the treatment floors for four months, from June to end of 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 production 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 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 will focus on the "intention to treat" effect, assuming that any start of a style that should have been treated was actually treated. Appendix C shows results when regressing productivity on recorded instances of the treatment.[18]

Randomization and Balance

Nine floors from a sample of 17 floors across the three factories were randomly selected to receive the treatment (randomization stratified at the factory level - Factory 3 requested to include only

design template, and supervise the switching process of lines to new styles, but do not stay on a line after the switch.

[15]We can think of the intervention lowering the costs of seeking and providing help between line chiefs, particularly two parts of the costs. First, the perceived cost of approaching someone else for help, as one exposes a lack of knowledge (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 reveals a lack of knowledge. Second, 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 fixed costs are borne, one would need to decide how much effort to spend analysing the problem and communicating a solution; this introduces a variable cost. While the randomized help provision eliminates neither the fixed nor the variable costs, it renders the fixed costs 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 of the line chief when deciding whether to provide help.

[16]26 further briefings of the 154 recorded 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. Possibly, line chiefs who had already produced similar styles were sent to give instructions.

[17]All three factories reported that prior to the intervention, they occasionally sent line chiefs with experience on a style to other lines to help co-workers if they started producing the same style. 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.

[18]Note 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.

Table 3: Balancing of Randomization across Sewing Lines

Variable	Control	Diff.	N
Line Characteristics:			
Nbr. worker	30.20	-1.39	137
Daily Runtime	9.156	0.43*	137
Efficiency	50.83	-1.42	137
Efficiency First Day	36.77	-6.24**	121
SMV	10.81	-0.99	137
Start Rank	3.859	0.45	137
Supervisor Characteristics:			
Age	29.53	0.45	79
Seniority Factory	65.67	-3.21	72
Seniority as Supervisor	35.92	-1.94	79
Sen. as SV on curr. line	26.02	-2.60	69
External Arrival as SV	0.333	0.05	72
Nbr. Social Ties	2.806	0.47	79
Education	15.33	-0.41	72

Notes: Line Chief characteristics from line chief surveys. Line characteristics from production data. "Control" column shows average values from control floors from April and May 2014. "Diff. " column shows deviation of average values on treatment floors from those from control floors. Statistical differences in comparisons: *** p<0.01, ** p<0.05, * p<0.1.

six of its 14 floors in the sample).[19] Table 3 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, controlling for factory fixed effects. No observable line chief characteristics differ significantly on conventional levels. However, lines on treatment floors have significantly longer daily runtime (operation hours); I will control for daily runtime in later regressions. While overall line productivity does not differ significantly between treatment and control lines, productivity on treatment floors is significantly lower on the first day lines on these floors produce new styles, which is of importance as explained further below.

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. This allows me to control the regression analysis of the management intervention for pre-treatment levels of line productivity and expand the sample size of the estimation to obtain more precise estimates. Figure 4 plots the average productivity over the first five days a line produces a new style for four different cases: treated lines before and during the intervention months, and control lines before and during the intervention months. I use data from the beginning of 2014 until the

[19] In 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.

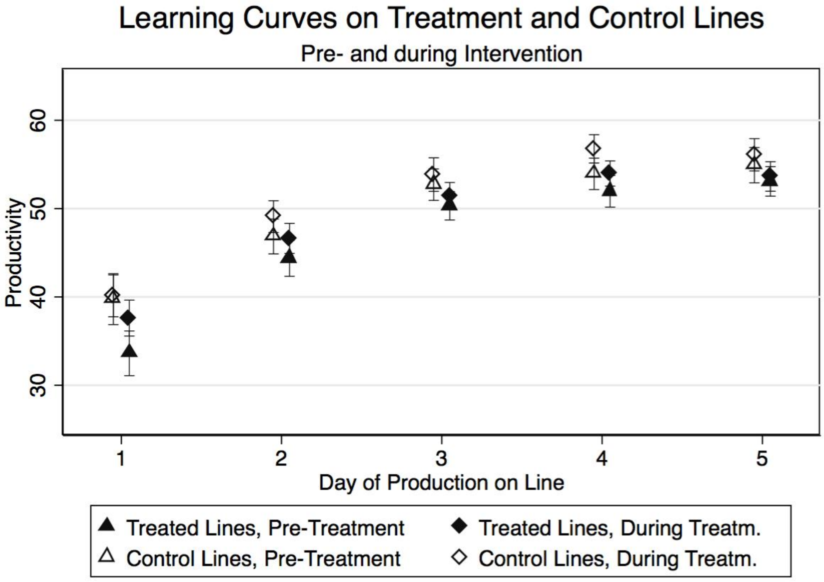

Figure 4: Pre-and Post- Treatment Learning Curves of Treated and Control Lines. The figure plots average productivity over the first five days a line produces a new style previously produced on other lines for four different cases: from treatment floors prior to start of treatment, treatment floors during experiment, control floors prior to start of treatment, and control floors during treatment. Productivity standardized at factory-level.

end of September 2014, when the initially agreed upon 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 3. This difference is not accounted for by observable characteristics of lines or line chiefs, and is driven by two 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, particularly on the first day a new style was produced.

Using a difference-in-difference design to estimate the effects of the intervention, the identifying assumption is that absent the intervention, productivity levels on the treatment floors would have followed the same trends as on the control floors. One way to assess the validity of this assumption is to determine whether productivity on treatment and control floors follow similar trends before the start of the intervention. Figure 5 plots, for the first day a line produces a new style that has already been produced by another line chief, monthly average productivity from January to September 2014, separately for lines selected for treatment (square symbols) and as control lines (triangle symbols). Given that data from Factory 3 is only available from

April 2014 onwards, the graphs are shown separately for Factory 3 on the right column of Figure 5. The solid vertical line in the graphs indicates the start of the treatment with June 2014, while the dashed line indicates the end of the agreed upon intervention period at the end of September 2014. However, log-book entries indicate that the briefings continued to some extent, but anecdotal evidence suggest that from then on they were also extended to control floors.

The upper two graphs show the productivity trends when only controlling for SMV, daily runtime and number of workers on the line, while the lower two graphs also control for style fixed effects. As suggested already in the balancing Table 3, first-day productivity was systematically lower on floors selected for treatment in the months before the start of the treatment. And while the data are noisy, leading to fluctuations in first-day productivity across months, productivity broadly follows the same trends at all factories in the month running up to the start of the intervention. The productivity difference is then offset after the start of the intervention due to an upward shift of first day productivity on treatment floors. Interestingly, during the first month of the intervention, if anything, a negative treatment effect appears; compared to control floors, productivity on treatment floors drops. However, in the remaining three months, it catches up and in some months at some factory overtakes productivity at the control floors. The fact 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.[20] Appendix C.2 shows parallel trends from Factory 1 and 2 for a longer time preceding the intervention, as from these two factories, more than one year of production line data before the start of the intervention is available.

Regression Specification

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 five style-days from each style start. I restrict the sample, as in Figures 4 - 5, to style-starts from observations from January to September 2014 and to style starts for which another line in the factory has already produced the same style, so that in principle the style start is "treatable" due to the availability of another line chief who has already produced the style. Using this sample, I run the following baseline regression, resembling the regressions from Table 2:

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

"$Treat_{fisn}$ is a dummy indicating that line i is a treatment line and that it started produc-

[20]The 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, in which the share was roughly twice as large

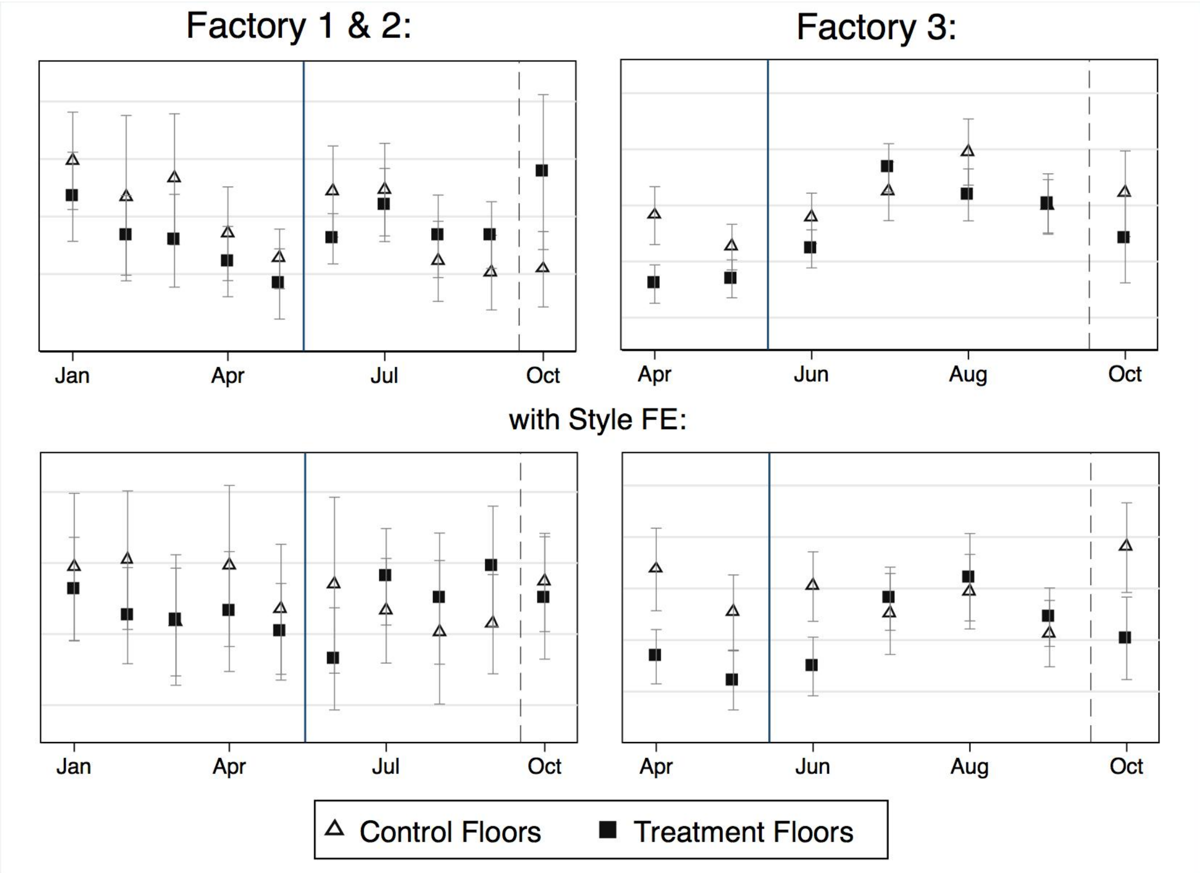

Figure 5: Pre-Post Intervention Start Trends for First-Day Productivity. The graph shows average monthly productivity of lines on the first day they start producing a new style that has already been produced on another line in the factory, separately for lines selected for treatment (solid squares) and lines not selected (hollow triangles). The solid vertical line indicates start of treatment from June 2014 on. The dashed line indicates the end of the initially agreed upon treatment time at the end of September 2014. Capped bars represent 95% confidence intervals.

ing style s for the first time during the treatment time June-September 2014, interacted with 'style-day' fixed effects. As in the previous section on general spill-overs, I control for fixed effects a_{fin} at the line-chief - 'style-day' level, and γ_{ft} on the factory - month - 'style-day' level. Furthermore, $Prev_{isn}$ stands for controls of the cumulative previous amount of the garment that has been produced on other lines before line i started producing it. Because the precise functional form with which to control for this previous output is not obvious ex ante, I experiment with two different approaches. In Panel 1 in Table 4 I use a set of four control variables, each interacted with style-day and factory fixed effects: Previous output, previous output on the same floor, log previous output, and log previous output on the same floor. In Panel 2, instead I use fixed effects on the number of line chiefs that have already produced the same garment, again interacted with factory and style-day fixed effects. All regressions control for SMV, daily runtime and number of workers on the line, interacted with factory and 'style-day' fixed effects.

Due to the small number (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 et al. (2008). I follow their suggestion to use wild cluster bootstrap to obtain standard errors at all regressions estimating the effects of the randomized intervention.

Results of Intervention

Table 4 shows the main results from the communication intervention. Column 1 shows the results from the initial regression described above. Regardless of whether using (log) previous output (on the same floor) (Panel 1), or fixed effects for the number of line chiefs who already produced the garment (Panel 2), a significant positive effect on productivity on the first day of production of a new style can be seen. The two estimates suggest that the effect on the first day is roughly five productivity units, or about a third of the average productivity drop on the first day lines produce new styles that they did not produce before. On the second day a line produces the style, the effect is somewhat smaller but still visible and marginally significant in Panel 2, but not anymore in Panel 1. From the third day onwards, no effects are discernible anymore in either panel.

In column 2, I again add style fixed effects (interacted with dummies for the first five style-days in the sample). This is a demanding specification, as it only uses variation in productivity within the same style and within the same style-day either across treatment and control floors or pre-treatment and treatment times. Given that on average a garment is being produced on four lines in the data, and the first line to produce the garment is excluded from the sample in this specification, each cell has on average only three observations. Nevertheless, the effect remains significant on the first style-day and marginally significant on the second in Panel 2,

Table 4: General Treatment Effects

					Reweighted	
Panel 1: (Log) Previous Output						
Treatment x						
Day 1	4.842**	(.045)	5.505**	(.040)	5.297**	(.0.40)
Day 2	2.120	(.290)	2.549	(.145)	3.098**	(.020)
Day 3	2.235	(.370)	2.201	(.495)	1.504	(.575)
Day 4	2.815	(.275)	5.016**	(.45)	4.438**	(.045)
Day 5	2.151	(.550)	3.575	(.255)	4.375	(.340)
N	4,903		4,903		4,652	
Panel 2: Nbr. Prev. LCs FE						
Treatment x						
Day 1	5.537***	(0.010)	5.386*	(.070)	6.061**	(.020)
Day 2	3.816*	(0.065)	3.204*	(.050)	3.800***	(.010)
Day 3	2.976	(0.330)	3.352	(.325)	2.522	(.470)
Day 4	1.955	(0.465)	4.061	(.215)	2.124	(.595)
Day 5	1.782	(0.670)	4.517	(.285)	5.719	(.255)
N	4,946		4,946		4,682	
Controls	YES		YES		YES	
Line Chief & Month FE	YES		YES		YES	
Style FE	NO		YES		YES	

Notes: Results from regressing an ITT treatment dummy of the line chief receiving a briefing on a new style on the productivity on each of the first five days his/her line produces the new style. Sample only includes line chiefs starting styles that had already been produced by other lines before. Column 2 and 3 include style fixed effects interacted with dummies for the first five "style-days" a line produces the style. Column 3 uses reweighting technique based on DiNardo et al. (1996) to control for differential pre-treatment productivity of treatment and control lines. Controls include SMV, daily runtime, and number of workers on line. All controls, as well as line chief and month fixed effects interacted with factory and style-day fixed effects. In column 2 and 3, due to computational constraints, separate regressions for each of the first five style-days run; N shown is sum across the individual N for these five regressions. Wildcluster bootstrap based p-values clustered at 17 randomization units shown in brackets: *** p<0.01, ** p<0.05, * p<0.1.

in line with the effects from the more basic specification in column 1 (a significant effects can also be seen on the fourth style-day in Panel 1, but it could be discounted as spurious, given the demanding specification).

As pointed out before, the fact that the effects of the intervention, similar to the general spill-over effects in the previous section, are only visible during the first days a line produces a new style supports the view that the intervention effects are indeed due to knowledge exchange, and not due to alternative mechanisms, such as the briefing simply informing the line chief that other lines also produce the style, thus triggering competitive peer effects. Such peer effects could be expected to continue for longer, particularly if the management discounts the first production days when comparing productivities of different lines producing the same style, given that lines need several days to achieve full productivity with a new style. Instead, the pattern suggests that treated lines receive an early boost in production related knowledge, reducing the initial productivity drop when lines switch to new styles with which they are not yet familiar.

This conclusion is also supported by the fact that most documented briefings (ca. 75 percent) happened between supervisors from the same floor. Floors typically contain only 5-10 production lines, and a supervisor should be well aware of which lines produce which garment, given frequent group meetings with their superiors and given that the style currently produced on a line is usually displayed at the head of a line. This speaks again against the hypothesis that the intervention merely informed the supervisors that other lines are also producing the same style.

The reduction in the difference of first-day productivity with the onset of the intervention between treatment and control floors, as shown in Figure 5, 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, in column 3 of Table 4, 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 Duflo 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 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. The results do not differ qualitatively from the ones in the previous columns. Appendix C.3 provides more technical details on this reweighting approach.

Figure 6 shows distributions of line productivities on the first day production lines 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 intervals. The increase in first-day productivity of treatment lines during the implementation of the intervention seems to be driven by a strong reduction in 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 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. 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.

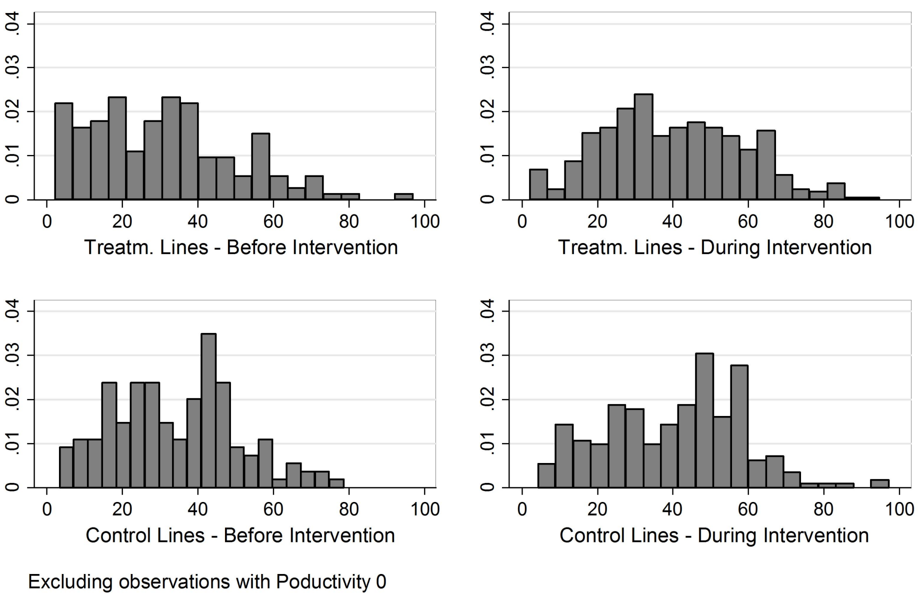

Figure 6: First Day Productivity Distribution before and during Intervention. The graph shows distribution of productivity on first day sewing lines produce new styles previously produced on some other line in the factory, on treatment floors, before and during implementation of intervention (top row), and on control floors, before and during the implementation (bottom row).

Return on Intervention

Thus, increasing the scope for knowledge exchange across production units within firms leads to increased productivity spill-over. To obtain an idea of the overall effect of this intervention on factory productivity and profits, I make the conservative assumption that the intervention only affected line productivity on the first day the line produced a new style, increasing productivity on average by 5 units on that day, as suggested by the results from Table 4. Sewing lines on average switch to a new style every 12 days, and at roughly every second of these starts, another line has previously produced the style. 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 $5/(12 * 2 * 47.4) = 0.44$ 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 0.44 percent of labor costs, this would translate into an increase in profits of almost 0.9 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 starting 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 0.9 percent increase would thus imply an increase in profits of more than $5,000, or a return on the intervention of more than 500 percent, given the interventions monetary costs. These high estimated returns are also in line with anecdotal reports from the production managers in the post-intervention survey, indicating that the monetary costs of the intervention were not deemed as a hindrance to its implementation.

5 Why Intervention was not done before

Given the positive estimated treatment effects and high returns of the intervention, the question remains as to why such a management routine was not 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). See also Hanna et al. (2014) on evidence of the failure of micro-entrepreneurs to notice possible improvements in operations by focusing attention on too-limited a set of dimensions for possible improvements); and resistance by involved workers, such as the line chiefs. A post-experiment survey was conducted with the three production managers of the factories, 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).

Resistance by line chiefs, on the other hand, could stem from at least two sources; 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)). During the intervention, the production managers charged with documenting the actual briefings should have recorded for each briefing whether they had encountered any resistance by line chiefs, or whether for any other reason they deemed the briefing to have been ineffective. However, this data did not yield much variation that could be exploited. On the other hand, 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, particularly by younger, less-experienced line chiefs. Exploiting the imperfect compliance in the implementation of the intervention, at least according to the recording of the briefings that took place, I do find some evidence that older line chiefs randomly selected to receive a briefing were less likely to actually receive one. Table 5 shows the marginal effects from probit regressions of a dummy indicating that a briefing was recorded, on the characteristics of the line chiefs who should have received the briefing, among the sample of the 377 briefings that should have happened. It shows that conditional on factory or floor fixed effects, and a set of controls, older line chiefs were less likely to receive briefings when randomly selected to receive one, an effect that is marginally statistically significant (standard errors clustered at line chief level). Interestingly, more educated line

Table 5: Who is Being Treated?

VARIABLES	(1) Treatm.	(2) Treatm.	(3) Treatm.
Age	-0.005*	-0.006*	-0.010*
	(0.003)	(0.004)	(0.006)
Seniority as Line Chief			0.002
			(0.002)
Female			-0.033
			(0.165)
Education			0.029*
			(0.016)
Nbr. Social Ties			0.014
			(0.013)
SMV			-0.002
			(0.003)
Productiv., Day			0.000
			(0.002)
Avg. Productiv., Year			-0.007
			(0.013)
Avg. Productiv., Year, First Style-Day			-0.007
			(0.009)
Observations	365	365	311
Factory FE	YES	NO	NO
Floor FE	NO	YES	YES

Notes: Among a sample of 377 briefings that should have occurred according to the intervention design, probit regressions were run of a dummy indicating that the briefing was recorded to actually have taken place on characteristics of line and line chief who should have received a briefing. 'SMV' is SMV of style on which briefing should have occurred. 'Productiv. Day' is the line productivity of the line chief who should have receive briefing, on the day the briefing should have occurred. 'Avg. Productiv., Year' is average line productivity of the line chief who should have received the briefing since the start of the year 2014, and 'Avg. Productiv., Year, First Style-Day' is average line productivity in 2014 on the first days the line chief produces new styles. Standard errors clustered at the line chief level: *** p<0.01, ** p<0.05, * p<0.1.

chiefs are marginally more likely to receive briefings conditional on being selected to receive one, indicating that increased education might make people more receptive to help from others.[21] Experience, in terms of years working already as a line chief has a similarly large effect as age, though being far from statistically significant (p-value 0.29). These results lend some credence to the hypothesis that status concerns among workers, possibly amplified by age, could interfere with the smooth implementation of a knowledge exchange routine as attempted in these three factories.

[21] Education is an ordinal measure of highest schooling level/degree obtained.

6 Conclusion

This paper presents novel evidence on organizational learning from three large and modern Bangladeshi garment factories. It shows strong evidence of productivity spill-over across production lines within these factories, which is not sensitive to different knowledge depreciation factors considered when proxying accumulated experience with a garment with past cumulative output. I find that even within factories, spatial closeness of production lines matters for the strength of the spill-overs, while garment complexity does not. Finally, exploiting a randomized knowledge exchange management intervention that ran for four months in the three factories, I show that increasing knowledge exchange leads to increased productivity spill-over. This supports the long held assumption in organizational economics that productivity spill-overs are a useful proxy for the otherwise hard to observe phenomenon of organizational learning.

Given the positive 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 had 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. 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 make time 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 some evidence that the factories were more likely to comply with the instructions to implement the briefings, if the line chief receiving the briefing was younger. 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 (2017). Input Allocation, Workforce Management and Productivity Spillovers: Evidence from Personnel Data. mimeo, McGill University.

Argote, L. (2013). *Organizational Learning.* Springer US.

Argote, L., S. l. Beckman, and D. Epple (1990). The persistence and transfer of learning in industrial settings. *Management Science 36*(2), 140–154.

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 (2017). Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan. *Quarterly Journal of Economics 132*(3), 1101–1164.

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 (2013). Team Incentives: Evidence From A Firm Level Experiment. *Journal of the European Economic Association 11*(5), 1079–1114.

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. C., J. Gelbach, and D. Miller (2008). Bootstrap based Improvements for Inference with Clustered Errors. *Review of Economics and Statistic 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). Experimential and Social Learning in Firms: The Case of Hydraulic fracturing in the Bakken Shale. mimeo, University of Chicago.

Darr, E. d., L. Argote, and D. Epple (1995). The Acquisition, Transfer, and Depreciation of Knowledge in Service Organizations: Productivity in Franchises. *Management Science 41*(11), 1750–1762.

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.

David, G. and T. Brachet (2011). On the Determinants of Organizational Forgetting. *American Economic Journal: Microeconomics 3*(3), 100–123.

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.

Duflo, 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.

Egelman, C. D., D. Epple, L. Argote, and E. R. Fuchs (2016). Learning by Doing in Multiproduct Manufacturing: Variety, Customization, and Overlapping Product Generations. *Management Science 63*(2), 405–423.

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 98*(1), 394–425.

Haltiwanger, J., R. S. Jarmin, R. Kulick, and J. Miranda (2016). High Growth Young Firms: Contribution to Job, Output and Productivity Growth. in: Measuring Entrepreneurial Businesses: Current Knowledge and Challenges. National Bureau of Economic Research, Inc.

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 Evidence from a Field Experiment. *Quarterly Journal of Economics 129*(3), 1311–1353.

Hidalgo, C. A., B. Klinger, A. Barabasi, and R. Hausmann (2007). The product space conditions the development of nations. *Science* (317), 482–487.

Hjort, J. (2014). Ethnic Divisions and Production in Firms. *Quarterly Journal of Economics 129*(4), 1899–1946.

Ichniowski, C. and K. Shaw (1997). The Effects of Human Resource Management Practices on Productivity: A Study of Steel Finishing Lines. *American Economic Review 87*(3), 291–313.

Kaldor, N. (1967). *Strategic Factors in Economic Development.* Ithaca, NY: New York State School of Industrial and Labor Relations, Cornell University.

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.

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.

Macchiavello, R., A. Menzel, and C. M. Woodruff (2015). Challenges of Change: An Experiment Training Women to Manage in the Bangladeshi Garment Sector. mimeo, CERGE-EI Prague.

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. (2007). How Much Did the Liberty Shipbuilders Forget? *Management Science 53*(6), 908–918.

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 workers 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 three to five 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 previously produced the same style. 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

Table 6: Worker Movement

VARIABLES	Efficiency
Additional Line	3.570***
	(1.029)
Constant	57.052***
	(3.379)
Observations	19,972
R^2	0.429
Controls	YES
Factory-week FE	YES
L.Chief-Style.Day FE	YES

Notes: The table shows the result of regressing, in the sample of the first lines in the factory to produce a style, daily line productivity on a dummy indicating that by that day, further lines also started producing the same style ("Additional Line"). Regressions control for factory-week fixed effects and line chief fixed effects interacted with style-day, and SMV, runtime and worker on line interacted with factory - 'style-day' fixed effects. Standard Error clustered at the line chief level in parentheses: *** p<0.01

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 of 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 also start producing the style, and workers are reallocated away to these lines. Table 6 below shows the results when regressing the daily productivity of the 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 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; they 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 of certain styles caused the observed productivity spill-over.

Appendix B1: Productivity Spill-Over and Product Complexity

Table 7 shows the results when interacting the variables for the previous output of the same style with the style complexity measure SMV. Columns 1 and 2 of Table 7 repeat columns 1 and 2 of table 1, but adding Log Previous Output on other lines (on the same floor) interacted with SMV, as usual interacted with dummy variables for each of the first ten style-days included in the sample. Note that SMV itself was already included in all regressions of Table 2, so no longer need be added. Column 3 interacts log previous output of the same garment (on the same floor) with a dummy indicating that the style's SMV is of above median SMV in the factory for robustness. As usual, Panel 1 shows the results only for the overall previous output (interacted with SMV), while Panel 2 additionally contains the previous output on the same floor (interacted with SMV). The results do not show any consistent effect of product complexity on spill-over strength. When not controlling for style fixed effects (Column 1), then the spill-over strength seems to be reduced on the first day lines produce a garment for more complex styles, but this effect becomes small and insignificant when controlling for style fixed effects (columns 2 and 3). On the other hand, on the third style-day, past output seems to have a stronger effect among more complex products, but given the large amount of coefficients estimated, this is probably a spurious effect. The overall effect of previous output is qualitatively changed little by the inclusion of the interaction terms. The effect remains highly visible during the first three to five days a line produces a garment, an effect that seems mainly driven by previous output on the same floor.

Appendix B2: Results with Knowledge Depreciation

A number of studies have shown that knowledge gained within a firm can depreciate (Argote et al. (1990); Benkard (2000); Thompson (2007); David and Brachet (2011)). To study whether knowledge depreciation can also be measured in the garment factories in this sample, I follow an approach laid out by, among others, Argote et al. (1990). I construct the central independent variables of interest, log past cumulative output of the same garment (on the same floor) using different daily depreciation factors, and re-run the central regressions of interest, Columns 1 and 2 of Table 2 using maximum likelihood estimation. In the subsequent estimation I look for the depreciation factor that yields the highest likelihood function value. I first broadly used daily depreciation rates of 0%, 10%, 20%, 30%, to establish that with these values, the likelihood value decreases as one uses higher depreciation factors, and then I used one percent steps between zero and ten percent to find more precise estimates. As in Table 2, for computational reasons, I estimate each regression individually for each style-day. Therefore, I obtain likelihood values for each style-day. When using style fixed effects (Panel 2 of Table 8 below), looking only at the data from the first style-day, then a depreciation factor of 5% yields the highest likelihood value. When summing up the likelihood values over the first three, five, or ten days of production, then a depreciation factor of 2% yields the highest likelihood value. On the other

Table 7: General Productivity Spill-Over & Product Complexity

	(1) SMV		(2) SMV		(3) SMV > med. SMV	
Panel 1: All Factory						
Cumul. Previous Output ...						
Day 1	0.986***	(0.15)	0.608***	(0.21)	0.267	(0.16)
Day 2	0.463***	(0.12)	0.283*	(0.16)	0.170*	(0.10)
Day 3	0.279***	(0.10)	0.168	(0.15)	0.220**	(0.09)
Day 4	0.328***	(0.10)	0.349**	(0.14)	0.261**	(0.11)
Day 5	0.366***	(0.11)	0.290**	(0.14)	0.249**	(0.11)
Day 6	0.277**	(0.11)	0.350**	(0.15)	0.247**	(0.11)
Day 7	0.301***	(0.11)	0.153	(0.17)	0.100	(0.14)
Day 8	0.287**	(0.12)	0.273	(0.19)	0.358**	(0.17)
Day 9	0.239**	(0.12)	0.169	(0.17)	0.186	(0.13)
Day 10	0.375**	(0.15)	0.167	(0.22)	0.248	(0.15)
Cumul. Previous Output x SMV ...						
Day 1	-0.034***	(0.01)	-0.025	(0.02)	0.130	(0.20)
Day 2	0.001	(0.01)	-0.002	(0.01)	0.174	(0.12)
Day 3	0.010	(0.01)	0.011	(0.01)	0.140	(0.13)
Day 4	0.010	(0.01)	-0.001	(0.01)	0.150	(0.13)
Day 5	0.002	(0.01)	-0.009	(0.01)	-0.115	(0.15)
Day 6	0.005	(0.01)	-0.013	(0.01)	-0.080	(0.15)
Day 7	-0.002	(0.01)	0.003	(0.01)	0.160	(0.15)
Day 8	0.009	(0.01)	-0.001	(0.01)	-0.167	(0.19)
Day 9	-0.004	(0.01)	-0.007	(0.01)	-0.162	(0.16)
Day 10	-0.000	(0.01)	0.009	(0.02)	0.030	(0.21)
Panel 2: All Factory & Same Floor						
Cumul. Previous Output ...						
Day 1	0.569***	(0.21)	0.273	(0.24)	0.033	(0.18)
Day 2	0.138	(0.17)	-0.016	(0.18)	0.018	(0.13)
Day 3	-0.004	(0.14)	-0.104	(0.16)	0.053	(0.10)
Day 4	0.189	(0.15)	0.357**	(0.17)	0.255**	(0.13)
Day 5	0.106	(0.14)	0.075	(0.17)	0.072	(0.13)
Day 6	0.185	(0.17)	0.193	(0.22)	0.075	(0.16)
Day 7	0.179	(0.16)	0.093	(0.22)	0.033	(0.18)
Cumul. Previous Output, same Floor ...						
Day 1	0.579***	(0.21)	0.436**	(0.21)	0.351**	(0.15)
Day 2	0.469***	(0.17)	0.429**	(0.18)	0.222**	(0.11)
Day 3	0.427***	(0.15)	0.403**	(0.17)	0.251***	(0.10)
Day 4	0.216	(0.17)	-0.0170	(0.17)	0.009	(0.13)
Day 5	0.401***	(0.14)	0.310	(0.19)	0.263*	(0.14)
Day 6	0.083	(0.19)	0.207	(0.20)	0.244*	(0.14)
Day 7	0.157	(0.16)	0.076	(0.17)	0.097	(0.15)
Cumul. Previous Output x SMV ...						
Day 1	-0.025*	(0.01)	-0.014	(0.02)	0.155	(0.27)
Day 2	0.011	(0.01)	0.014	(0.01)	0.269	(0.17)
Day 3	0.021*	(0.01)	0.026**	(0.01)	0.317**	(0.16)
Day 4	0.015	(0.01)	-0.002	(0.01)	0.145	(0.18)
Day 5	0.013	(0.01)	0.002	(0.01)	0.061	(0.21)
Day 6	0.002	(0.01)	-0.007	(0.02)	0.075	(0.22)
Day 7	0.000	(0.01)	0.005	(0.01)	0.207	(0.22)
Cumul. Previous Output, same Floor x SMV ...						
Day 1	-0.007	(0.02)	-0.012	(0.02)	-0.024	(0.26)
Day 2	-0.012	(0.01)	-0.025	(0.02)	-0.130	(0.19)
Day 3	-0.015	(0.01)	-0.023	(0.02)	-0.264	(0.20)
Day 4	-0.009	(0.01)	0.002	(0.01)	0.010	(0.21)
Day 5	-0.018	(0.01)	-0.016	(0.02)	-0.254	(0.26)
Day 6	0.013	(0.01)	-0.007	(0.01)	-0.201	(0.20)
Day 7	0.001	(0.01)	-0.001	(0.01)	-0.060	(0.21)
N	33,153		33,153		33,153	
Controls	YES		YES		YES	
Month FE	YES		YES		YES	
Line Chief FE	YES		YES		YES	
Style FE	NO		YES		YES	

Notes: Column 1 reports regression of daily output on the 1st - 10th day a line produces a new style on previous cumulative output of the same style on other lines, and on other lines on the same floor, and both interacted with the style's SMV. Column 2 add style fixed effects interacted with 'style-day' fixed effects. Column 3 interacts with dummy for above median SMV instead of absolute SMV value. Controls are SMV, daily runtime, and number of workers on line. All Controls and FE in turn interacted with indicators for n'th day of production of style at the line. Panel 2 shows only coefficients for 1st-7th day of production due to space constraints. Due to computational constraints, regressions were estimated separately for each of the first ten days a line produces a new style. 'N' refers to summed 'N' for all ten of these regressions. Standard errors clustered on line chief level in brackets; *: $p < 0.1$, **: $p < 0.05$, ***: $p < 0.01$

hand, when not using style fixed effects, then the likelihood value is always maximised when not depreciating at all (Panel 1 of Table 8 below). However, using different depreciation values does not change the general results on productivity spill-overs in any significant way, as shown in Table 8, which displays results when using depreciation values of 0%, 3%, 10%, and 20%.

Given that we find some evidence that daily knowledge depreciation could be around 2-5%, I also re-run the central regressions evaluating the randomized management intervention from Table 4, but controlling for depreciated past cumulative output at 2 and 5 percent. The results are shown in Table 9 below, and do not differ qualitatively from the results in Table 4.

Appendix C1: Results using Recorded Treatments

The main analysis of the randomized communication intervention 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 five 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 had 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 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 before 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 Table 4 and in Figures 4- 8), the only type of non-compliers left in the sample are style starts which should have been treated but were not.

Table 10 shows the results when repeating Table 4 but with treatment now indicating recorded treatments instead of ITT treatment. Using recorded treatments, we only find statistically insignificant and very small effects. The point estimates are even negative (though statistically completely insignificant) when not including style fixed effects (column 1). When using style fixed effects and absolute and log previous output of the same garment on other lines (on the same floor) as controls, the effect becomes somewhat positive, but remains statistically insignificant.

The insignificant and small estimated effects of recorded treatments are puzzling in light of the significant intention-to-treat effects. The following points could provide explanations

Table 8: Productivity Spill-Over with Knowledge Depreciation

	(1) No Depr.		(2) Depr. 3%		(3) Depr. 10%		(4) Depr. 20%	
Panel 1: Without Style FE								
Cumul. Previous Output x ...								
Day 1	0.272**	(0.115)	0.266**	(0.122)	0.241*	(0.125)	0.247**	(0.125)
Day 2	0.271***	(0.090)	0.305***	(0.097)	0.311***	(0.102)	0.310***	(0.104)
Day 3	0.248***	(0.078)	0.277***	(0.082)	0.303***	(0.086)	0.300***	(0.091)
Day 4	0.375***	(0.079)	0.429***	(0.083)	0.437***	(0.086)	0.430***	(0.089)
Day 5	0.263***	(0.063)	0.289***	(0.067)	0.306***	(0.070)	0.300***	(0.075)
Day 6	0.206**	(0.080)	0.232***	(0.085)	0.238**	(0.092)	0.237**	(0.096)
Day 7	0.181*	(0.092)	0.210**	(0.099)	0.195*	(0.107)	0.183***	(0.112)
Cumul. Previous Output, Same Floor x ...								
Day 1	0.543***	(0.133)	0.589***	(0.141)	0.630***	(0.148)	0.59***	(0.153)
Day 2	0.327***	(0.102)	0.302***	(0.111)	0.306**	(0.118)	0.310**	(0.122)
Day 3	0.238***	(0.087)	0.243***	(0.092)	0.253**	(0.098)	0.263**	(0.106)
Day 4	0.101	(0.080)	0.081	(0.085)	0.090	(0.091)	0.092	(0.096)
Day 5	0.199***	(0.076)	0.205**	(0.081)	0.212**	(0.086)	0.227**	(0.093)
Day 6	0.208**	(0.098)	0.223**	(0.103)	0.251**	(0.111)	0.273**	(0.119)
Day 7	0.170*	(0.098)	0.169	(0.106)	0.182	(0.117)	0.209*	(0.124)
Likelihood Day 1	-21122.9		-21126.7		-21134.2		-21145.4	
Likelihood Day 1-3	-55152.6		-55162.6		-55176.3		-55195.7	
Likelihood Day 1-5	-82780.4		-82789.3		-82807.6		-82835.4	
Likelihood Day 1-10	-131311.5		-131316.4		-131340.2		-131375.7	
Panel 2: With Style FE								
Cumul. Previous Output x ...								
Day 1	0.112	(0.135)	0.087	(0.139)	0.059	(0.147)	0.053	(0.158)
Day 2	0.155	(0.101)	0.171	(0.106)	0.160	(0.109)	0.166	(0.109)
Day 3	0.212**	(0.085)	0.210**	(0.087)	0.205**	(0.092)	0.193*	(0.101)
Day 4	0.335***	(0.093)	0.340***	(0.097)	0.319***	(0.105)	0.317***	(0.115)
Day 5	0.091	(0.099)	0.080	(0.100)	0.070	(0.106)	0.057	(0.112)
Day 6	0.109	(0.091)	0.105	(0.098)	0.074	(0.109)	0.066	(0.118)
Day 7	0.148	(0.112)	0.135	(0.112)	0.090	(0.119)	0.091	(0.122)
Cumul. Previous Output, Same Floor x ...								
Day 1	0.339***	(0.121)	0.394***	(0.129)	0.422***	(0.140)	0.407***	(0.150)
Day 2	0.162	(0.098)	0.151	(0.105)	0.155	(0.115)	0.157	(0.111)
Day 3	0.130	(0.088)	0.149	(0.095)	0.153	(0.103)	0.158	(0.157)
Day 4	0.009	(0.087)	0.009	(0.094)	0.013	(0.101)	-0.004	(0.109)
Day 5	0.154*	(0.091)	0.192**	(0.093)	0.19*	(0.097)	0.189*	(0.104)
Day 6	0.154	(0.106)	0.166	(0.114)	0.170	(0.124)	0.163	(0.135)
Day 7	0.060	(0.104)	0.073	(0.112)	0.061	(0.122)	0.042	(0.127)
Likelihood Day 1	-18898.2		-18895.2		-18896.9		-18901.9	
Likelihood Day 1-3	-48336.2		-48332.1		-48337.1		-48345.5	
Likelihood Day 1-5	-72371.8		-72365.4		-72376.2		-72389.9	
Likelihood Day 1-10	-113652.3		-113650.3		-113674.7		-113696.1	
N	33153		33153		33153		33153	
Month FE	YES		YES		YES		YES	
Line Chief FE	YES		YES		YES		YES	
Controls	YES		YES		YES		YES	

Notes: The table repeats column 1 (in Panel 1) and 2 (in Panel 2) from Table 2, Panel 2. Column 2 uses log cumulative prior output (from the same floor) using a 3% daily depreciation factor. Column 3 uses a 10% and column 4 a 20% depreciation factor. Controls are SMV, daily runtime, and number of workers on a line. All Controls and FE in turn interacted with indicator for n'th day of production of style at the line. Panel 2 shows only coefficients for 1st-7th day of production due to space constraints. Due to computational constraints, regressions were estimated separately for each of the first ten days a line produces a new style. 'N' refers to summed 'N' for all ten of these regressions. 'Likelihood Day X' refers to likelihood function value, summed up over the estimations for the first X days line produced garment. Standard errors clustered at the line chief level in brackets; *: $p < 0.1$, **: $p < 0.05$, ***: $p < 0.01$

Table 9: Treatment Effects Controlling for Depreciated Prior Output

					Reweighted	
Panel 1: (Log) Previous Output, 2% daily depr.						
Treatment x						
Day 1	4.104*	(.060)	5.117**	(.045)	5.000**	(.045)
Day 2	1.765	(.335)	2.368	(.225)	2.764	(.045)
Day 3	1.717	(.530)	1.858	(.575)	1.472	(.630)
Day 4	2.729	(.300)	4.922*	(.065)	4.470**	(.040)
Day 5	1.277	(.750)	3.322	(.335)	4.263	(.410)
Panel 2: (Log) Previous Output, 5% daily depr.						
Treatment x						
Day 1	4.216**	(0.050)	5.336**	(.045)	4.868*	(.055)
Day 2	1.600	(0.370)	2.353	(.225)	2.707*	(.060)
Day 3	1.181	(0.660)	1.733	(.590)	1.197	(.680)
Day 4	2.385	(0.385)	4.765*	(.070)	4.423*	(.055)
Day 5	0.624	(0.850)	2.821	(.475)	3.701	(.490)
N	4,903		4,903		4,652	
Controls	YES		YES		YES	
Line Chief FE	YES		YES		YES	
Month FE	YES		YES		YES	
Style FE	NO		YES		YES	

Notes: The two Panels of this Table repeat Panel 1 of Table 4, but using (log and absolute value of) depreciated cumulative previous output of the same garment on other lines (on the same floor) as controls. Panel 1 uses depreciated previous output using a daily depreciation factor of 2 percent, and Panel 2 depreciated previous output with a daily depreciation factor of 5 percent. Controls are SMV, daily runtime and number of worker on line. All controls and fixed effects interacted with factory and 'style-day' fixed effects. Wildcluster bootstrap based p-values clustered at 17 randomization units shown in brackets: *** p<0.01, ** p<0.05, * p<0.1.

Table 10: General Treatment Effects, Recorded Treatments

					Reweighted	
Panel 1: (Log) Previous Output						
Treatment x						
Day 1	-1.800	(.395)	1.698	(.490)	1.844	(.455)
Day 2	-0.960	(.680)	-0.157	(.890)	-0.764	(.690)
Day 3	3.786	(.190)	4.245**	(.015)	5.254**	(.050)
Day 4	1.508	(.465)	-0.507	(.750)	-0.527	(.770)
Day 5	1.615	(.790)	0.865	(.530)	0.333	(.825)
N	4,920		4,920		4,626	
Panel 2: Nbr. Prev. Line Chief FE						
Treatment x						
Day 1	-2.050	(.380)	0.502	(.895)	0.872	(.825)
Day 2	-0.940	(.585)	-1.722	(.555)	-2.178	(.360)
Day 3	2.442	(.310)	0.965	(.375)	0.532	(.600)
Day 4	0.792	(.630)	-0.969	(.535)	-1.429	(.485)
Day 5	-0.776	(.850)	-2.944**	(.035)	-2.739*	(.080)
N	4,920		4,920		4,626	
Controls	YES		YES		YES	
Line Chief FE	YES		YES		YES	
Month FE	YES		YES		YES	
Style FE	NO		YES		YES	

Notes: The table repeats Table 4, but instead of the ITT specification, it uses actual recorded treatments (except those, in column 3, that were recorded for lines starting a new garment that no other line had previously produced, and were thus not in line with the intervention protocol). Columns 2 and 3 use style fixed effects interacted with 'style-day', while column 3 uses the reweighting technique based on DiNardo et al. (1996) to control for differential pre-treatment productivity of treatment and control lines. Controls are SMV, runtime and number of workers on a line. All fixed effects interacted with factory and 'style-day' dummies. Wildcluster bootstrap based p-values clustered at 17 randomization units shown in brackets: *** p<0.01, ** p<0.05, * p<0.1.

for this pattern. First, as indicated by the factory managements, more style starts were likely treated than indicated as "treated" in the data. This would dilute the estimated effect of the treatment, particularly if the likelihood of the treatment of not being recorded was uncorrelated (or even negatively correlated) with the effect of the particular treatment on productivity. Second, actually administered treatments were probably directed towards style starts at which the factory management expected 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 effects becoming relatively larger when using style fixed effects as compared to when not, as well as with the reduction in the left tail of the productivity distribution on treated lines with the onset of the treatment as shown in Figure 6.

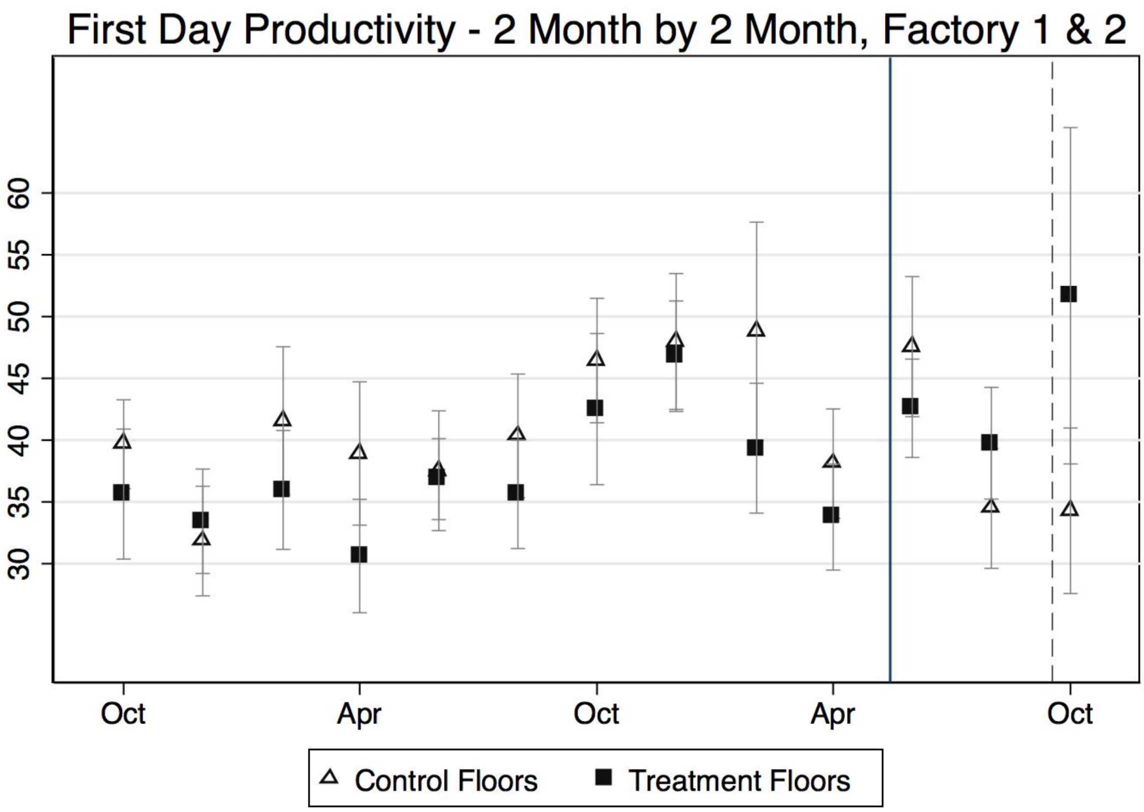

Figure 7: Pre-Post Intervention Start Trends for First Day Productivity, Factory 1+2.

Appendix C2: Long Term Productivity Trends at Factory 1 and 2

At Factory 1 and 2, we have more than one year of continuous production data before the start of the intervention. We can therefore investigate for a longer time period whether productivity on randomly selected treatment and control floors followed parallel trends, which would support the identifying assumption of the difference-in-difference estimation of the treatment effect of the knowledge exchange intervention. Figure 7 shows aggregate production data over two-month intervals from the first days lines produced styles that had already been produced by another line chief. I use two-month intervals instead of one-month intervals to reduce noise in the data and the amount of intervals shown, due to the longer time period included. Figure 7 shows the graph when only controlling for SMV, runtime and number of workers on line (interacted with factory FE). Figure 8 additionally controls for style fixed effects. The figures indicate that also over the longer time period considered here, first day line productivity largely follows parallel trends on treatment and control floors, and even more so when controlling for style fixed effects.

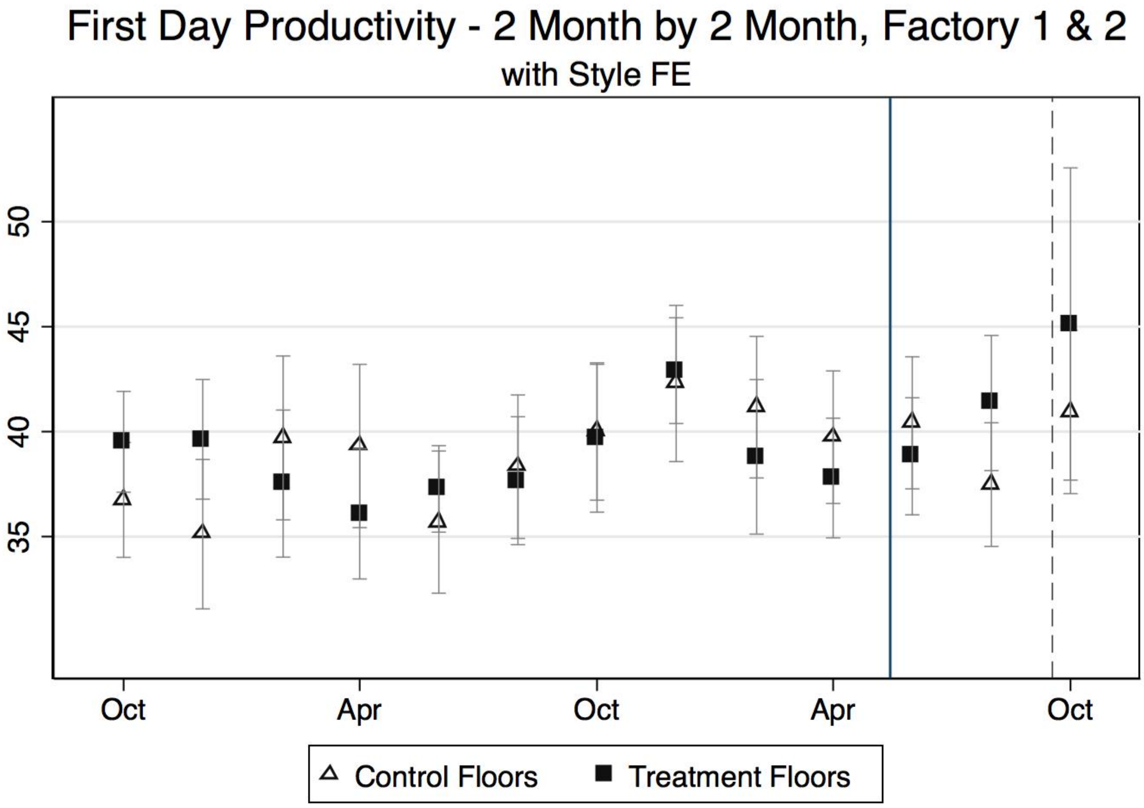

Figure 8: Pre-Post Intervention Start Trends for First Day Productivity, Factory 1+2. The graph shows average productivity within two-month intervals of lines on the first day they start producing a new style that other line chiefs have previously produced, separately for lines selected for treatment (solid squares) and lines not selected (hollow triangles). The solid vertical line indicates the start of treatment from June 2014 on, while the dashed line shows the end of the contracted upon four months treatment period on treatment floors. Capped bars represent 95% confidence intervals.

Appendix C3: 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 the treatment or control sample, respectively, are then used to calculate weights w_i for each unit i according to:

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

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 they produced new styles that have previously been produced on other lines, 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 use these weights for each production line based on equation 3 in column 3 of Table 4 to re-run the specification from column 2 in the same table, only with the data reweighted (control units are not reweighted in this approach, therefore weights w_i for lines from control floors are set to 1).

Abstrakt

Přelévání produktivity v rámci jedné firmy se běžně používá jako proxy proměnná pro proces učení se v rámci jedné organizace. V tomto článku používám data z více než 200 produkčních linek ve třech bangladéšských oděvních závodech a rozvíjím stávající literaturu dokumentující tyto přelivy, a to ze dvou hledisek. Prvním zjištěním je to, že fyzická vzdálenost uvnitř jedné firmy hraje důležitou roli pro přelévání produktivity, naproti tomu složitost výrobku nemá na přelévání téměř žádný vliv. Toto zjištění vede k velmi důležitým implikacím pro firmy v rapidně se rozvíjejících zemích jako je Bangladéš: zatímco vylepšování jednodušších výrobků na složitější má na přelévání pramalý vliv, růst firmy má vliv značný. Dalším hlediskem, kterým přispívám ke stávající literatuře je to, že využívám náhodné intervence v komunikaci ve třech závodech, abych zjistil, do jaké míry je přelévání produktivity opravdu mírou výměny znalostí uvnitř dané organizace a do jaké míry je způsobeno jinými druhy peer efektu, jako je například konkurence. Intervence spočívala v tom, že vedoucí pracovníci přikázali vedoucím výrobních linek, aby sdíleli své znalosti o výrobě v případě, kdy jejich linky vyráběly stejné textilie. Tato intervence zvýšila míru přelévání produktivity mezi cílenými výrobními linkami. Toto zjištění podporuje dosavadní pohled na přelévání produktivity jako na míru výměny znalostí uvnitř firem.

Working Paper Series
ISSN 1211-3298
Registration No. (Ministry of Culture): E 19443

Individual researchers, as well as the on-line and printed versions of the CERGE-EI Working Papers (including their dissemination) were supported from institutional support RVO 67985998 from Economics Institute of the CAS, v. v. i.

Specific research support and/or other grants the researchers/publications benefited from are acknowledged at the beginning of the Paper.

Published by
Charles University, Center for Economic Research and Graduate Education (CERGE)
and
Economics Institute of the CAS, v. v. i. (EI)
CERGE-EI, Politických vězňů 7, 111 21 Prague 1, tel.: +420 224 005 153, Czech Republic.
Printed by CERGE-EI, Prague
Subscription: CERGE-EI homepage: http://www.cerge-ei.cz

Phone: + 420 224 005 153
Email: office@cerge-ei.cz
Web: http://www.cerge-ei.cz

Editor: Jan Zápal

The paper is available online at http://www.cerge-ei.cz/publications/working_papers/.

ISBN 978-80-7343-414-4 (Univerzita Karlova, Centrum pro ekonomický výzkum a doktorské studium)
ISBN 978-80-7344-450-1 (Národohospodářský ústav AV ČR, v. v. i.)