Gender Wage Gaps and Worker Mobility: Evidence from the Garment Sector in Bangladesh.

Andreas Menzel (CERGE-EI Prague) and Christopher Woodruff (Oxford)[#]

May, 2020

Abstract:

Data from 70 large export-oriented garment manufacturers in Bangladesh show that gender wage gaps are similar to those found in higher-income countries. Women's wages are 20 percent lower than men's and are 8 percent lower among narrowly-defined production workers; a significant gap remains even after controlling for very precisely measured skills. Longer careers of men in the sector explain around half of the wage gap, with the other half due in roughly equal parts to differences in internal and across-factory promotions. Our results are most consistent with broader gender norms, beyond gendered household responsibilities, driving the gap.

Keywords: Gender Gaps, Export Manufacturing, Garment Sector

JEL Code: O14, O15, J7

[#] Christopher.woodruff@qeh.ox.ac.uk; Andreas.Menzel@cerge-ei.cz. The data used in this paper were collected from garment factories in Bangladesh as a part of several projects led by the authors and other researchers. The fully anonymized data and replication-files can be obtained from the authors. We thank the UK Department for International Development-Economic and Social Research Council Growth Research Programme, the Growth and Labour Markets in Low Income Countries Programme, the International Growth Centre, and the Small and Medium Enterprises Program at Innovations for Poverty Action (IPA) for funding of projects that enabled the collection of the data. Funding for the analysis was also provided by the European Research Council Advanced Grant 669746 (RMGPP). Andreas Menzel also thanks the Czech Academy of Science (GACR) for financial support through Standard Grant 17-26395S. We thank Chris Burningham for research assistance and seminar participants at PSE, Oxford, Michigan, HKU, SITE– Stockholm, CERGE-EI Prague, ESMT- Berlin and Yale for comments. Errors remain our own. Finally, the work would not have been possible without the cooperation of the factories providing the data, and the amazing data and survey teams based at IPA-Bangladesh who collected and made the data usable.

1. Introduction

A bourgeoning literature seeks to determine the underlying causes of gender pay gaps (Blau and Kahn 2017). The explanations of the gaps can be broadly classified into three strands. First, different preferences of women and men may lead them to make different job choices, or to behave differently on the job (Cook et al. 2018, Azmat and Ferrer 2017, Flory et al. 2015). Second, women may be subject to greater constraints on their choices of, or flexibility within, jobs, due to, for example, household and family demands (Goldin et al. 2014, Bertrand et al. 2010, Corcoran et al. 2005, Black 1995, Manning 2013, Card et al. 2016). Third, the gaps may reflect outright discrimination (Sarsons 2017, 2019, Egan et al. 2017, Mengel et al. 2018, Boring 2017, MacNell et al. 2015, Hengel 2018, Card et al. 2018, Goldin and Rouse 2000).

However, most of the literature on gender pay gaps is based on data from high-income countries. Evidence from developing countries is much sparser. This owes in large part to the lack of high-quality data in developing countries. Nevertheless, we may expect gender wage gaps to be as prevalent in developing countries. Indeed, many of the related potential drivers, such as norms, preferences and expectations, are even more pronounced in many developing countries (Jayachandran 2015, 2019, Giuliano 2017, Duflo 2012, World Bank 2012).

We address the data issue by using administrative records for more than 80,000 workers from 70 large garment export factories in Bangladesh. The salary data come from electronic wage records of the factories, and have several advantages over more widely available labor survey data. First, our administrative data identify the worker's precise occupation in the factory, and record wages as well as overall income. The administrative data are also likely more precise than data reported in labor surveys.[1] Moreover, for a subset of workers, we have

[1] The 2017 Bangladesh Labor Force Survey (LFS) includes 1,295 workers reporting salaries in the garment sector. Around 75 percent report income to the nearest 1,000 BDT, and the data mix regular and overtime pay. This adds noise as overtime pay varies across months and workers. Huynh (2016) and Abras (2012) use labor survey data to estimate wage gaps in several Asian countries, including

very detailed skills assessments that allow us to measure wages conditional on skills. Finally, our sample of more than 80,000 workers from a single sector is much larger than is typical for labor force surveys. This allows us to conduct a much more precise analysis of the gender wage gap in a sector that is crucial from a development policy perspective.

We use the data to make three contributions to the literature. First, we show that the wage gap in these data is quite similar to that found in the large literature from higher-income countries. Men earn, on average, around 20 percent more than women. When we exclude the four percent of workers in supervisory positions, the wage gap drops to around eight percent. These levels are very similar to those found in the U.S. for the unadjusted wage gap, and for the gap after controlling for worker's industry and occupation (Blau and Khan 2017).

Second, our unusually precise skills data allow us to control for worker ability to an extent that non-experimental papers have so far struggled to do. Our ability measure captures the raw skills that the workers offer the employer, allowing us to study the extent to which the same skills provide men higher wages than women.[2] We estimate that skills explain up to a half of the wage gaps, but significant gaps remain conditional on skills, suggesting that women capture less of the surplus their work creates.

Third, we combine evidence from surveys of a nearly representative subset of the workers and a model to study the underlying mechanisms that drive the gap. The model allows us to estimate otherwise unobservable wage increases when workers move between factories.

Our evidence is most consistent with gendered norms limiting women's commuting options and possibilities to move between employers, hurting their bargaining position with factories

---

Bangladesh. Both find positive wage gaps for women in the Bangladeshi sector. We discuss this in more detail below. Boudreau et al. (2019) use retrospective labor histories from a household survey to show that garment workers in Bangladesh tend to move to factories with better working conditions but lower salaries.

[2] Our ability measures differ from the productivity measures used in other sector-specific case studies, such as for example billable hours among lawyers in Azmat and Ferrer (2017). A concern with these measures is that they may be affected by factors such as customer discrimination or household constraints, even conditional on the worker's underlying skills.

(Manning 2013), in line with other recent research on the constraints that commuting puts on women in South Asia (Borker 2018, Sajjad et al. 2017). We find little evidence that gaps change with the marital status of women, and no evidence that gaps are larger for women with children. Though marital or parental status is not exogenous, these results are not in line with recent evidence from higher income countries (Kleven et al. 2018,19).

Our paper also speaks to the growing literature that focuses on the dynamics of wage gaps over workers' careers (Barth et al. 2019, Goldin et al. 2017, Bronson and Thoursie 2019, Albrecht et al. 2018). While the literature from higher-income countries uses data that allow tracking of individual workers across time, our data track workers only within their current factory. However, around 5 percent of all male workers, and 4 percent of all female workers in our sample switch factories each month.[3] Anecdotal evidence indicates that these movements often come with wage increases. Because we do not follow workers when they change factories, we do not observe directly any cross-factory salary changes. However, by combining the administrative and survey data with some structure, we can decompose total wage growth into the share associated with movement across factories and promotions within factories, while controlling for selection of workers out of the sector. We find that early in a worker's career, wages grow almost twice as fast for men. Moreover, the share of external promotions among all promotions is higher for men than for women. These cross-factory promotions are concentrated at early stages of workers' careers for both women and men, consistent with established findings from the United States (Topel and Ward 1992). At later stages of workers' careers, almost all wage growth is within-factory, and wage growth is no longer significantly different for women and men.

[3] These high mobility rates are consistent with evidence from export manufacturing sectors in other developing countries. In Ethiopian factories, Blattman and Dercon (2018) find that 77 percent of new hires in unskilled factory jobs leave the factories within one year of employment. The garment sector consultancy Impactt reports monthly worker exit rates of 7-12 percent in garment factories in China, Bangladesh and India (Impactt 2011, 2012, 2013).

The fact that women in the sector earn less even conditional on fine skills measures suggests the possibility that they have weaker bargaining positions vis à vis factories. Weaker bargaining positions can stem from a number of sources. For example, women may have fewer employment options outside the garment sector. Indeed, data from the 2017 labor force survey of Bangladesh show that 30 percent of women in full-time non-agricultural employment in Bangladesh work in the RMG sector, compared with only 6 percent of men. However, if a lack of employment options outside the sector drives the results, we might expect women to pursue their careers at least as vigorously inside the sector as men, for example by moving as much as men between factories to gain promotions. The results from our data and model do not suggest that this is the case.

Within the sector, men may face lower costs of changing factories and, indeed, we do find that men change factories more often. However, in contrast to recent findings in a number of high-income countries (Kleven et al. 2018, 2019), we do not find evidence that either the gender mobility or wage gap is driven by the marital or parental status of women. Instead, we find that, compared to men, women change factories more often in areas with higher density of factories. This suggests that women are more likely to move between factories when the move does not require longer commuting times, and where changing factories does not require changing residence.

Thus, our results are more consistent with women's bargaining power being weakened by broader gender norms that restrict women's access to labor markets regardless of their marital of parental status. Such norms are still widespread in the South Asian context, and are often motivated by concerns for the safety and "modesty" of women when moving in the public (Dahr et al. 2019, Dean and Jayachandran 2019, Field et al. 2019, Munoz-Boudet et al. 2012).

The fact that marital status is correlated with male career behavior but not women's also suggests that gender norms may directly affect work effort and career ambitions of workers of both sexes. Bertrand et al. (2015) find that many women in the U.S. curtail their career ambition to not out-earn their spouses. For the Bangladeshi garment sector, Macchiavello et al. (2020) show that managers and workers of both sexes believe that women are less capable supervisors in the factories, even though objective measures do not support these beliefs. Such dynamics could reduce the career ambitions of women in the sector, resulting in lower wage growth both internally and through movement across factories, consistent with recent research showing that minority workers in French supermarkets reduce work effort when working under biased managers (Glover et al. 2017).

We begin by describing the data we use in the analysis in more detail. We next examine the wage gaps between the genders, first using the full administrative sample, and then the subsample of the data that allows us to control for the worker skills. We then briefly describe the model that allows us to back out wage growth rates when workers move between factories. We conclude by discussing evidence on how marriage, childbearing, and the density of factories around worker's workplaces affect the wage gaps and worker mobility of men and women in the sector.

2. Data

The ready-made garment sector is the largest manufacturing sector in Bangladesh, accounting for 80 percent of Bangladesh's exports and around 12 percent of GDP. The 4,000 factories in the sector employ 4 million workers, more than half of them women. With growth compounding at an annual rate of around 15 percent over more than two decades, the garment sector dramatically increased the share of women working in full-time wage jobs. We study workers in a sample of 70 large export-oriented garment factories located in and around Dhaka, the larger of two major production areas in Bangladesh. All of the factories in our sample produce woven or light knit (e.g., t-shirts) garments. Production is typically organized into three sections: cutting, sewing, and finishing. The sewing sections employ around two-thirds of the workers in these factories, and women represent the large majority of workers in the sewing sections, around 80 percent in our data. The organizational hierarchies in the sewing sections are highly comparable across factories, making this a uniquely suitable sector to study the interplay between management practices, workforce dynamics and firm-level outcomes. We thus focus on sewing workers for the remainder of the paper.[4]

We draw data from three sources: administrative payroll records from the factory's human resources (HR) departments, surveys of randomly selected workers, and skills assessments of workers conducted by the industrial engineering departments of the factories. We briefly discuss each of these three types of data in turn.

2.1 Payroll records

We collected monthly payroll records from the factories. Each factory's data cover a period of at least six months, with a mean and median of 11 months. The data come from factories participating in various research projects, with the earliest records from January 2012 and the latest from December 2017. The payroll records have one observation per worker-month and contain all workers employed by the factory for at least one day of the given month. For the majority of factories, our data are limited to the sewing sections and, for around half of the factories, to non-supervisory workers. These workers are classified by pay grades ranging from 7 (entry-level helpers) to 3 (highly skilled operators). Grade 2 typically captures line supervisors, while grade 1 higher-level supervisory staff. The Bangladeshi minimum wage law for the garment sector proscribes a minimum wage for each worker grade, though workers on a given grade are often paid slightly more than the minimum level. Throughout the text we use the term "promotion" to represent a movement up in grades – for example from grade 5 to

[4] The proportion of women in the cutting and finishing sections is typically lower. Some factories also have knitting, dying, and embroidery sections. These sections are typically more capital intensive, and the share of male workers higher.

grade 4. The monthly payroll records usually contain the worker's name, an ID number, the date the worker joined the factory, wages for regular hours, overtime earnings, some measure of absenteeism during the month, and a designation of the job performed by the worker.

The lowest grade (7) is assigned to unskilled, entry-level, workers. In the sewing section, these are typically referred to as "helpers" whose job is to cut thread or line up fabric for sewing operators. Workers operating sewing machines ("operators") are assigned to grades 6 through 3, with 6 being the lowest-skilled and 3 being the highest-skilled operators. Essentially all workers enter the sector as a grade 7 workers and advance to higher grades as they gain skills. The minimum wages in effect from December 2013 through the end of our data were around US$65 per month for grade 7 workers (for 6 days per week of 8 hours) and around US$85 for grade 3 workers. Overtime pay beyond 8 hours per day is paid 1.5 times the hourly rate implied by the minimum wage. Panel 2 in Table 1 shows the distribution of female and male workers in grades 7 through 3, indicating that men are overrepresented on higher grades.

We track the workers over the repeated rounds of monthly pay records within factories with the worker ID numbers assigned by the factories.[5] The administrative records typically do not include the gender of the worker, so we use names to code the gender of each worker. Because some names may be either male or female, we drop around 10 percent of observations for which we are not able to designate a gender.

We use the payroll data to define variables for promotion and exit from the factory. A worker is promoted in month t if her/his grade is higher in month t+1 compared with month

[5] Some factories reuse ID numbers from workers that have left the factory, so more than one worker may have the same ID at different points in time. Factories may also assign a new ID number when a worker is promoted to a new grade. Thus, the same worker can have more than one ID. One or the other of these problems arose in around 30 of the 71 factories in our sample. For these factories, we used worker-name and join-date combinations to assign a single unique identifier to each worker, dropping the few observations with the same name / date-of-join combination.

t.[6] A worker is deemed to have exited a factory if s/he disappears from the data before the last month for which we have data from the factory. By these definitions, we find that the monthly promotion rate is around one percent, while the monthly exit rate is 4.9 percent. A naïve interpretation of the promotion rate implies that workers move up one grade every eight years, a rate of progression much lower than that implied by our survey data. We return to this issue below. Meanwhile, 4.4 percent of workers on grades above entry level (grade 7) are new to the factory in a given month. These workers almost certainly have previously worked at another factory, reflecting the substantial rate of movement between factories in the sector.

One potential issue with both of these measures is that we lack data on supervisors in 34 of the 70 factories. In these factories, workers promoted to grade 2 or 1 will disappear from our dataset and will be recorded as having exited the factory rather than being promoted. In aggregate, we view it as a minor issue because promotion above grade 3 is a rare event. In the 36 factories for which we have supervisory data, only about 7 percent of all male (and 0.5 percent of female) promotions end with the worker in a supervisory grade or higher. Thus, given that we miss these promotions in roughly half our sample, 3.5 percent of all male promotions may not be recorded (and just short of 1 percent of male exits incorrectly recorded).

The salary records provide us with wage and grade data, and we use both of these in the analysis. However, we note that wage and grade are not independent of one another. A common practice in this sector is for the production staff, industrial engineers in particular, to determine an appropriate wage for a worker, and then for the human resources staff to set an appropriate grade conditional on that wage. We find it convenient to focus on wages for part of the analysis and on grades for another part. But the two should be viewed as co-determined.

[6] We observe a small number of cases in which a worker appears to have been demoted. Most often, the worker soon reverts back to the previous grade. We assume in these cases that the initial demotion reflects a mistake by the factories. In the promotion analysis, we drop these few cases, though the results are not affected if we leave them in the data.

2.2 Skills assessments

Production in all of our factories is organized on production lines, with each sewing operator performing a single stitch in the sequential production process. Almost all of our factories employ industrial engineers who assign a "Standard Minute Value" (SMV) to each process. The SMV represents the time a fully efficient worker would take to complete the task. SMVs are typically based on international standards, adjusted for factory conditions – for example, the presence or absence of a helper, or whether the machine has automatic thread cutting or not. One measure of worker skill that factories care about is efficiency measured as the number of minutes of sewing output a given operator is able to produce divided by the number of minutes a fully efficient worker would produce.

Most factories conduct regular skills assessments of operators, and we obtained assessments from 20 factories.[7] We use the skills assessments to calculate four measures for each worker. First, we record the number of sewing processes on which a worker is tested as a measure of flexibility. Second, an industrial engineer with substantial industry experience working for our project team coded the "criticality" of each process on a scale of one to seven.[8] We create a variable which indicates the level of the most critical process on which the operator is tested. Third our industrial engineer also flagged the processes that require physical strength to complete. We create a variable indicating the worker is tested on at least one process requiring a high level of physical strength. Finally, we measure the average efficiency of the worker on all skills tested. These four skills measures are, according to factories, highly correlated with the productivity of workers. The skills map more closely to the value of

[7] Most of these are from factories participating in a project implementing a consulting intervention which includes training on conducting regular skills assessments.

[8] Some factories record processes as "A" (most critical), "B", or "C" (least critical). Our industrial engineer separated the A ratings into high-A, middle-A and low-A, and created similar tiers for the B ratings. In his judgment, the C-rated processes were all of the same level of difficulty. This measure is highly negatively correlated with the target for the process, indicating that processes rated as more critical are more complex and take longer.

marginal product of a given worker than measures more commonly available in data, such as wages, providing a more precise control for productivity.

2.3 Survey data

In addition to the salary records and skills measures, we have survey data from a sample of randomly selected sewing machine operators in each of the factories. The surveys were conducted to support other projects, and yield a total sample of 2,170 surveyed workers. The surveys contain several variables useful in our analysis: tenure in the garment sector, the number of factories in which the respondent has worked, and demographics such as age, years of schooling and marital status. Administrative factory records often don't record worker age and almost never record any of the other measures available in the survey. A subset of 1,037 surveys from 48 factories also contain information on children of the respondents.

The survey samples contain operators on grades 6 through 3 selected randomly from given production lines. Workers of the lowest grade (7) are not included because they were not part of the original studies for which the data were collected. The production lines from which the operators are sampled were selected by the factory to participate in these studies and so may not represent average lines in the factories. However, as shown in Table 1.b, the differences in the administrative variables between the subsample of surveyed workers and the much larger sample of workers not surveyed are very small in magnitude. In seven out of 10 comparisons the difference is smaller than 0.1 standard deviations, and it is never larger than 0.15 standard deviations. Furthermore, as we show below, the wage gap among surveyed workers is almost identical to that found among all workers in the HR data.[9]

[9] In most factories, workers are not sorted by skills across lines, but rather to tasks within lines. Lines are kept interchangeable to maximise flexibility to finish orders coming in on short notice, and to deal with frequent production interruptions. This is reflected in the fact that line fixed effects explain less than 3.5 percent of the variation in HR variables aside from overtime, where the line fixed effects explain 7.3 percent of the variation. Given that usually in the factories whole lines are assigned to overtime, as opposed to individual workers, the larger explanatory power for overtime is less surprising.

2.4 Summary Statistics

The factories in our sample are all direct suppliers of major European and American brands. They are largely locally owned and managed. Brands contract directly or through intermediaries with the factories. Table 1 provides summary statistics for our data. The top panel shows summaries on the factory-level. The factories in our sample have an average (median) of 1,294 (1,071) sewing section workers. The smallest factory has 133 sewing workers and the largest around 4,400. In total, our data contain 144,500 individual sewing-line workers, 91,800 of whom are present in the data from the first month from each factory.[10] For around 81,500 of them we have non-missing information on sex, grade and wage in the first month. Thus, these workers form our primary sample for our analysis. Among these workers, 79 percent are female, a share that varies between 54 percent and 98 percent across the factories in our sample. Though information on the full distribution of factories in Bangladesh is limited, Figure 1 compares our sample factories with all factories registered through Accord and Alliance.[11] These data indicate that the full distribution of Accord / Alliance factories is represented in our sample, though larger factories are over-represented.

Panel 2 of Table 1 shows summary statistics on the worker level from the administrative records, using the sample of workers present in the first month of data available from each factory. Workers have spent on average about two years in the current factory, and their monthly base pay, not including overtime pay, bonuses, or deductions for absenteeism, is 5,950

[10] In total, we have around 1.1 million observations at the worker-month level from workers of grades 7-3. Since our data do not allow us to track workers across different factories in our sample, we cannot rule out that some of the 144,500 unique workers are indeed the same person working at different factories in our sample. Our factories are a very small part of the industry, however, and are geographically scattered, so we expect that this would be rare.

[11] The Alliance for Bangladesh Worker Safety and the Accord on Fire and Building Safety in Bangladesh were both formed after the Rana Plaza factory collapse in 2013. These two organizations cover around 2,000 factories (roughly half of the factories active in the country) that are primary suppliers to Accord or Alliance buyers. Boudreau (2019) provides a more in-depth discussion of these two organizations.

BDT, a bit less than US$ 70.[12] About one-third (34 percent) of the workers that we observe in the first month of data exit the factory before the last observed month, while 8.2 percent received an internal promotion to a higher grade, with Table 1 showing these numbers disaggregated by gender and grade.

The third panel of Table 1 shows worker-level statistics by gender from the sample of workers for which we have surveys. Our surveys all come from sewing machine operators in grades 6 through 3. The surveyed workers are 72.6 percent female (the equivalent share in the HR data for workers in grades 6 through 3 is 74.5 percent). Their average age is 25.2 years, 78 percent are married, and 76 percent report having at least one child. They report having worked on average 3 years in the current factory, 6.3 years in the sector. Men report having worked in 2.5 other garment factories prior to the current factory, while women report having worked in only 1.5 other factories.

For comparison, we show data from the 2017 Bangladesh Labor Force Survey (LFS). We take the sample of respondents in the LFS who report working full time in private firms in the apparel sector (ISIC 1410) in firms with 250 or more employees. We limit the sample to those reporting occupation codes corresponding to "sewing and embroidery machine operators." The LFS sample is quite comparable to ours, though our sample is slightly more female (73 vs. 66 percent) and just more than a half year younger. Males in our sample have slightly fewer, and females slightly more, years of schooling.[13] Our sample of workers also has characteristics very similar to those in the large-scale representative sample described in Haque et al. (2015).

[12] Throughout the period, the exchange rate was around 80 BDT per US$.

[13] Curiously, a simple regression on the LFS sample of machine operators shows a substantial and significant *positive* wage premium for females, consistent with the results reported by Abras (2012) and Huynh (2016). It is unclear why the household data should suggest wage gaps so different from those we find in our administrative data from factories, particularly since a comparison of the data on years of schooling suggests that our sample of women is relatively more positively selected compared with the LFS sample.

While the workers in our sample are comparable to garment workers in the country's labor force survey, selection of workers into the garment sector as opposed to other sectors is also a concern for measuring the gender pay gap. Using data from the 2011 Bangladesh Population Census, Figure 2 shows that among male garment workers, average education rates coincide with average education rates from broader population, not just in the cross section, but also for men of different ages. For women, average education rates are slightly below average rates from the census for the different birth cohorts, though the differences are minimal.

3. Gender Wage Gaps

We follow the empirical approach that is standard in the literature on gender wage gaps, running a regression of the form:

$$LnWage_{if} = \alpha Female_i + \delta_f + \epsilon_{if}$$

where *Female* is a dummy variable indicating that worker i is female. We include fixed effects, and cluster standard errors, at the factory level. We use the data for all workers recorded as working during the first month we have data from each factory. Results, however, are qualitatively unchanged throughout if we use the data from all months or for all 140,000 workers who are present in any month.

We start in Table 2 with the restricted set of 36 factories for which we have wage data for all sewing section workers, including supervisors. The grade 2 supervisor workers are 93 percent male and constitute 4.1 percent of all sewing workers on grades 2 through 7 in these 36 factories. Column 1 reports a basic regression with only factory fixed effects. We find that female workers earn about 20 log points less than male workers. On the right half of the table, in column 5, we repeat the same regression among the sample of workers in grades 3 through 7 using data from all 70 factories. Removing the (mostly male) supervisors leads to a drop of the wage gap to around 8 log points. The smaller gap is not due to the factories in the expanded sample of 70 being different than those in the sample used in column 1; even in the sub-sample of 36 factories for which we have grade 2 workers in the data, the wage gap is 7.7 log points when the sample is limited to non-supervisory workers (column 4). The 8 percent earnings gap among our production workers is very similar to the 8.4 percent wage gap conditioned on industry and occupation that Blau and Khan (2017) report for the U.S.

In columns 2 and 6 of Table 2 we add grade fixed effects to regressions using the restricted sample with supervisors (column 2) and the full sample (column 6). Given the highly structured minimum wage laws governing the sector in Bangladesh, we should expect to find that controlling for worker grade will reduce the wage gap significantly. Indeed, we find the gap drops to only 1.2 percent in the sample of 36 factories (column 2) and 1.5 percent in the larger sample (column 6). These within-grade differences are not affected by interacting the factory and grade fixed effects. Thus, the wage gap is substantially a "grade gap". Men work on higher grades, but conditional on the grade, the wage gap is small (even though it remains statistically significant). Columns 3 and 7 confirm this by estimating the "grade gap" directly. We regress the ordinal grade level of each worker – reversed, so that a positive coefficient reflects a higher grade level – against a dummy indicating the worker is female plus factory fixed effects. Using the restricted set of factories that includes supervisory workers, we find that the grade of women is, on average, 1.18 levels lower (column 3); limiting the sample to grades 7 through 3 and using the full set of 70 factories, the gap is slightly smaller, 0.79 levels (column 7).

3.1 Controlling for Absenteeism and Overtime

Men may earn more because they are perceived to be more reliable or more flexible in working hours. The administrative data for 27 of the 70 factories contain data on both worker absenteeism and overtime hours. Appendix Table A.1 shows that women are on average absent on 50 percent fewer days per month and are more likely to receive an attendance bonus.[14]

[14] Attendance bonuses are typically small bonuses of around 5 percent of the monthly base wage, paid each month the worker was not absent more than 0, 1, or 2 days, depending on the factory.

However, women also work 4.7 percent fewer overtime hours than men. Women thus seem to be more reliable workers, albeit maybe less available for certain overtime needs of factories, resembling situations also observed in high-income country settings (Goldin 2014). In Appendix Table A.2 we show that controlling for the average absenteeism or overtime hours of workers leaves the estimated wage and grade gaps unchanged, suggesting that neither absenteeism nor overtime hours explain any part of the gender gaps shown in Table 2.[15]

3.2 Controlling for Worker Skills

We have worker skills data for sewing operators (Grade 3 through 6) from 20 factories for which we have HR data. Skills captured by one or more of these measures is the most often mentioned criterion for promotion in written promotion policies available from several factories. In the few factories that assign weights to the criteria, these skills account for half or more of the weight. The promotion policies also routinely mention factory tenure and attendance. As discussed above, Appendix Table A.2 shows that attendance and overtime are not related with wage or grade gaps.

In Column 1 of Table 3 we replicate the basic grade gap regression of column 7, Table 2, for operators from these 20 factories. This shows a grade gap of 0.34 for this sample.[16] In column 2, we verify that the three main skills measures that we have for all 20 factories – the number of processes a worker can do, the complexity of the most complex process a worker can do, and whether he or she can do processes that require physical strength – predict the grade of the worker. The results show that particularly the first two are significantly associated with higher grades. A one standard deviation increase in number of processes a worker can do is associated with a 0.21 higher grade level, whereas the same number for the complexity of

[15] It also does not affect the estimated wage gap conditional on grade fixed effects.
[16] The drop from 0.79 grades shown in column 7 of Table 2 is almost all due to dropping grade 7 workers from the sample. The coefficient for grade 6 through 3 workers from the full sample of 70 factories is -0.376. As the data in Table 1 show, women are particularly over-represented on grade 7 positions.

the most complex process the worker masters is 0.16. Workers able to do at least one process requiring physical strength have a 0.11 higher grade level, which however is not statistically significant. As we cluster standard errors on a relatively small number of 20 factories, we use wild-cluster bootstrap, as suggested by Cameron et al. (2008). The skill measures therefore seem to capture meaningful variation in worker ability.

We next ask how much of the gender gap is accounted for by differences in skills. We add gender to the regression including the measures of skills in column 3, plus a set of indicator variables for the types of machines the worker is skilled on. A comparison on columns 1 and 3 shows that adding the skills measures reduces the estimated grade gap by a third, to 0.24 grades. The drop is not affected by adding our fourth skill measure, efficiency of a worker on those processes s/he is skilled on, a measure that is available for only 16 factories (column 4). The results are virtually unchanged when we add the squares of the skill measures, and, as we show in Appendix B.1, are also visible among workers within narrowly defined skill groups. The drop in the gender gap in column 3 suggests that females are less skilled in at least some dimensions. We examine this in columns 5 to 8. We find that male workers are reported to be able to perform significantly more complex tasks, and tasks that require physical strength.

The bottom half of Table 3 replicates the results from columns 1-4 using log wage instead of grade as the dependent variable. A significant wage gap of around 3.5 percent is visible among the workers for which we have skill data. Once we control for the different skill measures, the remaining wage gap drops by around a quarter, slightly less than for the grade gap. The grade hierarchy provides a more homogeneous measure across factories of the career progression of workers, and so we focus in the next sections on understanding in more detail the evolution of the grade gap over the average career of men and women in the sector.

Collectively, the results on Table 3 show that as much as a third of the gender gap is explained by the lower skill acquisition of women, but the larger share comes from lower grades or wages conditional on the level of skills. Thus, equally skilled women seem to be on lower paid grades on average in the factories.

4. Grade Gap Decompositions

Because almost all workers in the sector start their career at grade 7,[17] a worker's current grade indicates the number of promotions s/he has received to reach that grade.[18] This allows us to decompose the gender grade gap along two dimensions. First, we can decompose the gap into the part coming from differences in average time spent in the sector and the part coming from differences in the rate of promotion. Second, we can decompose the part coming from differences in promotion rates into that arising from differences in promotion rates within factories and promotions received when moving from one factory to another. We refer to promotions in the factory as internal, and promotions when moving across factories as external. These decomposition exercises are useful as they can inform on the potential mechanisms driving the wage and grade gaps we observe, which we will discuss further in the next section. Furthermore, they inform on the evolution of these gender gaps over the course of worker careers in the sector, which is of policy interest in its own right (Barth et al. 2019, Goldin et al. 2017, Bronson and Thoursie 2019, Albrecht et al. 2018).

4.1 Time in sector vs Promotion rate

The administrative data report factory tenure but not sector tenure. We know sector tenure only for the subset of workers who were surveyed. To estimate the portion of the gender grade gap coming from differences in the average sector tenure, we run the grade gap regression on the survey sample and add a control for sector tenure. Table 4, column 1 shows that the basic gender grade gap among the surveyed workers is 0.375 grades, almost exactly the 0.376 grade

[17] Of all workers in the survey, 91 percent report to have started their career on grade 7.
[18] We occasionally observe workers in the HR data rising more than one grade in a single promotion. We treat these as multiple promotions occurring at the same time.

gap among grade 6 to 3 workers in the overall HR data. Column 2 of Table 4 adds years in sector as reported by the surveyed workers, and its square. The inclusion of these variables reduces the estimated grade gap to 0.251, or by around 33 percent. Adding additional worker observables as controls reduces the grade gap only marginally (column 3, Table 4).

Note that the sample of surveyed workers from our factories does not include grade 7 workers. This is unfortunate because a large share of the grade / wage gap is accounted for by differential promotion rates from grade 7. We address this by randomly sampling from the HR data a number of female and male grade 7 workers so that when these grade 7 workers are added to the sample of surveyed grade 6 to 3 workers, the share of grade 7 workers in that combined sample matches the share in the HR data. We then assume that sector tenure and factory tenure is the same for grade 7 workers – that is, that workers do not move between factories without promotions on grade 7. Column 4 of Table 4 shows a grade gap of 0.80 in the combined sample of grade 7 to 3 workers without controlling for sector tenure, close to the grade gap of 0.79 estimated from the full HR data in column 7, Table 2. Column 5 shows that adding again sector tenure (approximated by factory tenure for grade 7 workers) and its square explains about 50 percent of the overall grade gap we found in column 4.[19]

The skills data and differences in sector tenure appear to explain around 35 and 50 percent of the grade gap, respectively. Does that imply that together they explain 85 percent of the gap? The answer to the question is almost surely "no", because skills accumulate with experience. The skills gap and the tenure gap are likely two ways of measuring related factors. We have both skills and survey data only for a small sample of 154 workers in nine factories, and the grade gap is not statistically significant in this small sample, though it is of comparable

[19] Data from surveys of 190 grade 7 workers in four factories not included in our data suggest that grade 7 workers do sometimes make lateral movements across factories. In appendix B.2 we show that, using reasonable assumptions, adjusting for these movements does not materially affect the results from Table 4. If anything, the 50 percent of the grade gap explained seems to be an upper bound to the true share explained by sector tenure (see column 6, Table 4).

magnitude than the gap in the overall survey sample. Controlling for sector experience reduces the wage gap by around 25 percent (column 2) among workers of grades 6 to 3 in this sample, while controlling for both skills and sector experience reduces it by around 50 to 60 percent (Table C.1, comparing columns 1 and 5).

4.2 Internal vs External Promotions

We observe internal promotion frequencies directly in our HR data, and we begin by presenting the evidence on them. Table 5 shows the monthly internal promotion rates, overall and by grade, for male workers, and the difference for female workers. Promotion frequencies decline as workers climb the grade hierarchy. Internal promotion rates out of grades 7 and 6 are around 2.25 percent per month for men (column 1), falling to 1.28 percent at grade 5 and to less than 1 percent per month for the more skilled grades. Internal promotion rates are 35 percent (grade 7) to 10 percent (grade 4) lower for women. This gender difference is larger and more statistically significant when controlling for factory-month fixed effects. This suggests that female workers select into factories with higher internal promotion rates, though this effect is not statistically significant at conventional levels. Furthermore, the lower promotion rates for women also depend on conditioning on worker grade. Overall, at any point in time, the share of female and male workers in the sector being promoted is the same. However, this reflects the fact that women make up a larger share of workers on lower grades (Table 1), where promotion rates are higher.

4.3 Backing out external promotion rates

We do not observe external promotions in our data, as our data do not trace workers across factories. Moreover, simple back-of-the-envelope calculations of external promotion rates are likely to be severely distorted. For example, simply dividing the number of promotions implied by a worker's current grade by sector tenure in the survey data indicates a monthly (annual) promotion frequency for both women and men of around 4 percent (50 percent). Given that around 1 percent of workers receive an internal promotion each month, a naïve comparison would thus suggest that three quarters of all promotions are external. However, this comparison fails to account for the fact that the surveys do not include workers exiting the sector. To take an example, suppose at the beginning of each period of time, a share s of workers is promoted to the next grade, while the remaining share *(1-s)* leaves the sector at the end of the period. A survey of a cross-section of workers conducted at any point in time would capture only surviving workers, and show an average promotion rate of 1 rather than s. With precisely measured internal promotion rates, this will result in an overestimate of external promotion rates unless the probabilities of exit from the sector and promotion are independent of one another, which seems unlikely given that promoted workers are less likely to exit the sector.

For this and other likely confounders in simple back-of the envelope calculations of external promotion rates, we write down a model that explicitly incorporates these. As we show in more detail in Appendix D, we can then recover (otherwise unobserved) external promotion rates of workers on each grade by fitting the model to the empirical distribution of (female and male) workers across grades 7-3, and the average and variance of sector tenure of workers on each grade reported in our survey data. We summarize here the basic results obtained from fitting the model.

Table 5 shows the external promotion rates for men for each grade obtained from fitting the model, and the differences to those of women, alongside the internal promotion rates from the HR data we discussed above. Comparing the two, we see that the external promotion gaps are larger than the internal promotion gaps, indicating that external promotions explain more than half of the grade gap that is not explained by longer careers of men in the sector. External promotions predominantly occur at the early stages of careers in the sector, when workers are on the lower grades 7, 6, and 5, resembling established evidence from higher income countries that career advancement through movements between employers predominantly occurs at earlier stages in the careers of workers (Topel and Ward 1992). On grades 7 to 5, external and internal promotions are roughly similar for female workers, while for men external promotion rates are around 60 percent higher than internal rates on grade 7, and 10-30 percent higher on grades 6 and 5. The overall (external + internal) promotion rate at the lowest grade 7 is almost twice as high for men, with both the external promotion rate and the gender gap in overall promotion rates gradually decreasing to more or less zero for promotions from grade 4. Thus, there is a strong correlation between the gender gap in overall promotion rates on a grade, and the share of external promotions among all promotions on the grade, in particular for men.

4.4 Selection out of Garment Sector and the Wage Gap

We conclude this section with a brief discussion of how selection out of the sector affects the estimated gender gap. If lower-ability workers leave the sector faster (slower), average pay will correlate more (less) strongly with tenure in a cross-section of workers. Differences in the relationship between ability and sector exit could differ for women and men could thus account for (parts of) the widening grade gap over time. A simple test for differential selection is to compare worker observables for workers on the highest grade 3 with the average worker, and to see if that pattern differs between men and women. Table 6 shows such a test on education, being married and having children. Generally, we do not find strong effects, though for education, the point estimate of the difference between grade 3 workers and other workers is positive for men while negative for women, with the difference significant at the 10 percent level. This indicates that men who remain in the sector are more selected in education than women (column 1, Table 6). Note, however, that including education as control has very little effect on the estimated grade gap, as shown in Table 4. We do not find evidence for differential selection of male and female workers on marriage or children. Also, we do not find any evidence that the selection process has changed over the time covered in our sample: when we interact the gender indicators with time trends (columns 2, 4, and 6 of Table 6) for education, being married, and having children, the interaction effects are insignificant.

5. **Drivers of Grade and Promotion Gaps**

The finding that female workers are on lower pay grades even when controlling for fine measures of skills suggests that they capture a lower share of the surplus their employment generates, possibly due to weaker bargaining positions of female workers. As we show in Appendix C.2, skills explain a very similar share of the gender gap in internal promotions as they do for grade levels, though statistical power is weaker for promotions. This underlines the notion that women are less able to translate skills into promotions and higher wages in negotiations with factory management. Bargaining power is determined by outside options of workers in standard bargaining models. The main outside option of garment workers is employment in other factories. Movement to other factories could be costlier for women due to gendered household or parenting burdens. These could for example restrict their maximum commuting time, thereby limiting the number of alternative factories at which they could look for work.

5.1 External Constraints on Women: Household and Childcare duties

In Table 7, we first examine the extent to which marriage can explain the grade gap by comparing married and unmarried workers. Around 70 percent of male workers and 80 percent of female workers in our sample report being married. We expect to find that married women have more household responsibilities than single women and so face higher mobility frictions. Column 1 of Table 7 regresses worker grade on a female worker dummy, controlling for tenure in the sector and the factory, average years per factory, age and education level. Column 2 adds a control for being married, interacted with worker gender to allow differential effects for women and men. The marriage controls themselves are highly insignificant and their inclusion reduces the wage gap by less than 10 percent.

Columns 3-4 and 5-6 of Table 7 repeat columns 1 and 2 using internal promotions and years spent per factory over a worker's career, respectively, as dependent variables. We use average years per factory over the worker's career as a proxy for otherwise unobserved external promotions, as this captures the combined rate of external promotions and lateral movements between factories.[20] In the sample of surveyed workers, women are 0.42 percentage points less likely to receive an internal promotion each month, conditional on grade.[21] This gap falls slightly to 0.38 percentage points (column 3) after controlling for basic worker observables. Women change factories at a 30 percent slower pace than men, and hence spend an average of 0.7 more years in each factory over their career (column 5). Controlling for being married, allowing the effects of marriage to differ by gender, reduces the baseline gender gap in internal promotion rates by almost half, leaving it statistically insignificant (column 4). If anything, though, this effect is driven by married men receiving more promotions then unmarried men; married and unmarried women receive internal promotions with equal frequency. The pattern is similar for our proxy of external promotions, as shown in column 6. Married men move significantly more frequently than unmarried men, while married and unmarried women move at roughly the same rate. The data suggest marriage leads to changes in career behavior for men rather than women. This result resembles those from an older literature from high-income

[20] The number of factories a worker has worked in over his or her career is a highly statistically significant predictor of grade. This suggests a strong relationship between mobility and external promotions. This relationship is marginally stronger for women (p-value = 0.1).

[21] This compares to 0.50 percentage points in the HR data. The worker's grade is an outcome of her accumulated promotions, so grade is endogenous to promotions and may therefore be a "bad control". Running Columns 3-6, Table 7 without grade fixed effects does not change the results fundamentally. On the other hand, the distribution of female and male workers across grades differs significantly (as shown in Table 1), as does the promotion pattern across grades. Therefore, not including grade fixed effects potentially confounds the coefficient estimates by strong compositional effects.

settings that showed that being married boosts the careers for men (for recent summaries of that literature see de Linde Leonard and Stanley, 2015, and Sobel, 2012).

A further potentially important factor is whether workers have children. A growing literature using data from high-income countries shows that the gender wage gap is increasingly found to be a mother wage gap (Kleven et al. 2018, 2019, Adda et al. 2017). Our surveys included questions on children in 48 of the factories, which contribute around 45 percent of all surveyed workers. Thus, columns 7-12 of Table 7 repeats columns 1-6 of the same Table on the sample from these 48 factories, adding controls for having any child and any child younger than five years. We again allow the effects of children to differ by gender in columns 8, 10 and 12.[22] None of the gaps on the three outcome variables grade, internal promotions, and years per factory, is affected in a meaningful way by adding these controls.[23]

5.2 Gender Norms beyond household roles: Movement restrictions

Female worker's outside options may also be reduced by norms that affect women regardless of their parental or marital status. For example, they may have fewer employment options outside the garment sector, due to gendered norms on what employment is appropriate for women (Dean and Jayachandran 2019). Data from the country's 2017 Labor Force Survey show that, among women without tertiary education, 30 percent of those working full-time outside of the agricultural sector work in the garment sector; the comparable number for men is only 6 percent. However, this channel does not explain why women move less often between

[22] 73 percent of men and 77 percent of women report having at least one child, while 53 and 34 percent, respectively, report having a child younger than 5 years.

[23] Children might affect the grade gap indirectly by inducing some women to exit the sector. This exit could be differentially selected on ability. Given that we do not observe workers who have left the sector, this channel is difficult for us to quantify. However, we note that in our worker surveys, 80 percent female workers have children by age 25. This is close to the 90 percent share found in LFS data. Furthermore, age 25 is the peak of the age distribution among women in the sector, suggesting that women do not generally drop out of the sector with the birth of their first child.

factories within the sector, and accumulate fewer external promotions in that process, as indicated by the results from our model.

This result is more consistent with general gendered norms limiting the access to, and movement within the labor market for female workers. Such norms are still widespread in the South Asian context, and are often motivated by concerns for the safety or "modesty" of women in the world outside the household (Dahr et al. 2018,19, Munoz-Boudet et al. 2012). In particular, lengthier commutes could present safety concerns for both female workers and their relatives (Borker 2018; Sajjad et al. 2017), limiting the number of factories women can access for employment. Similarly, family households may be less willing to move to follow women if they switch employer, particularly if their jobs provide less income (Costa and Kahn, 2000). We test for these mechanisms in Table 8 by regressing again our key outcome variables on a female worker indicator variable and factory fixed effects, plus an interaction term of the female worker dummy with the number of factories within a 2km radius around the factory, which is the maximum distance female workers typically report to commute to work. We use the addresses of factories that are members of the Alliance or Accord to geo-code these almost 2,000 factories. For the 63 factories from our sample that we could locate in the maps, there are between zero and 66 Alliance and Accord factories within 2km distance. We compress the highly right-skewed distribution using a log hyperbolic sine transformation, and control for interactions of factory age and workforce size with the female worker variable. We find that women have relatively higher exit rates from the factories, and relatively higher external arrival rates (the share of women on grades 6 through 3, who likely already worked at other garment factories before, joining the factory in a given month) when there are more factories located nearby.[24] Thus, relative to men, women are less mobile across factories if there are fewer

[24] Results are qualitatively similar when using a 1km or 3km radiuses instead of 2km.

factories located in close vicinity. These apparent constraints to commuting might limit female worker's bargaining power with factories.

6. Discussion and Conclusion

We use detailed administrative records from 70 large garment factories in Bangladesh to show that, even within a very narrowly defined set of occupations, men are paid on average around 8 percent more than women. We conduct a series of analyses that reveals several clear patterns in the data. First, women are paid less than men even conditional on very detailed skills. Second, male workers both have faster internal wage growth, via within-factory promotions, as well as external wage growth, from moving between factories and entering new factories at higher pay grades. This external wage growth is concentrated at early stages of worker careers in the garment sector, when men move more frequently between factories. Third, the mobility of female workers is related to the number of factories in walking distance to their current employer, but does not appear to be related to marital status or childbearing.

These three findings are best tied together by "monopsony in motion" mechanisms (Manning 2013). Female workers in the factories have fewer outside options, both in other factories in the sector, and outside the sector. This reduces their bargaining power with factories and thus their ability to translate skills into higher wages. The lower returns on skills could reduce incentives to acquire skills in the first place, which in turn reduces the career prospects and time spent in the sector. At higher skill levels, at which factories rely mostly on internal promotions, the gender gap in promotions disappears. However, the early-career differences across gender have important cumulative effects: an estimated 59 percent of males but only 32 percent of females that enter the sector ultimately reach one of the two highest operator grades, as suggested by our model when fitted to the data (see Appendix D, Table D.1).

This mechanism is not inconsistent with gender norms also affecting career ambitions directly, through channels other than constrained outside options of women (Bertrand et al.

2015). Our results that married men purse their careers more vigorously indicates this possibility. Macchiavello et al. (2020) document widespread negative beliefs that women are less able garment supervisors, even though objective measures do not support these beliefs. But strong glass ceilings persist in the sector, with women making up the majority of workers on each operator grade, but less than 7 percent of supervisors. This could signal to women that efforts in career advancement carry lower returns for them, curtailing their career ambitions.

It is less obvious that taste-based discrimination plays a primary role, at least at the level of operators.[25] The garment sector worldwide has always employed primarily women for the positions studied in this paper, due to their perceived higher dexterity or "docility" (Bhattacharya and Rahman 1999, SOMO 2011). This makes it unlikely that bias against women *per se* leads the managers to offer lower wages. Taste-based discrimination is also inconsistent with the lack of gender differences in internal promotion rates at the highest-skilled, highest-paid operator grade, where external promotions are virtually absent. Thus, the faster career advancement of men on lower grades seems more connected to their ability or willingness to change factories more frequently on these grades, and less to intrinsic taste for male operators.

Export manufacturing is viewed as one of the few development strategies that has generated sustained economic growth in low-income countries. However, there are also widespread criticisms of export manufacturers for not safeguarding labor (and environmental) standards. Given this, we consider this setting to be of particular policy interest.[26] The size of the sector, and its employment of large numbers of female workers in a context where female labor force participation rates are otherwise low is representative of export sectors more generally in many

[25] The experiment in Macchiavello et al. (2020) also suggests that beliefs rather than taste-based discrimination underlie the lack of promotion of women to supervisory positions.
[26] Blattman and Dercon (2018) and Boudreau (2019) provide excellent summaries of the academic literature on labor standards in export manufacturing. For the broader debate on the role of export manufacturing in development, see also the ISID (2019).

developing countries. This makes our setting a particularly interesting place to study the effects of export manufacturing on many broader societal outcomes of interest, such as the wage gap.

References

Abras, A.L.G. 2012. "Success and upgrading after the end of the MFA." In G. Lopez-Acevedo and R. Robertson (eds): Sewing success? Employment, wages, and poverty following the end of the Multi-fibre Arrangement (Washington, DC, World Bank), 87–135.

Adda, Jerome, Christian Dustmann, and Katrien Stevens. 2017. "The Career Costs of Children," *Journal of Political Economy*, 125 (2), pp. 293–337

Albrecht, Jim, and Mary Ann Bronson, Peter Thoursie, and Susan Vroman. 2018. "The Career Dynamics of High-Skilled Women and Men: Evidence from Sweden." *European Economic Review* 105: 83-102

Azmat, Ghazala, and Rosa Ferrer. 2017. "Gender Gaps in Performance: Evidence from Young Lawyers*." Journal of Political Economy* 125 (5): 1306-1355

Barth, Erling, Sari Kerr, and Claudia Olivetti. 2019. "The Dynamics of Gender Wage Differentials: Evidence from Establishment Data." NBER WP 23381

Belloni, Alexandre, Victor Chernozhukov, and Ying Wei. 2016. "Post-Selection Inference for Generalized Linear Models with Many Controls." *Journal of Business & Economic Statistics* 34 (4): 606-619

Bertrand, Marianne, Claudia Goldin, and Lawrence F. Katz. 2010. "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors." *American Economic Journal: Applied Economics* 2 (3): 228–55

Bertrand, Marianne, Emir Kamenica, and Jessica Pan. 2015. "Gender Identity and Relative Income within Households." *Quarterly Journal of Economics* 130 (2): 571-614

Bhattacharya, Debapriya, and Mustafizur Rahman. 1999. "Female Employment Under Export-Propelled Industrialization: Prospects for Internalizing Global Opportunities in the Apparel Sector in Bangladesh", United Nations Research Institute for Social Development (UNRISD) Occasional Paper

Black, Dan A. 1995. "Discrimination in an Equilibrium Search Model*." Journal of Labor Economics* 13 (2): 309-333

Blattman, C., and S. Dercon. 2018. "The Impacts of Industrial and Entrepreneurial Work on Income and Health: Experimental Evidence from Ethiopia*." American Economic Journal: Applied Economics* 10 (3): 1-38

Blau, Francine, and Lawrence Khan. 2017. "The Gender Wage Gap: Extent, Trends, and Explanations*." Journal of Economic Literature* 55 (3): 789-865

Boring, Anne. 2017. "Gender biases in student evaluations of teaching." *Journal of Public Economics* 145: 27-41

Borker, Girija. 2018. "Safety First: Perceived Risk of Street Harassment and Educational Choices of Women," Working Paper, World Bank

Boudreau, Laura. 2019. "Multinational Enforcement of Labor Law: Experimental Evidence from Bangladesh's Apparel Sector." Working Paper, UC Berkeley.

Boudreau, Laura, Rachel Heath, and Tyler H. McCormick. 2019. "Migrants, Information, and Working Conditions in Bangladeshi Garment Factories." Working Paper, UC Berkeley

Bronson, Mary Ann, and Peter S. Thoursie. 2019. "The Wage Growth and Within-Firm Mobility of Men and Women: New Evidence and Theory." Working Paper, Georgetown University.

Cameron, Colin, Jonah Gelbach, and Douglas Miller. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *Review of Economics and Statistics* 90: 414-427

Card, David, Ana Rute Cardoso, and Patrick Kline. 2016. "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on Relative Pay of Women." *Quarterly Journal of Economics* 131 (2): 633-686

Card, David, Stefano DellaVigna, Patricia Funk, and Nagore Iriberri. 2018. "Are Referees and Editors in Economics Gender Neutral?" Working Paper, UC Berkeley

Cook, Cody, Rebecca Diamond, Jonathan Hall John A. List, and Paul Oyer. 2018. "The Gender Earnings Gap in the Gig Economy: Evidence from over a Million Rideshare Drivers." Working Paper, Stanford University

Corcoran, Mary E., Paul Courant, and Mary C. Noonan. 2005. "Pay Differences among the Highly Trained: Cohort Differences in the Sex Gap in Lawyers' Earnings." *Social Forces* 84 (2): 853–72

Costa, Dora L., and Matthew E. Kahn. 2000. "Power Couples: Changes in the Locational Choice of the College Educated, 1940–1990," *Quarterly Journal of Economics* 115 (4): 1287–1315

Dahr, Diva, Tarun Jain, and Seema Jayachandran. 2018. "Reshaping Adolescents' Gender Attitudes: Evidence from a School-Based Experiment in India", Working Paper, Northwestern University

Dahr, Diva, Tarun Jain, and Seema Jayachandran. 2019. "Intergenerational Transmission of Gender Attitudes: Evidence from India", *Journal of Development Studies* 55 (12), pp. 2572-2592

Dean, Joshua, and Seema Jayachandran. 2019. "Changing Family Attitudes to Promote Female Employment", *American Economic Association Papers and Proceedings* 109, pp. 138-142

de Linde Leonard, Megan, and Tom Stanley. 2015. "Married with Children: What Remains when Observable Biases are Removed from the Reported Male Marriage Wage Premium." *Labor Economics* 33: 72-80

Duflo, Esther. 2012. "Women Empowerment and Economic Development," *Journal of Economic Literature* 50 (4), pp. 1051-79

Egan, Mark L, Gregor Matvos, and Amit Seru. 2017. "When Harry Fired Sally: The Double Standard in Punishing Misconduct." NBER Working Paper No. 23242

Field, Erica, Rohini Pande, Natalia Rigol, Simone Schaner, and Charity Troyer Moore. 2019. "On her account: How strengthening Women's Financial Control affects Labor supply and Gender norms," NBER Working Paper 26294

Flory, Jeffrey A., Andreas Leibbrandt and John A. List. 2015. "Do Competitive Work Places Deter Female Workers? A Large-Scale Natural Field Experiment on Gender Differences in Job-Entry Decisions." *Review of Economic Studies* 82 (1): 122-155

Giuliano, Paola. 2017. "Gender: An Historical Perspective." NBER Working Paper 23635

Glover, Dylan, Amanda Pallais, and William Pariente. 2017. "Discrimination as a Self-Fulfilling Prophecy: Evidence from French Grocery Stores." *Quarterly Journal of Economics* 132 (3), pp. 1219-1260

Goldin, Claudia, and Cecilia Rouse. 2000. "Orchestrating Impartiality: The Impact of "Blind" Auditions on Female Musicians." *American Economic Review* 90 (4): 715-741

Goldin, Claudia. 2014. "A Grand Convergence: Its Last Chapter." *American Economic Review* 104 (4): 1091-1119

Goldin, Claudia, Erling Barth, Sari Kerr, and Claudia Olivetti. 2017. "The Expanding Gender Earnings Gap: Evidence from the LEHD-2000 Census." *American Economic Review P&P* 107 (5): 110-114

Haque, A.K. Enamel. 2015. "Garment Workers in Bangladesh: Social Impact of the Industry." Dhaka: Asian Center for Development

Hengel, Erin. 2018. "Publishing while Female." Working Paper, University of Liverpool

Huynh, Phu. 2016. "Assessing the gender pay gap in Asia's garment sector." ILO Asia-Pacific working paper series

Impactt. 2011. "Finding the Sweet Spot: Smarter ethical trade that delivers more for all." https://impacttlimited.squarespace.com

-----. 2012. "Nice Work? Are workers taking the strain in the economic downturn? 2006-2012." https://impacttlimited.squarespace.com

-----. 2013. "Nicer Work? Impactt's Benefits for Business and Workers Programme 2011-2013." https://impacttlimited.squarespace.com

ISID. 2019. "Working in Export Manufacturing: A Blessing or a Curse?", ISID Policy Brief PB-2019-03, by F. Amodio and A. Menzel

Jayachandran, Seema. 2015. "The Roots of Gender Inequality in Developing Countries." *Annual Review of Economics* 7. Pp. 63-88

Jayachandran, Seema. 2019. "Social Norms as a Barrier to Women's Employment in Developing Countries", Working Paper, Northwestern University

Kleven, Henrik, Camille Landais, and Jakob Egholt Sogaard. 2018. "Children and Gender Inequality: Evidence from Denmark." *American Economic Journal: Applied Economics* 11 (4): 181-209

Kleven, Henrik, Camille Landais, Johanna Posch, Andreas Steinhauer, and Josef Zweimuller. 2019. "Child Penalties across Countries: Evidence and Explanations." *AEA Papers and Proceedings vol. 109: 122-126*

Macchiavello, Rocco, Andreas Menzel, Atonu Rabbani, and Christopher M. Woodruff. 2020. "Challenges of Change: An Experiment Training Women to Manage in the Bangladeshi Garment Sector." Working Paper, CERGE-EI

Manning, Alan. 2013. "Monopsony in Motion: Imperfect Competition in Labor Markets." Princeton: Princeton University Press.

MacNell, Lillian, Adam Driscoll, and Andrea Hunt. 2015. “What’s in a Name: Exposing Gender Bias in Student Ratings of Teaching.” *Innovative Higher Education* 40 (4): 291–303

Mengel, Frederike, Jan Sauermann, and Ulf Zölitz. 2018. “Gender Bias in Teaching Evaluations.” *Journal of the European Economic Association* 17 (2): 535-566

Munoz-Boudet, Ana Maria, Patti Petesch, Carolyn Turk, and Maria Angelica Thumala. 2012. “On Norms and Agency: Conversations about gender equality with women and men in 20 countries,” Washington D.C.: The Worldbank.

Sajjad, Fizzah, Ghulam Abbas Anjum, Erica Fiedl and Kate Vyborny, 2017, “Gender Equity in Transport Planning: Improving Women’s Access to Public Transport in Pakistan,” International Growth Centre Policy Paper.

Sarsons, Heather. 2017. "Recognition for Group Work: Gender Differences in Academia." *American Economic Review*, 107 (5): 141-45.

Sarsons, Heather. 2019. “Interpreting Signals in the Labor Market: Evidence from Medical Referrals.” Working Paper, University of Toronto

Sobel, Michael E. 2012. “Does Marriage Boost Men’s Wages?: Identification of Treatment Effects from Fixed Effects Regression Models for Longitudinal Data.” *Journal of the American Statistical Association* 107: 521-529.

SOMO (2011). Gender aspects in the Latin American garment industry, www.somo.nl/wp-content/uploads/2011/04/Gender-aspects-in-the-Latin-American-garment-industry.pdf

Topel. Robert H., and Michael P. Ward. 1992. “Job Mobility and the Careers of Young Men.” *Quarterly Journal of Economics* 107 (2): 439-479

World Bank. 2012. “World Development Report 2012. “Gender Equality and Development,” World Bank.

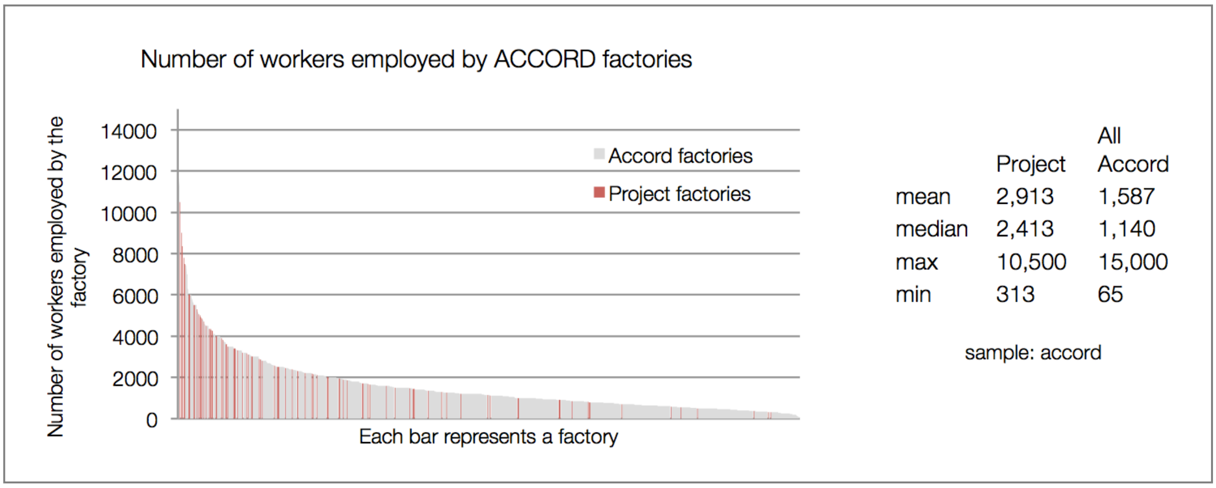

Figure 1: Project factories among the distribution of Alliance and Accord Factories: Graph shows the distribution of number of workers per factory among the around 2,000 factories organized in the buyer groups Alliance and Accord, with the factories participating in the project marked in red.

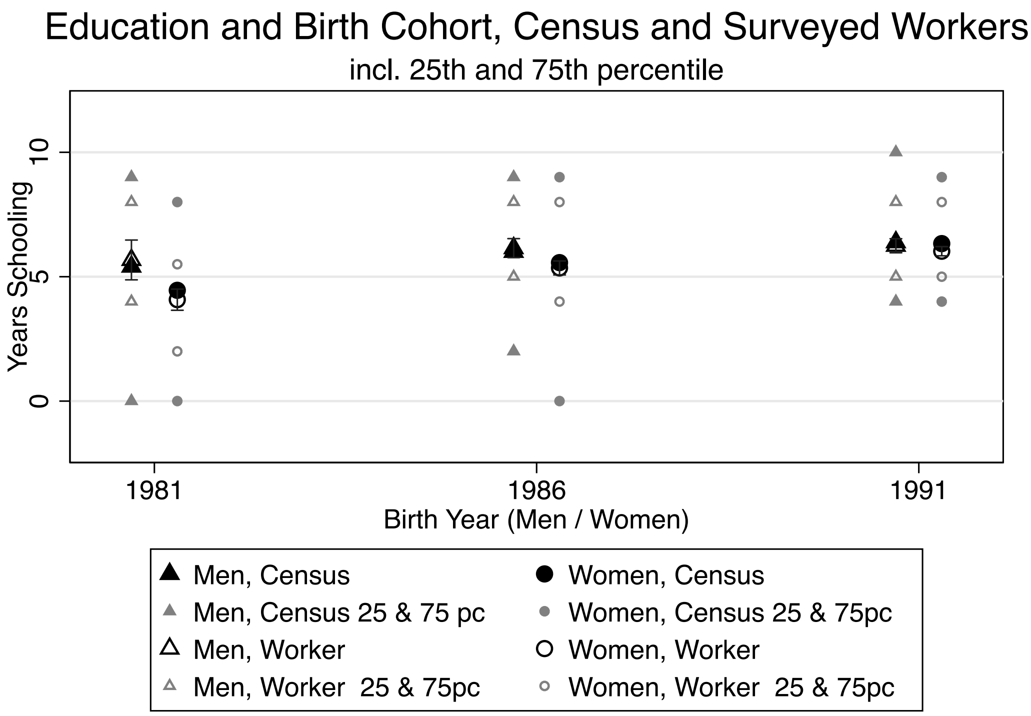

Figure 2: Selection into Garment Sector: Graph shows average years of schooling from the Bangladeshi census (solid symbols), and from our surveys of sewing workers of Grades 6-3 (hollow symbols), separately for workers born 1979-1983 ("1981"), 1984-1988 ("1986"), and 1989-1993 ("1991"). For both census and survey data, mean years of schooling, and 25th and 75th percentiles are shown. Data for men to the left (triangular symbols) and for women to the right (circle symbols). Mean years of schooling from surveys shown with 95% confidence intervals.

Table 1: Summary Statistics

Panel 1: Factory Level Statistics

Variable	N	median	mean	min	max
Factory Size, Sewing Section	70	1,071	1,294	133	4,414
Share Female Worker	70	0.80	0.800	0.545	0.983
Tenure in Factory (years, HR data)	66	2.149	2.263	0.891	6.373
Monthly Promotion rate (Not demoted)	70	0.005	0.008	0.000	0.060
Monthly Exit rate	70	0.051	0.053	0.005	0.128

Panel 2: Worker Level, HR Records, those present in first round of data

Sewing Worker Sample, Grades 7-3, with non-missing gender, grade and pay:	N	median	mean	min	max
Female	81,508	1	0.794	0	1
Tenure in Factory (years, HR data)	72,053	1.234	2.062	0	31.3
Base Wage (monthly, no overtime, BDT)[27]	81,508	6,151	5,946	752	14,023

Grade	3	4	5	6	7
Male, Share	26%	34%	15%	13%	12%
Female, Share	11%	24%	18%	19%	29%
Male, Promoted internally within 12 months	NA	3.8%	14%	17.5	16%
Female, Promoted internally within 12	NA	2.0%	8.1%	7.6%	8.6%
Male, Left Factory within 12 months	34%	35%	36%	35%	49%
Female, Left Factory within 12 months	32%	34%	35%	36%	42%

Panel 3: Worker Survey Data (73% of respondents female)

	(Mean LFS)	N	median	mean	min	max
Male Respondents:						
Age	(26.2)	594	25	25.61	18	67
Years Schooling	(6.95)	592	6	6.21	0	15
Married	(0.74)	595	1	0.70	0	1
At least one child		234	1	0.73	0	1
Tenure in Factory (years)		595	2.13	2.82	0	26
Tenure in Garment Sector (years)		595	6.25	6.90	0.16	26
Nbr. of previous Factories		595	2	2.52	0	19
Female Respondents:						
Age	(25.9)	1,572	25	25.16	16	46
Years Schooling	(5.49)	1,564	5	5.72	0	15
Married	(0.84)	1,575	1	0.81	0	1
At least one child		803	1	0.77	0	1
Tenure in Factory (years)		1,573	2.19	3.02	0	24
Tenure in Garment Sector (years)		1,573	5.03	5.88	0.11	25.1
Nbr. of previous Factories		1,573	1	1.45	0	15

27 At the time of data collection, US$ 1 ≅ 80 BDT

Table 1.b: Selectivity of Surveyed Workers

	(1)	(2)	(3)	(4)	(5)
	Grade	Absent Days	Gross Pay	Overt. Hours	Tenure
Surveyed Men	0.148***	-0.074***	0.074***	0.068***	-0.086***
	(0.017)	(0.023)	(0.018)	(0.018)	(0.018)
Observations	188,736	133,124	188,548	164,711	165,245
Surveyed Women	0.078***	-0.070***	0.125***	0.110***	-0.049***
	(0.010)	(0.011)	(0.010)	(0.011)	(0.011)
Observations	550,470	392,871	550,027	486,795	481,608

Notes: Table shows regressions of indicator variables that worker was surveyed among all sewing workers of grade 3-6 in HR data of factory, on different worker observables. Data standardized at Factory-Gender-Month Level. Thus coefficients resemble deviations of surveyed sample in standard deviations.

Table 2: General Gender Wage- and Grade Gaps

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Sample	Grades 2-7			Grades 3-7			
Dep. Variable	Log Wage	Log Wage	Grade	Log Wage	Log Wage	Log Wage	Grade
Female Worker	-0.206***	-0.012***	-1.184***	-0.077***	-0.080***	-0.015***	-0.791***
	(0.017)	(0.003)	(0.078)	(0.007)	(0.007)	(0.005)	(0.062)
Observations	46,134	46,134	46,201	44,267	81,508	81,508	81,581
Nbr. Factories	36	36	36	36	70	70	70
Factory FE	YES	YES	YES	YES	YES	YES	YES
Grade FE	NO	YES	-	NO	NO	YES	-

Notes: Table shows the results from regressing log monthly basic wage (excluding overtime and bonus payments, and deductions from missing days), and inverted grade (7 – grade) on an indicator variable that worker is female, and on factory fixed effects. Columns 2 and 6 also control for worker grade fixed effects. Unit observations are all worker in first month of data available from factory. Columns 1-3 based on sample of all workers in factory's sewing sections of grades 2 to 7, from the 36 factories from which consistent data on supervisor workers of grade 2 is available in HR data. Columns 4-7 based on sample of all workers in factory's sewing sections of grade 3-7, from all 70 factories in dataset. Standard errors clustered at Factory level. Asterisks indicate significance at the .10 (*), .05 (**), and .01 (***) levels.

Table 3: Grade Gaps with Worker Skill Data

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep. Variables	Grade				Avg. Efficiency	Nbr. Processes	Highest Complex.	Phys. Strength
Female Worker	-0.344***		-0.235*	-0.245**	-0.020	-0.071	-0.403***	-0.339***
	(0.119)		(0.123)	(0.117)	(0.145)	(0.103)	(0.140)	(0.118)
Nbr. Processes		0.209***	0.195***	0.190***				
		(0.000)	(0.000)	(0.000)				
Highest Complexity		0.157***	0.135***	0.134***				
		(0.000)	(0.000)	(0.000)				
Physical Strength		0.111	0.195***	0.219***				
		(0.089)	(0.074)	(0.000)				
Average Efficiency				0.104**				
				(0.052)				
Observations	3,583	3,583	3,582	3,424	3,425	3,583	3,583	3,583
Nbr Factories	20	20	20	16	16	20	20	20
Factory FE	YES	YES	YES	YES	YES	YES	YES	YES
Machine FE	NO	NO	YES	YES	-	-	-	-
Replicate with Log Wage as Dep. Variable								
Female Worker	-0.035***		-0.026***	-0.028***				
	(0.012)		(0.009)	(0.010)				
Nbr. Processes		0.016***	0.016***	0.015***				
		(0.000)	(0.000)	(0.000)				
Highest Complexity		0.012***	0.011***	0.011***				
		(0.000)	(0.000)	(0.000)				
Physical Strength		0.011**	0.016***	0.017***				
		(0.005)	(0.000)	(0.000)				
Average Efficiency				0.009***				
				(0.003)				
Observations	3,577	3,577	3,576	3,418				

Notes: Columns 1-4 in the upper part of the table show results from regressing reversed grade (7-grade) on an indicator variable for a female worker, and four measures of worker's skill or productivity: "Nbr. Processes" is the number of sewing processes on which the worker is certified and tested by the factory; "Highest Complexity" is the complexity of that skill among which worker is certified that has the highest complexity on a seven-point scale; "Physical Strength" is indicator that the worker is certified on a skill classified as requiring physical strength; and "Avg. Efficiency" is the worker's average efficiency on those processes on which she/he is certified. Columns 5-8 test whether there is a gender gap on any of these four skill measures. The bottom part of the table replicates columns 1-4 from the upper part, using log wage as outcome. Level of observation is worker. All regressions control for Factory fixed effects. Wild-cluster bootstrap standard errors clustered at Factory level. Asterisks indicate significance at.10 (*), .05 (**), and .01 (***) level.

Table 4: Grade Gaps with Survey Controls

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Variable:	**(Reverse) Grade:**					
	Grades 6-3	Grades 6-3	Grades 6-3	Grades 7-3	Grades 7-3	Grades 7-3
					Fact.Ten.	Adj.Sect.Ten.
Female Worker	-0.375***	-0.251***	-0.236***	-0.803***	-0.388***	-0.427***
	(0.054)	(0.045)	(0.047)	(0.106)	(0.068)	(0.072)
Years in Sector		0.197***	0.200***		0.455***	0.427***
		(0.018)	(0.019)		(0.029)	(0.028)
Years in Sector^2		-0.006***	-0.006***		-0.016***	-0.015***
		(0.001)	(0.001)		(0.002)	(0.002)
Years in Factory			0.059**			
			(0.025)			
Years in Factory^2			-0.045***			
			(0.014)			
Age			-0.001			
			(0.005)			
Years Schooling			0.019***			
			(0.006)			
Observations	2,152	2,152	2,152	2,864	2,864	2,864
Nbr. Factories	70	70	70	70	70	70
Factory-FE	YES	YES	YES	YES	YES	YES

Notes: Columns 1-3 show results from regressing reversed grade (7-grade) of workers on an indicator variable for a female worker, and worker observables from surveys of a representative sample of sewing workers of grades 6-3. "Years in Sector" is years worked in any garment factory, "Years in Factory" is years worked in current factory, while "Years per Factory" is years in the sector divided by number of factories worked in. Columns 4-5 add sampled observations from grade 7 workers from HR data, such that share of sampled grade 7 workers in combined dataset is equal to their share in the overall HR data, and uses factory tenure for the grade 7 workers as sector tenure. Colum 6 shows same specification as column 5, but for grade 7 workers instead uses simulated sector tenure values, with the simulation approach explained in Appendix B.2. Level of observation is on the worker level. All regressions control for factory fixed effects. Standard errors clustered at Factory level. Asterisks indicate significance at the .10 (*), .05 (**), and .01 (***) levels.

Table 5: Internal and External Promotion Rates (Monthly)

	(1)	(2)	(3)	(4)	(5)
	HR Payroll data statistics: Internal Promotion Rates:			**Fitted model parameters:** External Promotion Rates:	
	Men	Diff. Female Workers	w/ Fact.-Month FE	Men	Diff. Female
Grade 3	0.0016	-0.0023**	-0.0006**	-	-
Grade 4	0.0054	-0.0005	-0.0028***	-0.0001	-0.0007
Grade 5	0.0128	-0.0027	-0.0057***	0.0167	-0.0045
Grade 6	0.0200	-0.0046	-0.0070***	0.0225	-0.0046
Grade 7	0.0225	-0.0082***	-0.0096***	0.0377	-0.0193
Overall	0.0090	+0.010	-0.0007		

Notes: The first three columns show internal promotion rates estimated directly from the HR data for men, and for the difference between men and women. We show both the overall exit rate and the rate for each grade. The second column shows the raw gender difference, the third shows coefficients from regressions of promotion rates on female indicator variable, controlling for factory-month fixed effects. The fourth and fifth column show fitted parameter values from the model (described in more detail in Appendix D), for external promotion rates on grades 4-7, and the difference of these numbers between female and male workers. The model does not provide estimates for (external) promotion rates out of grade 3 to grade 2, as we lack necessary data from grade 2 workers. The statistical significance of differences of internal promotion rates based on standard errors clustered at factory level.

Table 6: Selection out of Sector

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Variable	Education	Education	Married	Married	Child	Child
Grade 3, Male	0.303	0.305	-0.015	-0.020	-0.048	-0.154
	(0.233)	(0.235)	(0.036)	(0.036)	(0.056)	(0.119)
Grade 3, Male x Sur. Year		-0.010		-0.023		0.121
		(0.218)		(0.034)		(0.129)
Grade 3, Female	-0.255	-0.263	-0.050*	-0.054*	-0.037	0.025
	(0.182)	(0.183)	(0.028)	(0.028)	(0.037)	(0.074)
Grade 3, Female x Sur. Year		0.122		-0.023		-0.081
		(0.175)		(0.027)		(0.082)
Survey Year x Female		-0.005		-0.032		-0.020
		(0.143)		(0.022)		(0.078)
Survey Year		-0.032		0.061***		0.047
		(0.124)		(0.019)		(0.071)
Year of Birth	0.127	0.131	-0.213***	-0.219***	-0.206***	-0.208***
	(0.119)	(0.119)	(0.018)	(0.018)	(0.029)	(0.029)
Year of Birth x Female	0.456***	0.457***	0.111***	0.113***	0.035	0.036
	(0.134)	(0.135)	(0.021)	(0.021)	(0.032)	(0.032)
Female	-0.399***	-0.398***	0.136***	0.131***	0.085**	0.102
	(0.141)	(0.142)	(0.022)	(0.022)	(0.034)	(0.069)
Constant (Male Avg.)	6.112***	6.109***	0.685***	0.692***	0.689***	0.652***
	(0.123)	(0.123)	(0.019)	(0.019)	(0.031)	(0.062)
Observations	2,156	2,156	2,170	2,170	1,037	1,037
Nbr. Factories	70	70	70	70	48	48

Notes: The table shows regressions of years of education, marriage status and parental status on female worker indicator variable, and an indicator variable for the worker being on the hightest (non-supervisory) grade 3, with controls for year of birth and survey year. Each of the independent variables is interacted to allow the effects to differ for female and male workers. Columns 2, 4 and 6 interact the grade 3 variables with survey year variables. Asterisks indicate significance at .10 (*), .05 (**), and .01 (***) level.

Table 7: Grade, Promotion Gap, Marriage, and Children

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Dep. Variable:	***Part A*: Married**						***Part B*: Children**					
Dep. Variable:	**(Reverse) Grade:**		**Internal Promotions:**		**Years per Factory:**		**(Reverse) Grade:**		**Internal Promotions:**		**Years per Factory:**	
Female Worker	-0.236***	-0.215**	-0.376***	-0.197	0.694***	0.424***	-0.196***	-0.291**	-0.616**	-0.720*	0.909***	0.874**
	(0.047)	(0.084)	(0.137)	(0.193)	(0.145)	(0.135)	(0.067)	(0.139)	(0.262)	(0.389)	(0.328)	(0.363)
Married x Male		0.011		0.235		-0.412**						
		(0.076)		(0.234)		(0.190)						
Married x Female		-0.016		-0.020		-0.022						
		(0.049)		(0.176)		(0.147)						
Married							0.067	0.062	-0.152	-0.124	-0.278	-0.148
							(0.133)	(0.137)	(0.307)	(0.300)	(0.713)	(0.704)
Child								-0.640***		-0.661		-0.166
								(0.181)		(0.656)		(0.383)
Child under 5								0.468***		0.036		0.055
								(0.165)		(0.529)		(0.378)
Child x Female								0.421**		0.189		0.391
								(0.186)		(0.552)		(0.406)
Child under 5 x Fem.								-0.299		0.036		-0.765*
								(0.179)		(0.614)		(0.447)
Observations	2,152	2,152	2,128	2,128	2,152	2,152	1,030	1,030	1,011	1,011	1,030	1,030
Nbr. Factories	70	70	70	70	70	70	48	48	48	48	48	48
Worker Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Factory-FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Grade-FE			Yes	Yes	Yes	Yes			Yes	Yes	Yes	Yes

Notes: The dependent variable in columns 1-2 and 7-8 is the current reverse grade (7-grade), in columns 3-4 and 9-10 the internal promotion rates of workers (the difference between the current and next month's grade of a worker, multiplied by 12), and in columns 5-6 and 11-12 the ratio of years worked in the sector and the number of factories in which a worker has worked ("Years per Factory"). These are regressed on an indicator variable for a female worker, and worker observables from survey of workers of grades 6-3. "Child" is indicator variable for having any child, and "child under 5" for having any child younger than 5 years old. All regressions control for factory FE. "Worker Controls" include worker's Years in Sector and its square, Age, and Years of schooling, and in columns 1-4 and 7-10 Years in current factory and average Years per factory over career in sector (dropped in Cols. 5,6, 11,12 due to collinearity with outcome). Standard errors clustered at factory level. Asterisks indicate significance at the .10 (*), .05 (**), and .01 (***) levels.

Table 8: Number of Factories in Post-code

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Wage	Wage	Grade	Promotions	Exit	External Arrivals
Nbr Near Factories x Fem.	-0.002	0.004	-0.070	0.452	1.327*	0.029***
	(0.005)	(0.004)	(0.055)	(0.470)	(0.690)	(0.009)
Age Factory x Fem.	-0.001	0.003	-0.049	0.369	0.270	-0.017
	(0.010)	(0.007)	(0.061)	(0.546)	(1.358)	(0.017)
Size Factory x Fem.	-0.004	-0.002	-0.034	0.016	0.982	0.014
	(0.008)	(0.006)	(0.067)	(0.692)	(0.909)	(0.012)
Female Worker	-0.070***	-0.030**	-0.537***	-5.856***	-7.358***	-0.163***
	(0.018)	(0.013)	(0.193)	(2.129)	(2.394)	(0.033)
Observations	68,836	68,836	68,903	68,406	68,903	51,396
Nbr Factories	59	59	59	59	59	59
Factory-FE	YES	YES	YES	YES	YES	YES
Grade FE		YES	-	YES	YES	YES

Notes: The table shows results from regressing outcomes on the individual level on an indicator variable for female workers and an interaction of this variable with the log hyperbolic sine transformation of the number of other factories within 2 kilometres of the factory of the worker. Regressions control for factory fixed effects, and interactions of some factory characteristics (factory age, number of workers) with the female worker indicator. Standard Errors clustered at factory level. Asterisks indicate significance at the .10 (*), .05 (**), and .01 (***) levels.

Appendix A: The Wage Gap, Worker Absenteeism and Overtime:

For different sub-sets of the 70 factories in our main sample, the HR data included detailed data on days absent, the attendance bonus earned, and the number of overtime hours worked, on the worker-month level. Attendance bonus is a monthly payment to workers who missed no more than a specified number of workdays in a month, usually either one or two days. These attendance bonuses are typically around 5 percent of the overall monthly pay of a worker.

In Table A.1, we regress these three variables on an indicator variable for female worker, controlling for factory, month, and grade fixed effects, on the worker-month level. We cluster standard errors at the factory level. We can see that women miss on average 0.34 fewer days per month (of a mean of 0.67 absent days and a median of 0), earn on average 11.5 BDT higher attendance bonus per month (of a mean absent bonus of 336, or median of 400 (~US$ 4.8)), and work 2.36 overtime hours less per month (of a mean of 50.6 hours and a median of 50 hours). All differences are statistically significant.

Table A.1: Gender Differences in Absenteeism and Overtime:

	(1)	(2)	(3)
	Absent Days	Attendance Bonus	Overtime Hours
Female Worker	-0.0341**	11.5474***	-2.3638***
	(0.0154)	(2.2486)	(0.3324)
Observations	522,094	656,612	346,563
Nbr. Factories	57	63	42
Factory & Month FE	YES	YES	YES
Grade FE	YES	YES	YES

Notes: Column 1 shows regression of number of days worker was absent in a given month on dummy variable for female worker, while in column 2 the outcome variable is earned attendance bonuses (in BDT) by workers in a given month, and in column 3 the total overtime hours worked by the worker in a given month. Level of observation is worker-month level. All regressions control for factory, month, and grade fixed effects. Standard errors clustered on the factory level: Asterisks indicate significance at the .10 (*), .05 (**), and .01 (***) levels

In Table A.2, columns 1, 3, and 5, we proceed to replicate columns 5 to 7 of Table 2 on the baseline grade gaps in the sample of those 27 factories for which we have both the absenteeism and overtime data in the HR data. Both the basic wage gap and the grade gap is about 20 percent smaller in this subsample of factories. The inclusion of these three controls does not change the estimates of the wage gap (column 2), the grade gap (column 4), or the wage gap conditional on grade fixed effects (column 6) in any significant way.

Table A.2: Absenteeism, Overtime, and the Gender Wage Gap

	(1)	(2)	(3)	(4)	(5)	(6)
	Log Wage		(Reverse) Grade		Log Wage	
Female Worker	-0.064***	-0.064***	-0.621***	-0.628***	-0.013*	-0.014**
	(0.005)	(0.006)	(0.080)	(0.058)	(0.007)	(0.006)
Absent Days		20.069***		200.144***		2.838*
		(4.214)		(43.651)		(1.617)
Attend. Bonus		0.413***		3.811***		0.072**
		(0.089)		(0.902)		(0.031)
Overtime Hours		-0.732***		-10.389***		0.251
		(0.231)		(3.100)		(0.175)
Observations	28,107	28,107	28,165	28,165	28,107	28,107
Nbr. Factories	27	27	27	27	27	27
Factory-FE	YES	YES	YES	YES	YES	YES
Grade FE	NO	NO	NO	NO	YES	YES

Notes: Regression of log base wage (Columns 1-2, and 5-6) and reverse worker Grade (Columns 3-4) on an indicator variable for female worker in those 27 factories in the sample in which consistent attendance and overtime data is available in HR data. Columns 2 and 4, respectively, control for number of days in month worker is absent, monthly levels of attendance bonuses earned, and overtime hours of worker, averaged across month within workers. Columns 5-6 control for grade fixed effects. Level of observation is worker from the first month of data available from factory. All regressions control for factory fixed effects. Standard errors clustered at factory level. Asterisks indicate significance at the .10 (*), .05 (**), and .01 (***) levels.

Appendix B.1: Gender Gaps among workers within narrow skill-brackets

The fact that a grade gap remains after controlling for skills is not just due to the restrictive linear form in which we control for them. Table B.1 shows that the gaps are visible and significant even within sub-samples of workers of the same narrowly defined skill levels. For example, the grade gap remains economically and statistically significant among the top 25 or even top 10 percent of workers that are skilled on the highest number of processes. The same holds for the roughly 50 percent of workers who are skilled on a process of the highest complexity level, the roughly 50 percent of workers that can do a process requiring physical strengths, or among the top 25 or 10 percent when it comes to efficiency on the processes they can do. It is also visible for example among those that can do a process of the highest complexity *and* are among the top 25 percent workers that can do the largest number of processes (2nd Panel, Column 4).

Table B.1: Grade Gap by Skill Levels

	(1)	(2)	(3)	(4)
VARIABLES	Grade	Grade	Grade	Grade
Nbr. Processes	< 50th Pctl	> 50th Pctl	> 75th Pctl	> 90th Pctl
Female Worker	-0.351***	-0.326***	-0.504**	-0.506***
	(0.122)	(0.113)	(0.200)	(0.175)
Observations	1,986	1,597	935	418
Nbr Factories	20	20	15	7
Highest Complexity	< 50th Pctl	> 50th Pctl		> 50th Pctl & Nbr.Proc. > 75th
Female Worker	-0.313**	-0.290***		-0.462*
	(0.142)	(0.100)		(0.250)
Observations	1,878	1,705		755
Nbr Factories	20	20		14
Physical Strength	No	Yes		
Female Worker	-0.123	-0.366***		
	(0.180)	(0.127)		
Observations	1,847	1,736		
Nbr Factories	20	19		
Efficiency	< 50th Pctl	> 50th Pctl	> 75th Pctl	> 90th Pctl
Female Worker	-0.309***	-0.395***	-0.400***	-0.468*
	(0.107)	(0.137)	(0.139)	(0.240)
Observations	1,664	1,761	826	330
Nbr Factories	16	16	16	13
Factory FE	YES	YES	YES	YES

Notes: Table shows regressions of Grade on indicator variable for female worker and factory fixed effects, on varying subsamples, according to skill levels on the four different skill dimensions. Grade is reversed, so that negative coefficient implies lower grade. Wild-cluster bootstrap standard errors clustered at Factory level. Asterisks indicate significance at.10 (*), .05 (**), and .01 (***) level.

Appendix B.2: Simulating Sector Tenure for Grade 7 workers from HR data

To include grade 7 workers in the grade gap regressions controlling for sector tenure, we sample, as mentioned in the main text, 712 grade 7 workers from the HR data (80 male and 632 female). For these workers, we know factory tenure from the HR data, but not sector tenure. To account for the possibility that these workers may have longer sector tenure than factory tenure, we use information from a separate sample of 190 surveyed grade 7 workers from 4 factories outside the main sample of 70 factories used in this paper, to simulate sector tenure based on the joint distribution of factory and sector tenure among the 190 surveyed grade 7 workers. The left panel in Figure B.2 shows this empirical joint distribution among these 190 workers. To replicate this distribution among the 712 sampled workers, we use the following (non-parametric) approach. We first separate the 190 surveyed workers into ten deciles according to their reported tenure at their current factory. Then, for each decile, we take the share of workers for whom sector and current factory tenure differs, and then calculate the ratio between sector and factory tenure for those workers for whom the two variables differ in each decile.

We then separate the 712 sampled workers from the HR data into ten bins, based on the deciles of factory tenure from the 190 surveyed workers. We randomly sample in each of the bins a share of workers equal to the decile specific share of workers for whom sector and factory tenure differs in the survey data. For these selected workers from each bin, we then obtain a simulated sector tenure by multiplying their factory tenure values with the decile-specific ratio of sector and current factory tenure from the survey data. The right hand panel of Figure B.2 plots the joint distribution of the empirical factory tenure and one iteration of simulated sector tenure, among the 712 workers sampled from the HR data, indicating that it is qualitatively similar to the joint distribution from the 190 surveyed workers shown in the left hand panel of Figure B.2.

Replicating column 5, Table 4, with repeated simulated sector tenure for the grade 7 workers in column 6 of the same table shows that simulated sector tenure consistently explains around 10 percent less of the grade gap than if we just use worker's factory tenure as proxy for sector tenure. This is little surprising, because if we use a measure for an independent variable with less variation (factory tenure instead of simulated sector tenure) for a set of observations who all have the same value for the dependent variable (grade equal seven), then this measure will explain more of the variation in the dependent variable. We therefore consider using factory tenure as proxy for sector tenure for grade 7 workers as providing an upper bound for the share of the grade gap explained by differential sector tenure of men and women.

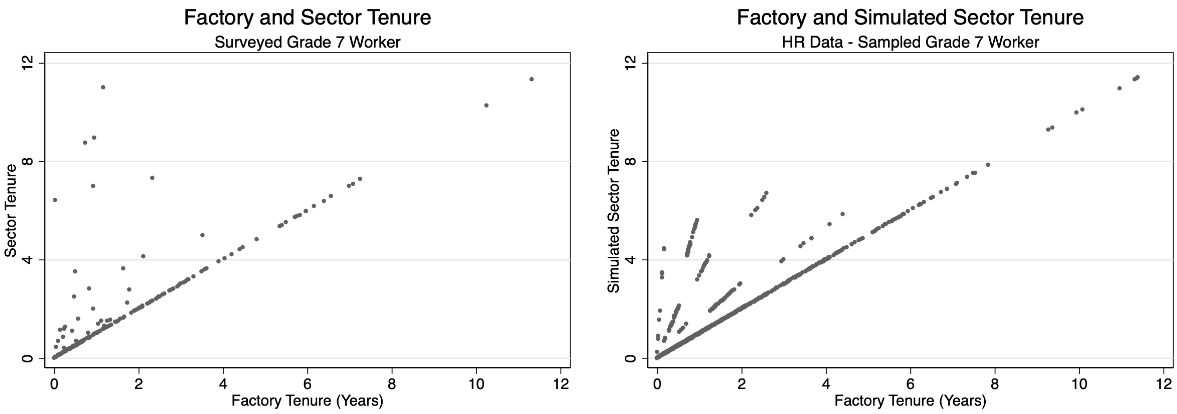

Figure B.2: Simulation of Sector Tenure for Grade 7 Workers. Left panel plots empirical joint distribution of sector tenure and tenure at current factory among a sample of 190 grade 7 workers surveyed at three factories outside of main sample of factories in this paper. Right panel shows joint distribution of simulated sector tenure and empirical factory tenure in sample of 712 grade 7 workers, randomly sampled from HR data, using the simulation algorithm laid out in Appendix B.2.

Appendix C.1: Controlling Grade Gap for Survey and Skill Observables

Table C.1 below adds both skill and survey controls into the grade gap regression, in the sample of workers for which both types of data is available. These are 154 workers from 9 factories, 131 of which are female. Given the small sample and the large number of possible controls, overfitting of the regression may be a concern. This would overstate the share of the grade gap explained by the two sets of control variable. Therefore, we use PDS Lasso (Belloni et al. 2016), to discipline the selection of control variables from both datasets in the regression. The pool of variables from which the Lasso can chose are all variables included in Tables 3 (skill data, including machine fixed effects), and 7 (survey data), and the squares of the continuous variables among them. All variables chosen by the PDS Lasso shown are in Table C.1.

Table C.1 Controlling Grade Gap for Survey Observables and Skills

	(1)	(2)	(3)	(4)	(5)
Dep. Variable	Grade	Grade	Grade	Years/Sector	Grade
Female Worker	-0.422	-0.313	-0.248		-0.198
	(0.287)	(0.240)	(0.257)		(0.255)
Years in Sector		0.065**			0.138***
		(0.027)			(0.038)
Nbr. Factories		0.092			
		(0.077)			
Years in Sector^2					-0.005**
					(0.002)
Nbr. Processes			0.061**	0.026	0.058**
			(0.029)	(0.053)	(0.024)
Physical Strength			0.769***		0.651***
			(0.170)		(0.202)
Observations	150	150	150	150	150
Nbr. Factories	9	9	9	9	9
Factory-FE	YES	YES	YES	YES	YES
Machine FE	-	-	Yes (4 sel.)	Yes (3 sel.)	Yes (5 sel.)

Notes: Table shows results from regressing reversed grade (7-grade) of workers on an indicator variable for a female worker, worker observables from survey of representative set of sewing workers of grades 6-3, and on four worker skill measures, from workers from those 9 factories at which survey and skill data overlap. Column 2 introduces survey controls only, column 3 skill controls only, and column 5 both. Controls from both sets of variables chosen by PDS Lasso, to prevent overfitting due to large number of control variables and small sample. Level of observation is worker level. All regressions control for factory fixed effects. Standard errors clustered at factory level. Asterisks indicate significance at the .10 (*), .05 (**), and .01 (***) levels.

Column 1 shows the basic grade gap in the regression when only controlling for factory fixed effects. Column 2 adds survey-based controls only, as selected by the PDS Lasso, and Column 3 skill controls only. While survey-based controls explain about 25 percent of the gap in this sample alone, the two selected skill controls explain about 40 percent. Skills are positively correlated with years in sector, at least the number of processes a worker can do, though the correlation is not statistically significant (column 4). Finally, jointly, skills and survey observable explain about 60 percent of the grade gap in this sample (column 5). Adding the skills controls to the controls for sector tenure increase the share of the grade gap that is explained by around 13 percent ((0.248 – 0.198)/ 0.422).

Appendix C.2: Worker Skills and Internal Promotions

We check whether the skills data explain not just parts of the gaps in grade levels (Table 3), but also of internal promotions. Table C.2, column 1, shows that the gender gap in internal promotions in the sample of 20 factories with skills data is not statistically significant, at least with standard errors clustered on the 20 factories using wild-cluster bootstrap. However, its point estimate of 0.6 percent per month is larger than the highly statistically significant internal promotion gap of 0.50 percent in the data from all 70 factories (when controlling for grade fixed effects). In column 2, we regress promotions on the three skill measures we have from all 20 factories and grade fixed effects. Task complexity is the only skill that predicts promotions. Adding the female worker dummy (column 3) shows that, relative to column 1, controlling for skills reduces by half the gender gap in promotion rates, with the gap remaining statistically insignificant. Adding average efficiency on the processes a worker can do as an additional control, in the 16 factories from which this data is available, does not change this result (column 4). Though the results are noisy, they suggest that skill differences between women and men explain roughly half of the internal promotion gap, similar to the share of the overall grade gap they explain.

Table C.2: Internal Promotions with Worker Skill Data

	(1)	(2)	(3)	(4)
	Int. Promotions			
Female Worker	-0.600		-0.283	-0.290
	(0.435)		(0.299)	(0.325)
Nbr. Processes		-0.097	-0.169	-0.160
		(0.335)	(0.167)	(0.148)
Highest Complexity		0.573*	0.538*	0.554*
		(0.340)	(0.295)	(0.312)
Physical Strength		0.123	-0.101	-0.048
		(0.245)	(0.326)	(0.265)
Avg. Efficiency				0.320
				(0.362)
Observations	3,527	3,527	3,526	3,370
Nbr. Factories	20	20	20	16
Factory FE	YES	YES	YES	YES
Grade FE	YES	YES	YES	YES
Machine FE	NO	NO	YES	YES

Notes: Table shows results from regressing internal promotions in the sample of factories with skill data available, on the four main skill variables. "Nbr. Processes" is the number of those processes on which the worker is officially trained on. "Highest Complexity" is the complexity of that skill among which worker is trained on that has the highest complexity on a seven-grade scale. "Physical Strength" is indicator variable that worker is trained on a skill classified as requiring physical strength. "Avg. Efficiency" is the workers average efficiency in those processes on which the worker is officially trained on. Level of observation are individual workers. All regressions control for Factory and grade fixed effects. Wild cluster bootstrap standard errors clustered at Factory level. Asterisks indicate significance at.10 (*), .05 (**), and .01 (***) level.

Appendix D: Backing out External Promotion Rates

We already argued in section 4 that simple back-of-the-envelope calculations to obtain the share of external promotions among all promotions, based on the HR and survey data available to us, may be severely biased if we do not account for the possibility that sector exit and promotion rates are correlated. Also a second naïve approach of multiplying the average gender gap in internal promotion rates with worker's average career length, and comparing it against the grade gap from the HR data is not valid, because it ignores the fact that promotions are more frequent early on in the careers of workers in the sector. There may be no gender differences in *average* promotion rates over worker careers but still a grade gap in cross-sectional HR data, if the promotions of men occur earlier in their careers.[28] Thus, we cannot claim that any grade gap not explained by differences in *average* internal promotion rates must be due to different external promotion rates.

We show here that, under the assumption that all workers started their career on grade 7, plus some additional modelling assumptions, we can back out external promotion and sector exit rates for each grade using only the sector tenure data for different grades, and the relative number of workers on each grade. To do this, we write down a model that includes promotions and sector exit rates as parameters, and that generates at least as many moment predictions that can be matched to aggregate moments in our data as it has parameters. The model explicitly allows workers that will reach different highest grades during their career to be promoted, and to exit the sector, at different rates. Furthermore, by allowing the rate at which workers exit the sector once on their final grade to differ from the rate at which they got promoted before reaching that grade, we can model arbitrary levels of concavity of promotions over worker careers, independently for different worker types.

A.D.1 The Model

We construct an exactly identified model with 13 parameters generating predictions for 13 different moments we also observe in the data. The model can be fitted to the data separately for male and female workers. For that purpose, all parameters can be indexed by gender $\{m, f\}$. To simplify notation, however, we suppress the gender subscripts here.

[28] To see this more formally, consider a stylised model, in which there are only two grades, 0 and 1. Each worker starts her/his career on grade 0 and exits the sector after T periods of time. Assume that men get promoted from grade 0 to grade 1 after $t_m < T$ periods, while women after $t_m < t_f < T$. If new cohorts of male and female workers of constant size enter the sector each period in time, the average grade of all workers of gender g at any point in time would be $(T - t_g)/T$. Thus, if t_g differs for men and women, a grade gap would be estimated between these workers. But average promotion rates would not differ between male and female workers.

Let each worker's type k be indexed by the highest grade i, 7 through 3, that the worker reaches during her career in the sector. A share s_k of workers of type k on all lower grades $i > k$ than their highest ultimate grade have an (external or internal) promotion each time period. Recall that grade 7 is the entry level grade and grade 3 is the highest operator grade. So a promotion is a move from a grade i to a grade $i - 1$. Having reached their highest grade $i = k$, a share p_k of workers of type k exit the sector each time period. The assumption that workers of type k move at a constant rate through all grades on the way to the highest grade $i = k$ is strong, but central for us to be able to identify unobserved overall promotion frequencies from our data.[29]

We use these assumptions to characterize the steady state, in which the number of workers on each grade is constant over time. Assume a cohort of M_k workers of type k enter the sector each period of time at the lowest grade 7. The number of workers of each type on each grade will remain constant only if the number exiting the grade is the same as the number entering. Let W_i^k be the steady state number of workers of type k on grade i. For $W_i^k, i < k,$ to be constant, we then need $M_k = W_i^k s_k$, or $W_i^k = M_k / s_k$. Similarly, for the number of workers of type k on their final grades $i = k$ to be constant we need $W_{i=k}^k = M_k / p_k$, This implies that, at any time, the number of workers W_i on any grade i is:

$$W_i = \sum_{k=3}^{i} W_i^k = \left[\sum_{k=3}^{i-1} \frac{M_k}{s_k}\right] + \frac{M_i}{p_i} \tag{D.1}$$

[29] One way we can test for the sensitivity of our results to this assumption is to look specifically at promotions from grade 4 to grade 3. Anecdotal evidence from the sector shows that grades 7-5 are "learning" grades, through which workers pass to reach the "final" grades 4 and 3. This implies that the promotion process from grade 4 to 3 differs structurally from those to the previous grades, and may occur at a lower rate for type 3 workers than the previous promotions up to grade 4. We can show that assuming for type three workers that their promotion rate from grade 4 to 3 is λs_3, $1/4 < \lambda < 1$, does not qualitatively affect the results derived from the model. Thus, they are insensitive to assuming that the last promotion for type 3 workers occurs at a rate that is up to four times slower than for the prior promotions.

While we do observe the number of workers on each grade 3 through 7 in our HR data, it will be easier to target the ratio of workers on adjacent grades i and $i+1$, or $r_{i,i+1} = W_i/W_{i+1}$. For this purpose, it will be useful to choose one worker type k as the numéraire and express the size of entry cohorts of other worker types M_k relative to the size of that cohort. Ultimately it does not matter which worker type's cohort size we choose as numéraire, but it turns out that the parameter estimates are more stable when we choose $M_3 = 1$.

We can model the average and variance of sector tenure of all workers on grade i in the steady state by making additional assumptions on how the share s_k of workers who are promoted are selected among the workers on the grade, and similarly the share p_k among workers who exit from their final grade. We consider three different assumptions that could be made on how this selection process works. The first assumes that all workers of type k spend the same time on a given grade i. That is, workers exit non-final grades after exactly $1/s_k$, and final grades after $1/p_k$ time periods. In a *cross-section* of workers of a given type on a given grade in the steady state, the time these workers will have already spent on the grade would then follow a uniform distribution, with mean $1/2\mathrm{s_k}$ on non-final grades, and mean $1/2p_k$ on final grades. We thus name this modeling option the "uniform" modeling assumption. A second possible assumption is that the workers promoted or exiting the sector are independently selected with respect to the time they already spent on the grade. This memory-less selection process will lead to an exponential distribution of time already spent on the grade among a cross section of workers of a certain type on a certain grade, with means $1/s_k$ or $1/p_k$. We refer to it as the "exponential" modeling assumption. A third option assumes that the workers promoted or exiting at each period of time are selected such that the distribution of the time they have spent on the grade, *at the time they exit the grade*, is uniform. This implies that the exit or promotion probability strictly increases in the time workers have spent on the grade. This assumption is thereby an intermediary case between the other two cases, in which the exit or promotion probability either jumps from zero to one at one point in time (uniform assumption), or remains constant over time (exponential assumption). In a cross section of workers of a certain type on a certain grade, the third modeling option leads to a linear, "triangular" distribution of time spent on the grade, with means $2/3\mathrm{s_k}$ or $2/3p_k$.

How do we obtain the average and variance of sector tenure for all workers on a given grade in the steady state (not just of those of a certain type on a grade)? The average tenure T_i of workers on grade i is the weighted average of the average tenures of workers of different types k on grade i, where the weights are the share of workers of each type on the grade in the steady state. For example, with uniform selection, this is:

$$T_i = \frac{\sum_{k=3}^{i-1} \frac{M_k}{s_k}\left((7-i+0.5)\frac{1}{s_k}\right) + \frac{M_i}{p_i}\left((7-i)\frac{1}{s_i} + \frac{0.5}{p_i}\right)}{\sum_{k=3}^{i-1} \frac{M_k}{s_k} + \frac{M_i}{p_i}} \tag{D.2}$$

M_i/p_i is the relative number of worker of type $k = i$ on grade i at any point in time, and $(7-i)\frac{1}{s_i} + \frac{0.5}{p_i}$ is their average sector tenure. $(7-i)\frac{1}{s_i}$ is the time spent on average on all preceding grades on the way to grade i, and $0.5/p_i$ is the average time spent on the final grade i. Similarly, $(7-i+0.5)\frac{1}{s_k}$ is the time spent by workers of higher types $k < i$ so far in the sector, and M_k/s_k is their relative number on grade i.

The variance of sector tenure of all workers on a grade in the steady state, on the other hand, cannot be simply derived as the weighted average of the variances within the subsamples of workers of different types on the grade. Instead we need to use variance decomposition: the variance of a sample is the average of the variances within all subsamples, plus the average of the squared deviations of the sub-sample means from the overall mean of the sample, with both averages weighted by the sizes of the sub-samples. Equation D.3 shows this equation again under the uniform selection rule. An advantage of using the uniform modeling assumption for promotions on non-final grades is that the variance of sector tenure is equal to the variance of tenure on the current grade. This is because under the uniform assumption all workers of the same type spend the same amount of time on a given grade. So, for example, all workers of type 3 currently on grade 5 have spent the same amount of time on grades 7 and 6, those grades they already have passed through. Only on grade 5 would there be variation in time spent on that grade; some have just entered it while others are already close to the time they are promoted to the next grade. This property greatly simplifies the derivation of variance of sector tenure. Note that this property does not hold under the exponential or triangular assumptions as these both generate variation in the time that workers of a given type have spent on previous grades. Thus, under the uniform selection assumption (here for both promotion and sector exit) the variance of sector tenure of all workers on grade i in the steady state is:

$$V_i = \frac{\sum_{k=3}^{i-1} \frac{M_k}{s_k}\left(\frac{1}{12s_k^2} + \left(\frac{7-i+0.5}{s_k} - T_i\right)^2\right) + \frac{M_i}{p_i}\left(\frac{1}{12p_i^2} + \left(\frac{7-i}{s_i} + \frac{0.5}{p_i} - T_i\right)^2\right)}{\sum_{k=3}^{i-1} \frac{M_k}{s_k} + \frac{M_i}{p_i}} \quad \text{(D.3)}$$

Once we obtained estimates for the parameters s_k, p_k, and M_k, we can back out *overall* promotion rates P_i on any grade i by the following equation:

$$P_i = \frac{\sum_{k=3}^{i-1} \frac{M_k}{s_k} s_k}{W_i} = \frac{\sum_{k=3}^{i-1} M_k}{W_i} \quad \text{(D.4)}$$

where W_i is the steady state number of workers on a given grade, while the nominator sums over the promotion rates s_k of workers of higher types $k < i$, weighted by their relative numbers on grade i. We then obtain external promotion rates on each grade by subtracting the internal promotion rates estimated directly from the HR data for each grade, as shown in the first two columns of Table 5, from these backed out overall promotion rates.

A.D.2 Fitting the model to data

Now we have 13 unknown parameters: five sector exit probabilities for each worker type once they have reached their final grade, p_7-p_3; four relative sizes of each worker type (except one that is set to 1 as numéraire), M_7-M_4; and four promotion probabilities, s_6-s_3 (type 7 workers exit the sector before being promoted, so $s_7 = 0$).

We also have 13 observed moments in the data: four ratios of workers on adjacent grades, $r_{6,7}$- $r_{3,4}$; five average sector tenures of workers on grades 7-3, T_7-T_3; and four tenure variances V_6-V_3. The four size ratios $r_{i-1,i}$ come from factory HR records, and the sector tenure and variance data come from our surveys of workers on grades 6 through 3.[30]

We fit the model to the 13 data moments separately for females and males using each of the nine possible combinations of exit and promotion rules. For males, we find that only the combination of uniform promotion and exit probabilities yields parameter values in the feasible range – for example, positive values for the sector exit or promotion rates. The same combination yields feasible values for females, as does the combination of uniform promotion rates and "triangular" exit rates. The first two columns of Table D.1 below show the 13 fitted parameter values under the uniform selection assumption for promotion and sector exit for male and female worker. The promotion and sector exit rates are monthly rates. We note that

[30] A note about the average sector tenure for grade 7 workers is merited, which are not included in the workers surveys for the factories used in this paper. However, we do have survey information of grade 7 workers, including on sector tenure, from four other factories, as already discussed in Appendix B.2. This data includes surveys data from more than 230 female grade 7 workers, with an average reported sector tenure of 20.4 month, which is close to the average factory tenure of female grade 7 workers of 18 months in our administrative data. This indicates that not many (female) workers move between factories while not leaving grade 7. Only around one in eight do so. We thus set average sector tenure on grade 7 workers to 20.4 months. However, the data include only 9 surveys of male grade 7 workers, with average sector tenure of 23 months, which is subject to outlier observations, though. We thus rather assume in our baseline specification that as indicated by the data around one in eight male grade 7 workers worked at another factory before. Assuming for our baseline specification that they spent twice as much time in the previous factory than the current one (which would be predicted by the uniform exit assumption laid out above), average sector tenure of men would be 1.25 times their average factory tenure (12 months according to our HR admin data), or 15 months. However, our results are not sensitive to varying this sector tenure for male grade 7 workers between 12 and 23 months.

while all values are positive, that is in the feasible range, for women some of the values appear unrealistically high. Specifically, promotion rates of women of types 5-3 are very high, implying that these women are promoted to a higher grade almost every two months, and reach grade 3 after only little more than eight months in the sector. For men of type 3, we estimate a more realistic, but also much longer 25 months till they reach the highest grade. Given higher internal promotion rates for men, it strikes us as unlikely that overall promotion rates (internal + external) are much higher for women. Also, more than 50 percent of women entering the sector are estimated to never get promoted out of the entry level position of grade 7, which again strikes us as likely too high.

While for men the model only has feasible solutions if both promotions and sector exit are modeled under the "uniform" specification, for women the model also has a solution when modeling sector exit under the "triangular" model assumption, while maintaining the "uniform" assumption for selection to promotion. In this version of the model, average sector tenure of workers on grade *i* is given by:

$$T_i = \frac{\sum_{k=3}^{i-1} \frac{M_k}{s_k}\left((7-i+0.5)\frac{1}{s_k}\right) + \frac{M_i}{p_i}\left((7-i)\frac{1}{s_i} + \frac{2}{3\,p_i}\right)}{\sum_{k=3}^{i-1} \frac{M_k}{s_k} + \frac{M_i}{p_i}} \qquad (\mathrm{D.5})$$

Meanwhile, the variance is given by

$$V_i = \frac{\sum_{k=3}^{i-1} \frac{M_k}{s_k}\left(\frac{1}{12 s_k^2} + \left(\frac{7-i+0.5}{s_k} - T_i\right)^2\right) + \frac{M_i}{p_i}\left(\frac{2}{9 p_i^2} + \left(\frac{7-i}{s_i} + \frac{2}{3\,p_i} - T_i\right)^2\right)}{\sum_{k=3}^{i-1} \frac{M_k}{s_k} + \frac{M_i}{p_i}} \qquad (\mathrm{D.6})$$

The derivation of these formulas is relatively straightforward. At the heart of the triangular assumption is that among all workers who enter a grade together, a share e_i of the initial cohort exits the grade each period of time. Thus, after $1/e_i$ time periods, all workers of that cohort have left the grade. These assumptions leads to a distribution of time t spent on the current grade in a cross-section of workers described by the linear pdf $2e - 2e^2t, 0 \leq t \leq 1/e$, with an expectation of $1/3e$, and a variance of $1/18e^2$. Given that we can show $s_i, p_i = 2e_i$, we arrive at equations D.5 and D.6 above. Note that as we continue to use the uniform modeling assumption for promotions, the variance of sector tenure continues to be equal to the variance of tenure on the current grade.

Fitting the model using the “uniform" selection assumption for promotions and the “triangular” selection assumption for sector exit to the data for women, we obtain parameter values that are more realistic, as shown in column 3 of Table D.1. Women of type 3 are now promoted around every 6 months, reaching grade 3 in around 2 years after entering the sector, while now only a quarter of women who enter the sector never advance out of grade 7 (that is, are of type 7). Thus the parameter values we obtain for men under the first assumption listed above for both promotion and exit, and for women under this alternative set are our preferred specification.

Finally, Table D.2 shows the estimates we obtain for the external promotion rates based on equation D.4, after subtracting the internal promotion rates shown in Table 5. Note that the negative values for external promotion rates on grade 4, P_4, in Table D.2 are very small in absolute values (particularly for our preferred specifications), and are the results of subtracting the estimated internal promotion rates, as shown in Table 5, from the backed out overall promotion rates. We regard these small negative values to be not different from zero, once possible sampling error is taken into account in both the estimated internal promotion rates and the empirical moments fitted to the model. Note also that the relative results between men and women that we obtain for external promotion rates are the same for one important qualitative outcome regardless of the exit pattern assumed for women, uniform or triangular: external promotion rates are higher for men than for women on all grades but the last promotion from grade 4 to grade 3, and for this grade 3, external promotion rates are essentially zero for both men and women, as shown in Table D.2 below.[31]

Table D.1: Model Parameter Estimates

	(1)	(2)	(3)
	"Uniform" specification for promotion & sector exit		"Uniform" for promotion & "Triangular" for sector exit
	Men	Women	Women
Sector exit rates			
$p_{k=3}$	0.0072	0.0062	0.0101
$p_{k=4}$	0.0076	0.0078	0.0128
$p_{k=5}$	0.0087	0.0086	0.0149
$p_{k=6}$	0.0120	0.0107	0.0229
$p_{k=7}$	0.0202	0.0232	0.0225
Total promotion rates			
$s_{k=3}$	0.1537	0.4728	0.1721
$s_{k=4}$	0.1980	0.4583	0.1537
$s_{k=5}$	0.1091	0.5185	0.0875
$s_{k=6}$	0.0329	0.2255	0.0301
Relative cohort sizes (share)			
M_3	1.00 (25%)	1.00 (5%)	1.00 (8%)
M_4	1.37 (34%)	2.98 (15%)	2.93 (24%)
M_5	0.58 (14%)	2.40 (12%)	2.26 (18%)
M_6	0.61 (15%)	3.11 (16%)	3.10 (25%)
M_7	0.44 (11%)	10.12 (52%)	2.94 (24%)

[31] There are other valid reasons to assume that our preferred specification for sector exit for women is indeed more realistic for women than for men. For example, women may exit the sector for a more diverse set of reasons than men, such as childbirth, care for other family members, or other household-related reasons, which may arise in a less predictable manner over time. This should result in more variation in the time that women spend in the sector, a feature that is captured by the triangular specification for women as opposed to the uniform one, which assumed that workers (of a certain type) exit the sector after a constant time.

Table D.2: Estimated External Promotion Rates per Grade

	(1)	(2)	(3)
	"Uniform" specification for promotion & sector exit		"Uniform" for promotion & "Triangular" for sector exit
	Men	Women	Women
P_4	-0.0001	-0.0025	-0.0008
P_5	0.0167	0.0037	0.0122
P_6	0.0225	0.0056	0.0179
P_7	0.0377	0.0062	0.0184