**Working paper**
February 2016

**Azam Chaudhry, Rocco Macchiavello, Theresa Chaudhry and Christopher Woodruff**

International Growth Centre

# Understanding productivity in the Pakistani garment sector

## 1. Summary

We report on a pilot project in the ready-made garment sector in Pakistan. The pilot had two objectives. The first objective was to develop a methodology for benchmarking productivity in the garment sector in Pakistan to productivity in a group of garment manufacturers of roughly equal size in Bangladesh. The second objective was to explore issues constraining productivity in the sector in Pakistan with an eye toward developing a full-scale project on one or more constraints faced by firms in the sector.

As this report details, the benchmarking exercise proved more challenging than initially anticipated. In particular, cross-firm calibration of the productivity measures routine developed and used by firms in the sector has been challenging. However, we believe we have made substantial progress toward a reasonable methodology for doing this. We apply that methodology to the Pakistani factory with the highest degree of product overlap with factories in the Bangladesh sample, and show that the Pakistani manufacturer is about 16 percent more efficient than the most efficient of the Bangladeshi factories in the comparable sample. We discuss the current state of the work, and the next steps in the benchmarking project.

With regard to the second objective, numerous conversations with factory managers and other industry stakeholders indicated that quality issues were a concern in the sector. We conducted a detailed survey with line-level managers in the factories. The survey focused on quality issues, and particularly on quality control processes in the factories. We have only recently completed and compiled these data, and will analyse them in the coming months with the goal of identifying a specific set of issues and an intervention to address those issues.

**When citing this paper, please use the title and the following reference number: F-37119-PAK-1**

# 2. Introduction

Pakistan is ninth largest exporter of clothing outside the EU, with exports in 2014 totalling USD 5.0 billion.[1] Further development of the garment sector is viewed by the Pakistani government as beneficial because garment production is labour-intensive and hence exports result in the creation of jobs. One constraint on growth is that wage rates in Pakistan are higher than competing countries such as Bangladesh and Cambodia. To remain competitive, Pakistan will need factories which are more productive than those in the lower-wage countries.

The objective of this project was to pilot a method for benchmarking the productivity of a set of ready-made garment factories in Pakistan to the productivity of roughly similar producers in Bangladesh. With this goal in mind, we collected production data from nine large producers in Punjab province. In several factories, the data proved to be inadequate for analysis, either because the factories did not collect sufficient information or because they were reluctant to share information on key variables. Even for the six factories providing reasonably complete information, calibration proved more challenging than initially anticipated. In this report, we describe the process of calibrating production data across factories, first within the (larger) sample of factories in Bangladesh, and then between one factory in Pakistan and the sample in Bangladesh. Work on calibration continues, but we are now less sanguine about the prospects for producing highly accurate calibration across countries.

# 3. Measuring productivity: Introduction and issues

Sustained increases in income per capita depend on increases in productivity. Recent research suggests that in lower-income countries, higher variance in productivity co-exists with lower mean productivity. There is now a large literature addressing productivity differences across countries. (See, for example, Hsieh and Klenow, Quarterly Journal of Economics, 2009; De Loecker and Goldberg (Annual Review of Economics, 2013.) Firm-level data do not often include measures of physical productivity. As a result, most studies measure revenue productivity. As De Loecker and Goldberg discuss in some detail, differences in revenue productivity reflect both differences in physical productivity and differences in mark-ups across firms. De Loecker and Goldberg also discuss various structural methods for recovering physical productivity. A key advantage of this

1. WTO Trade Statistics for 2014. http://stat.wto.org/StatisticalProgram/WSDBStatProgramSeries.aspx?Language=E , accessed 31 January 2016.


DIRECTED BY
FUNDED BY

study is the ability to measure output quantities in a way that is comparable across factories. The key, which we discuss below, are measures of the sewing content in each product produced by each factory. Our key challenge, which this report describes in detail, is making the measures of sewing content fully comparable across factories. Once we are able to do that, we will be able to compare physical productivity without the need for any assumptions about the nature of the firms' optimization or their market power.

The ready-made garment sector is an interesting place to study productivity for other reasons as well. First, garment production is common in a large number of low- and middle-income countries. Garment production has been the starting point for export-oriented manufacturing in many countries experiencing rapid growth in recent decades. Internationally, the sector is characterized by very large retail brands. The brands are sometimes integrated backward into production, but more commonly source product through relationships with independently owned producers. Brands often source the exact same product from multiple suppliers, sometimes located across multiple countries.

Within typical woven and light-knit factories, the sewing section employs between half and two-thirds of the workforce. As we note above, in theory, the efficiency of sewing operations can be closely compared even across products. In both Pakistan and Bangladesh, sewing is typically organized in sewing lines. Production is sequential, with each operator performing a specific task. The product is assembled as it moves down the line.

In a series of projects, we have collected very detailed production data from several dozen factories in Bangladesh. Our data measure line-level production on a daily basis, and include information on the quantity of each style produced on the line during the day, the number of production workers on the line, hours worked, absenteeism and quality defects. In the Bangladeshi sample, the data are continuous for a period of around one year. We were able to collect roughly comparable data from six factories in Pakistan, but for a period covering at most three months. Moreover, while the data are comparable, differences in product mix make direct comparison with the Bangladeshi factories problematic for most of the six Pakistani factories, as described below. We are presently working on ways to address these problems.

## Calibrating productivity within factories: SMVs

For most of the factories, the style data can be matched to documents referred to as operation breakdowns that developed as the basis for laying out of the production line. The operation breakdown details every step of the production process and assigns an allowed time to complete each step undertaken by either a sewing machine operator or helper. The allowed time is referred to as the Standard Minute Value (SMV), or the Standard Allowable Minute (SAM). The SMVs are summed across all of the production steps to obtain an SMV for the garment as a whole. The garment's SMV provides a precise estimate of the number of minutes required to produce the garment at full efficiency. This allows for the calculation of a quantity-based measure of output.

A key challenge for factories is to "balance the line." Since each operator is dedicated to a single step in the production, the rate of production of the full product will depend on the speed of the slowest process. For example, if the operator sewing the side seam works more rapidly than the operator sewing the bottom hem, then the product will pile up behind the operator producing the bottom hem. Industrial engineers and line managers have several ways to balance the line. First, they can assign multiple operators to a given step. If the side seam operator produces half again as many operations as the bottom hem operator, then the factory may assign two operators to side seams and three to bottom hems. The factory may also assign more experienced, faster, operators to the bottom hem and newer, slower, operators to the side seam. Helpers – additional workers without machines – may be assigned before or after a given step, for example to line up two pieces of material that are to be sewn. Finally, more efficient machines, for example, those with automatic thread cutting, may be deployed on critical processes.

The standard measure of efficiency used in the sector takes the ratio of the minute value of production – the SMV multiplied by the number of completed pieces – and the available labour minutes – the sum of the number of minutes worked by each worker on the line.[2] In formula:

[Quantity of Output * SMV] / [(# Operators + # Helpers) * runtime in minutes]

Inefficiency arises from two fundamental causes: individual operators take longer per operation than the SMV allows, and unbalanced production means that some operators are idle as they wait for other operators to complete steps that have higher SMVs. Inefficiency can also arise from other sources, for example, because quality defects mean that processes need to be undone and redone.

The SMV / SAMs for each process can be estimated using a standardized database or through time-and-motion studies within the factory. International databases such as the General Sewing Data[3] are used by some factories. Customized databases which are nevertheless constant across products within the factory are used by some factories. Other factories perform time-and-motion studies using operators on special sample lines to determine the SMVs on a product-specific basis. Within a factory, the SMVs generally allow for very close comparison of productivity across products. There may be some noise in the product-specific productivity measures, for example, in factories using repeated time-and-motion studies. But after careful analysis of the productivity data in numerous Bangladeshi factories, we find that the product-level SMVs show no patterns along a number of

2. We discuss below concerns with this measure that complicate the calibration of productivity across factories.

3. See, for example, http://www.whichplm.com/news/gsd-introduces-standard-method-libraries-for-the-apparel-industry.html . Accessed 1 February 2016.

observable characteristics, such as the final buyers, product types or the line on which a garment is typically produced. However, different factories may use different standards, and these standards need to be taken into account when comparing productivity across factories.

## Efficiency and helpers

All factories for which we have data employ helpers on sewing production lines. The helpers perform a variety of tasks, most commonly cutting the thread after an operator completes the stitch, aligning the fabric before the operator starts the stitch, or moving product from one place to another. The helper is assigned an SMV for her tasks as well. The typical measure of efficiency reported by factories adds the helper to both the numerator – the output minutes associated with the helper's SMV – and the denominator. So long as the helper-to-operator ratio is reasonably constant across products within the factory, this will not affect the within-factory measure of efficiency. But although all factories employ some helpers, the helper-to-operator ratio varies quite widely across factories, from one helper to around 20 operators in one Pakistani factory to more than one helper for each operator in several Bangladeshi factories.

The different helper-to-operator ratios are an additional complication in the comparison of efficiency across factories. The importance of the helper-to-operator ratio can be seen in a simple example Suppose Factory A employs no helpers and Factory B employs a helper for each operator. Suppose the two factories produce exactly the same product, and consider one process, for example the bottom hem, for that product. In Factory A, the operator may be allocated 0.35 minutes (21 seconds) to complete the bottom hem and cut the thread. In Factory B, the helper cuts thread after the stitch is completed. This saves the operator time – she does not have to stop sewing, pick up scissors, cut the thread, and then return to the next stitch. Therefore, she may be allowed only 0.30 minutes (18 seconds) for the stitch. But the helper will also have an SMV and the combined SMV of the operator and helper is likely to be higher than the 0.35 minutes. In the standard measure of efficiency, this will increase the number of output minutes. So even though the two factories produce exactly the same process, Factory B will show a larger number of output minutes. In the efficiency calculation, it will be rewarded for using the helper.

Whether the use of helpers increases efficiency or not depends on whether the helper increases the efficiency of the operators by a sufficient amount. Since helpers are typically paid around 60 percent as much as a skilled operator, if a helper takes 15 seconds to line up fabric before a stitch and in doing so saves the operator 10 seconds by doing so, then the use of the helper will be cost-efficient. Within a factory, the ratio of helpers to operators is not likely to vary much across products. Hence, the standard efficiency measure may provide a reasonable comparison of efficiency across lines and products within the factory. However, the ratio varies quite substantially across factories, and hence complicates further the comparison of efficiency across factories. What we need, and what we aim for, is instead a measure of efficiency based on the cost per output minute.

Recall that one common task of helpers is cutting the thread after the operator completes a stitch. One reason that a factory may employ fewer helpers is that they use sewing machines with automatic thread cutting, eliminating the need for either the operator or a helper to perform this task. Thus, even a fully-calibrated labour cost per output minute might be criticized as reflecting only differences in capital used in production. We did not gather information on the sewing machines in Pakistan, but we did in Bangladesh. In the calibration there, we are able to control for whether machines have automatic thread cutting, for example.

## Calibration of data across factories

With these issues in mind, we proceed with a calibration using data from the line layouts constructed by factory industrial engineers. The operation breakdown document shows each step in the production process and the SMV allocated to that step. The document also includes the product type, style name, buyer and other information related to the order. We proceed in the following manner:

1. For each of several broad product categories – t-shirts, trousers, shorts, etc. – we identify a set of processes which are common and comparable to essentially every style in the product category. For example, in t-shirts, there is always a side seam, a bottom hem and a sleeve attach. There are also operations related to the neck, but we do not include these in the list because there is a difference between round necks and v-necks and while the operation breakdown usually indicates the type of neck, it does not always do so. We include about 65 percent of the processes in a typical t-short in the calibration exercise.
2. We next select a single factory to be the benchmark factory. For a few factories we have a very comprehensive database of processes and SMVs, and we select one of those factories as the benchmark factory. Since we do not yet know how the SMVs used by the benchmark factory compare with the international standard SMVs, we view the resulting calibration as relative rather than absolute. That is, we are able to compare the factories within our dataset to one another, but are not yet able to say how these compare to the international standards.
3. A key step in the calibration is to adjust for differences in the helper-to-operator ratios across factories. We proceed as follows: Each of the common processes selected for use in the calibration is sometimes performed without any helpers and sometimes performed with helpers. We use the information in the operation breakdown to assign helpers to specific process. Where a helper is shown between two processes, we use the description of the helper's task to assign the helper either to the process immediately before or immediately after the helper. For example, a helper assigned to "thread trim" and placed just after the bottom hem is assigned to the bottom hem process. A helper whose task is "align fabric" is assigned to the process that follows the helper. We then identify the styles where no helper is assigned to a give process. We use these styles to estimate a "helper-free SMV" for that particular process. For each style / process with helpers, we assign the "helper-free SMV" in place of the operator and any helpers associated with the process. This yields measures which are comparable across factories with

regard to the ratio of helpers to operators.

4. The final step is to adjust for the variation across factories in standards used to set the SMV. We aim to calculate a single adjustment factor which places the SMV of a given factory on the same scale as the SMVs used by the benchmark factory. To do this, we sum the original SMVs of all of the steps, and then sum the "helper-free" SMVs for the same set of steps. The ratio of these two becomes the adjustment factor for the SMVs for the entire product. Table 1 below gives a simple example, using only two steps. The key implicit assumption is that the ratio of the SMVs in the common steps is the same as the ratio of the SMVs in the style-specific steps (e.g., hanger loop attach).

**Table 1**

| Process | Original SMV | Standard "no-helper" SMV |
|---|---|---|
| Bottom Hem | 0.35 | 0.45 |
| Trimming after bottom hem | 0.15 | - |
| Sleeve Join | 0.95 | 0.84 |
| **Total:** | **1.45** | **1.29** |

In the above example only three process were matched and the sum of original and efficient SMVs is 1.45 and 1.29 respectively. Therefore, a ratio of 0.89 (1.29 / 1.45) is applied to the entire SMV for this style.

The steps lead to something akin to an exchange rate with which we can translate the output minutes from each factory into the output minutes of the benchmark factory. Having output minutes in a "common currency" allows us either to calculate the standard measure of efficiency or to calculate the more useful measure of cost per unit output minute.

## A note on identical styles

One of the projects in Bangladesh was carried out in a group of suppliers to a single large European buyer. That buyer routinely places orders for the same style with multiple factories, and we find some instances of the same buyer / style name in more than one factory. This allows us to compare the operation breakdowns and the SMVs across factories for the same product, though we should stress that we can do this only a limited number of products. The results are nevertheless surprising and somewhat sobering.

We take one example of two factories, which we label 3001 and 4002, where we find three styles that appear to match precisely, and for which each factory has more than 50 days of production. For each of the three styles, factory 4002 shows SMVs which are considerably higher than 3001. This likely reflects in part the fact that 4002 has around one helper for each operator, while 3001 has around 0.6 helpers per operator. But the ratio of SMVs is 1.33 for one product, 1.23 for the second product, and 1.08 for the third. The variance in this ratio across products gives us pause, and suggests that perhaps there is noise in the assignment of SMVs at the product level. It could also be the case that, while the descriptions of the styles are the same, the styles themselves vary in subtle ways. On the positive side, we note that the average ratio of SMVs across the three styles is 1.21, which is almost identical what we will find with the calibration exercise described above. This suggests, at least, even with noise at the product level, the calibration across a larger number of products could certainly be reasonable. But it does suggest that we will be more confident in the calibrations when they are based on a relatively larger number of products, an issue which will sometimes prove to be important.

## 5. Calibration of t-shirt factories in Bangladesh

The calibration method described above was implemented first among a set of factories in Bangladesh producing a large number of t-shirt styles. T-shirts are perhaps the easiest product to calibrate, both because many factories in one of our Bangladeshi datasets produce a large number of styles and because the basic processes are relatively consistent across styles.

Table 2 shows the results of the calibration exercise for six factories. Column B shows the raw efficiency for each factory, calculated using all of the available production data, which in this data set is typically daily for a period of one year. These are large factories, with an average of 30 production lines. Data for the first factory starts on January 1, 2013, with coverage from most factories from March 2013 through February 2014, and the last data point on May 31, 2014. As noted in column H, the factories produce a large number of different styles of t-shirts during the period covered by the study, with at least 40 in each factory. The calibration exercise uses SMVs for between 450 and 2823 different processes.

## Table 2

| A | B | C | E | F | G | H | I | J | K |
|---|---|---|---|---|---|---|---|---|---|
| fact_code | raw efficiency | helper-operator ratio | ratio* | sd | number of process used | number of available processes | number of availble styles | normalised efficiency | adjustment |
| 1004 | 0.471 | 0.494 | 1.037 | 0.078 | 2308 | 3200 | 135 | 0.488 | 0.02 |
| 1008 | 0.521 | 0.787 | 0.939 | 0.113 | 2823 | 4575 | 202 | 0.489 | -0.03 |
| 3001 | 0.491 | 0.404 | 1.113 | 0.107 | 450 | 983 | 46 | 0.546 | 0.06 |
| 3006 | 0.462 | 0.941 | 0.784 | 0.168 | 487 | 1035 | 40 | 0.362 | -0.10 |
| 4002 | 0.577 | 1.256 | 0.877 | 0.107 | 729 | 1487 | 47 | 0.506 | -0.07 |
| 4004 | 0.458 | 0.473 | 0.914 | 0.176 | 1021 | 2081 | 85 | 0.419 | -0.04 |

The raw efficiency data, calculated as the ratio of output minutes to available minutes, is shown in column B. The range is large, with the factory reporting itself as most efficient (4002) being 26 percent more efficient than the factory reporting itself to be least efficient (4004). However, column C gives a hint that the un-calibrated data might not reflect reality: the helper-to-operator ratio varies widely across the factories, with 4002 having the highest ratio of 1.26, indicating more helpers than operators. Factory 4004, on the other hand, has less than half a helper per operator. The correlation between the raw efficiency and the helper-to-operator ratio is +0.71 – factories with more helpers are also more efficient by this measure – though note that this is for only six factories.

The key calibration factor is the ratio of the helper-free SMV for the factory, relative to the benchmark factory, which we take to be factory 1004.[4] This is shown in column D. The calibration ratio shows that the standards for setting SMVs vary quite substantially across factories. The ratio is defined as the average SMV in the benchmark factory relative to the average SMV in the benchmark factory. Hence, the 1.11 for factory 3001 indicates that 3001 has a tighter standard than the benchmark factory, allowing less time for the matched processes. Factory 3006, on the other hand, has a much more lenient standard, with a ratio of 0.78 indicating that it sets SMVs some 22 percent higher than factory 1004. Column E shows the average within-style standard deviation for each factory. A larger number indicates that the factory changes the SMV for a given process across styles to a greater degree. Note that the standard deviations indicate particularly large within-style variance for factories 3006 and 4002.

4. The ratio in 1004 is not 1.0 because even in that factory we adjust the helpers out of the production.

Finally, column I shows the calibrated efficiencies, which is just the product of columns B and D. The variance in this small sample is actually increased by the calibration, with the range of 0.46 to 0.58 in the raw data increased to 0.36 to 0.55 in the calibrated data.

## Pakistani data

The preceding description indicates the complexity of comparing efficiency across factories, even when those factories produce very similar products. The first challenge we face in comparing the six Pakistani factories with more or less complete data is that the Pakistani factories are specialized in the production of blue jeans, a product which is very uncommon in our current Bangladeshi samples. As jeans are by far the largest product in Pakistan, the longer-term solution to this issue is to bring on board more factories producing jeans in Bangladesh. In the short-run, we aim to demonstrate the method of comparing efficiency using a single Pakistani factory producing t-shirts. Even with this factory, there are several issues which lead to some uncertainty in the comparison. However, all measures that we construct lead to a similar conclusion that the Pakistani factory is substantially more efficient than any of the sample of factories from Bangladesh. We conclude that while the uncertainty gives us many issues to work on, there is some promise even in the more difficult task of comparing factories across countries.

We identified 40 styles of t-shirts in the Pakistani factory. The operation breakdowns show no helpers, so all of the processes are already in the "no helper" format. The SMV calibration uses 322 processes. The SMV standards appear to be very similar to the benchmark factory in Bangladesh, with a calibration factor of 1.03, suggesting a very slightly tighter standard in Pakistan. To this point, we feel fairly confident of the calibration exercise: the styles and processes are similar, and the SMVs appear sensible and comparable. However, there are two issues with the production data from this factory which leave us with some uncertainty, and a range of estimates. We discuss each of these in turn.

## Limited production data

The factory has limited production data provided in the most complete format. There is data for only a single production line, and for only a single month. These data show a raw efficiency of 82 percent. Given the calibration factor of 1.03, this yields a calibrated efficiency of 84 percent, far higher than any of the Bangladeshi factories. But in general, there is quite a lot of variation in efficiency within factories, across lines and months. For example, monthly line-level efficiency in the benchmark factory ranges from 27 percent to 68 percent, with the 10 / 90 range being 37 percent to 58 percent. We've no way of knowing where in the Pakistani factory's overall distribution this particular line-month measure lies.

A separate file provides the output data for many line across three months. These data provide only one of the four components of the efficiency measure, so using them requires finding the other three. We are able to match operation breakdowns through a work order number for many but not all of the products shown in the more extensive data. If we take a very strict standard of a precise name match, we obtain 55 line-days of data for which we can calculate the output minutes. If we take a slightly more relaxed standard (for example, matching "140" with "140A"), we obtain 185 line-days of data.

### Available minutes data

The second issue relates to the data for the denominator of the efficiency calculation. The output file does not include either the number of workers or the runtime for the line. However, separate HR data indicate daily clock-in and clock-out times for each worker in the factory, including the worker's factory ID number. These can be matched to monthly salary data which indicates the line on which the worker is deployed. From these two data sets, we can find the number of available minutes at the line-day level. An initial comparison with the one line-month for which we have full data indicates the need for two adjustments. First, the HR data include managers as well as workers, and so we remove the managers. Second, the HR data indicate there are some helpers working on the lines, even thought the operation breakdown documents do not show any helpers. Since we have helper-free SMVs, we include the helper time in the denominator in our preferred estimate, but we also report the efficiency calculations without the helpers in the denominator.

### Benchmarks

With these issues in mind, we estimate the Pakistani factory to be 64 percent efficient using the most restricted sample of 55 line-days with exact product matches, and 68 percent efficient using the 185 line-days with product matches somewhat more liberally defined. These percentages compare with the range of 36-55 percent in the sample of Bangladeshi factories. If we exclude the helpers from the denominator, the calibrated efficiency increase to 75 percent (restrict sample of 55) or 76 percent (sample of 185). By any measure, then, we estimate that the Pakistani factory is at least 16 percent more efficient (64 / 55 percent) than the most efficient of the Bangladeshi factories.

## 6. Discussion and next steps

We report here on a pilot project which had the goal of testing a method of benchmarking productivity among garment factories in Pakistan against a set of factories in Bangladesh. The work revealed both that a reasonable calibration is possible and that the work required to make the calibration is much more difficult than we initially anticipated. With regard to the full project, we are currently processing the data from around 25 factories in Bangladesh across several product categories, including t-shirts, shorts, and dress shirts. These will provide us with calibrated data within product categories. Some factories produce many styles in each of more than one product category. We intend to use data from these large multi-product factories to calibrate the data across product groups. In doing so, we will need to assume that multi-product factories use the same SMV standard across products, and/or that they are equally efficient in each product category in which they are active producers. We continue to believe that the resulting dataset will be unique in its level of detail on productivity, and that there are a range of interesting question which these data will be uniquely suited to answer.

## 7. Factory-level surveys in Pakistan

As a part of the pilot project, we conducted surveys with supervisors in the participating projects, focusing on the issue of quality control. The survey instrument was composed of two parts. The first roughly half of the survey was modelled on questions asked of supervisors in Bangladesh surveys. These questions on basic demographics and labour history of the supervisors, practices on the line and attitudes and factory norms. The second half of the survey focuses on quality control issues and procedures. Here the questions are unique to Pakistan. We have now entered the data from this survey and are in the process of analysing these data. At this point we have nothing on this to report, but we expect the first half of the survey may provide a descriptive comparison of supervisor demographics and line-level practices in Pakistan and Bangladesh. We will examine the second half of the survey with a goal of understanding issues which might become the focus of a larger follow-up project.

The International Growth Centre (IGC) aims to promote sustainable growth in developing countries by providing demand-led policy advice based on frontier research.

Find out more about
our work on our website
www.theigc.org

For media or communications
enquiries, please contact
mail@theigc.org

Follow us on Twitter
@the_igc

International Growth Centre,
London School of Economic
and Political Science,
Houghton Street,
London WC2A 2AE

International
Growth Centre

Designed by soapbox.co.uk