



Workers in space: Evidence from Bangladesh

Julia Cajal-Grossi and Gabriel Kreindler

- Cities in developing countries house a large share of the world's population and are hubs of productivity and export activity. However, whether cities are a path to prosperity for workers in low-income countries remains an open empirical question.
- This brief uses data from the garments-for-exports sector in Bangladesh to study the spatial extent of urban labour markets in developing countries. From the perspective of firms, we study how demand shocks to firms spill over space, and from the workers' perspective, we study frictions that may limit their movements over space, in particular (i) commuting costs, and (ii) distance-dependent information frictions.
- We introduce several novel stylised facts, that show that home-work commutes, factory switches, and the information workers have on job opportunities are highly local and decay sharply with distance. We show that plants compete for workers locally, but workers would be willing to commute longer distances for good job opportunities.
- We propose a simple partial equilibrium model of worker location choice that allows us to explore the welfare implications of information and commuting frictions for workers. Using the model structure, our results suggest that approximately one third of the observed spatial concentration in labour markets follows from information frictions.
- Relative to a benchmark with perfect information, policies that reduce effective distance will have an effect on welfare different in magnitude and shape, due to the presence of distance-dependent information frictions.

Policy motivation

Cities in low-income countries exhibit contradictory forces. On one hand, they often achieve incredible levels of population density and are important hubs of industry and exporting activity, which generates economic opportunities for the poor. On the other hand, these cities often have inefficient transport infrastructure, are disorganised, and lack safety, forces that stifle the potential of connecting the poor with job opportunities. In the words of Lall et al. (2017), such cities are “*crowded yet disconnected*.” The role cities play in delivering prosperity to workers in low-income countries remains an open empirical question. In particular, we explore study three interrelated questions: (1) what the spatial extent of urban labour markets in is developing countries, (2) if any, what are the frictions that give these labour markets their shape, and (3) what are the welfare implications of those frictions.

This research offers empirical evidence answering these questions in the context of Bangladesh’s garment-exporting urban districts. The sector, which has been an engine of growth for the country over the last few decades, is a key driver of urban employment. It is also a setting where inequality in the labour market across locations is significant. On the firms’ side, we study how demand shocks to plants spill over space. On the workers’ side, we study frictions that may limit their movements over space: in particular (i) commuting costs, and (ii) distance-dependent information frictions about job availability. We contribute to the literature by documenting the extent of labour markets in a developing country context, and by unpacking the frictions that lie behind workers’ commuting patterns. Our main takeaway highlights that workers’ lack of information about job opportunities in different locations is a key determinant of the spatial concentration of labour markets in the garments-for-exports sector in Bangladesh.

Overview of the research

We use a unique ecosystem of data from Bangladesh, including three years of high-frequency (weekly) surveys of garment workers, whose plants we can match “in real time” with their exports activity, as recorded by customs offices. We combine these with human resources (HR) records from a sample of plants from Menzel and Woodruff (2021), which we can also match to customs records. Finally, we perform a number of information elicitation exercises, which allow us to characterise aspects of job search processes that are typically not observable to researchers, such as the information that workers have about job vacancies.

A simple partial equilibrium model of workers' job location choice allows for a decomposition of the gravity pattern into information (workers are less likely to be aware of and hence less likely to choose distant jobs) and commuting preferences (commuters dislike commuting far). We establish several stylised facts related to the spatial organisation of the labour market in the garment sector in urban Bangladesh, and the drivers of these spatial patterns. We then perform an illustrative decomposition exercise using the structure of the model, and the survey data. The research is ongoing, and the research design and the results reported in this brief are preliminary.

Data sources

We exploit four sources of data, each of which we describe in turn.

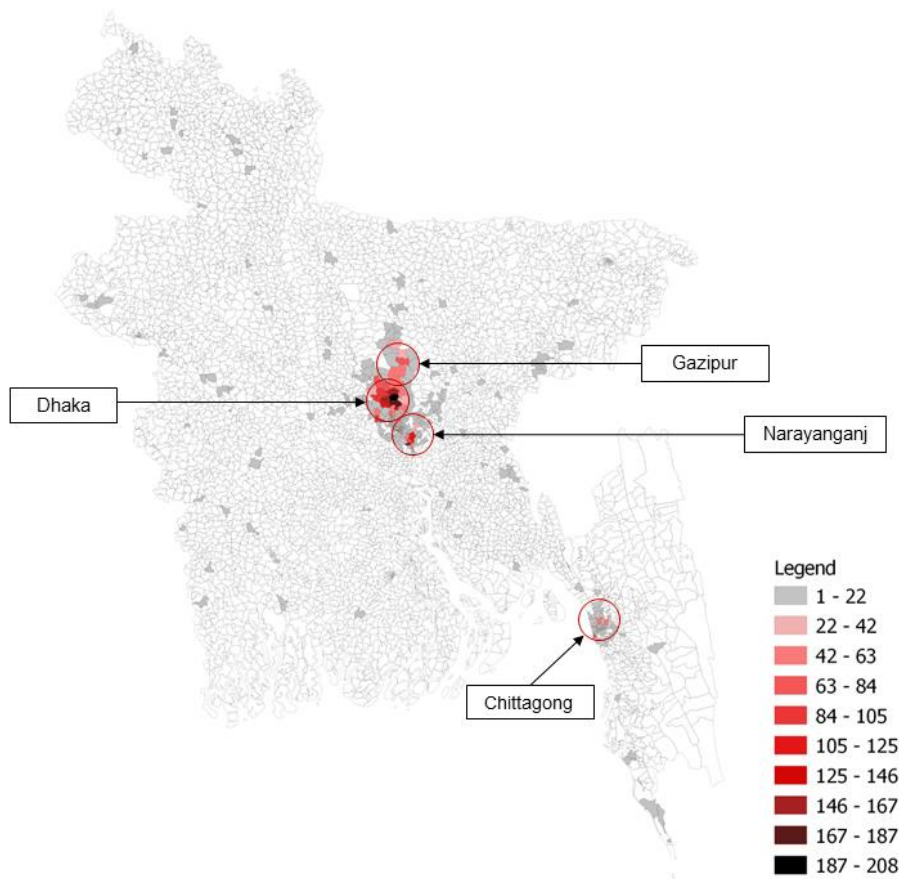
- **Customs Records.** We use administrative records from all customs offices in Bangladesh, detailing all garment exports transactions in the country, from January 2018 - March 2021. For each transaction, we observe the volumes and values, the date of the transaction and, crucially, the identity of the seller, which we can geo-locate and match with the other sources of data, described momentarily.
- **Workers High-Frequency Data and Preference Elicitation Surveys.** We draw from data collected by the International Labour Organization through their field partner, Micro-finance Opportunities (MFO). The main data collection instrument is a weekly survey of garment workers in urban Bangladesh, interviewed between April 2018 and December 2022. The survey accessed a total of 2,209 workers and includes 261,698 worker-week observations. In some of these weeks, workers answered special modules on their job search efforts and vacancy availability. In addition, we performed a randomised information provision and preferences elicitation exercise. Importantly, the workers included in the survey are geo-located on the basis of their residence and the location of the plant they work at, which is then matched to the customs records described above.
- **Other Data Sources.** This project has benefitted from the generosity of Andreas Menzel and Christopher Woodruff, who allowed access to the universe of HR payroll records of 83 Bangladeshi garment plants between 2012 and 2017 (Menzel and Woodruff, 2021). The plants in these data are matched with their customs records. In addition, we use the 2017 nationally representative Labour Force Survey (LFS) of the Bangladeshi Bureau of Statistics to benchmark our survey sample and characterise urban labour markets beyond our survey.

Context

We study the export-oriented, ready-made-garment sector in Bangladesh. The sector accounts for approximately 85% of the country's total exports, 15% of its

GDP and is a major driver of urban employment. Most garment plants in the country are located in a few districts, on which the sampling of workers for our survey took place (Figure 1). In these districts, the garment sector accounts for approximately 50% of the employment in manufacturing, based on computations using the LFS.

FIGURE 1: Garment-Exporting Plants Across Bangladesh



Notes: The figure shows the number of exporting plants in each union (Level 4 administrative unit) in Bangladesh, based on information in the Bangladesh Customs Records (2019). The four marked districts represent the districts from which workers in the High Frequency Survey (April 2018 – January 2020 and May 2020 to December 2022) were sampled.

In those districts, the density of garments-related activity is very high: using the LFS and our garment plants directory, we estimate an average of 5 garment plants and over 2,200 garment workers per square kilometre. As a result, a typical worker in our survey is surrounded by many other potential employers and garment workers.

On the demand side of the labour market, ready-made garment plants in Bangladesh are typically dedicated exporters that receive orders from international buyers and secure inputs on a buy-to-ship basis – they do not hold garment inventories or input stock; they procure materials by importing fabric through a drawback system, produce, and deliver. This feature of the environment, its near-just-in-time nature, implies that production and therefore

labour requirements are tied to shifts in the demand for garments that the plant's international buyers may experience.

On the supply side of the labour market, the typical worker in our survey is demographically similar with the typical garment worker in the sector: they are young – 26 years old on average –, predominantly female and married, and 89% are born in nearby rural villages outside of their current district of residence. The standard trajectory of these workers starts with them moving to the city as soon as they are of working age: they enter the sector at 18 and they stay in the sector for as long as they are physically able to cope with the demands of their activity. By the time of this survey, workers would have been in the sector for an average of six years, holding positions in two factories (including their current one). At any point in the survey, a small fraction of workers (15%) is unemployed.

The standard recruitment protocol in the sector begins with production chiefs estimating plants' labour requirements based on their current orders. The plant then announces vacancies through loudspeaker on production floors and, in some cases, posts a vacancy list on a notice board at the factory gate. These announcements normally take place on the last week of a month. On the first day of the following month, workers line up at the factory gate, are sorted according to their designation, and tested on a sewing machine, after which they are offered a job or not. This recruitment system, well-established in the industry, leads to most hiring taking place on the first few days of a calendar month, following orders received the previous month.

Key findings

We uncover five novel facts on worker mobility and labour market dynamics.


Fact 1. Gravity in commuting: workers work close to where they live.

Gravity in work-home commutes is documented in many settings, in both the developed and developing world, and more pronounced among low skilled workers (Tsvandis, 2022; Kreindler and Miyauchi, 2023). In our setting, this relationship is particularly steep in distance.

The vast majority of workers in our sample have short commutes and walk to work. We collect travel modes and commute times at three different points in time in our data collection period. Across the three instances, we find that approximately 80% of the workers commute on foot and another 10% combine walking with another mode of transport, typically rickshaw. Workers typically work six days a week, and half of the workers in our sample return home for lunch every day.

Based on workers' own estimations of their travel times, 75% of them have commutes that are shorter than 15 minutes (one way) in their chosen mode of transport, and shorter than 20 minutes in terms of walking time. Figure 2 shows that the median worker commutes less than a kilometre (0.86) and three quarters of the sample commute two kilometres or less. Regressions estimates, obtained by dividing the Dhaka and Chittagong metropolitan area into grids of 250 metres x 250 metres, support the steep decay in distance seen in Figure 2.

FIGURE 2: Gravity in Commutes: Residence-to-Factory Distances



Notes: Information on residential locations was collected in February 2022 and June 2023. We compute the distance (between geo-coordinates, in kilometres) to the plant the worker is employed in at the time of these survey blocks. The histogram presents results from 1,206 worker-plant pairs.

Fact 2. Gravity in job switches: job switches are valuable and highly local.

While worker commutes are local at any given point in time, it is possible that workers are mobile over space, across jobs. We show that when job switches occur, they are (i) profitable for workers, and (ii) highly localised.

A significant fraction of workers (30%) in our sample have never switched jobs, despite their tenure in the sector being high (5-6 years). Those that do switch, do so infrequently despite the observed high wage variability across plants. When workers are asked about the reasons behind their last (observed or not) plant switch, 95% of them report the switch being voluntary (i.e. the worker quitting rather than being laid off). Among those with voluntary switches, 58% indicate pay-related reasons for quitting (low wage, no promotion, late payment, overtime), 20% refer to personal reasons (family moving or health reasons),


14% mention commute related reasons (long or unsafe commute) and 9% report poor treatment of workers as a reason for quitting.

Having established that switches are infrequent and, in general, pay-motivated, it is possible to argue that there may not be significant differences in wages and amenities across potential jobs, to induce further switches. This argument is not supported by our data. Analysis using the LFS and HR data shows that even within upazilas (Level 3 administrative unit), there is high variability in residual wages, stemming from differences across plants.

When workers do switch plants, regressions using the High Frequency data show that they move to plants that are on average better in terms of pay-related characteristics. The new plants workers switch to on average pay over 400 BDT more a month (almost a 5% increase in pay), 1 BDT more per hour, offer work weeks that are 1.5 hours longer, and pay wages one day earlier than the plants workers leave.

Figure 3 shows the distribution of the distances between the origin (old) plant of the worker and the destination (new) plant of worker switches. The histogram shows that switching decays steeply over space, with a median distance between plants of 1.5 kilometres, and a 75th percentile of 7.3 km. While these between-factory distances are more spread over space than home-work commutes, our two sets of findings are broadly consistent. To see this, consider a scenario where the straight line joining the two plants intersects the worker’s residence at the mid-point of the line. In that case, commute patterns to both destination and origin plants are half of the distance between the two factories.

FIGURE 3: Gravity in Job Switches: Factory-to-Factory Distances



Notes: Switches are instances in which from one calendar month to the next, the worker is observed changing the plant in which they work. We do not consider switches from or into unemployment, or consecutive switches (for example, the worker is in plant A in March 2022, in plant B in April 2022 and in plant C in May 2022). Distance is computed in kilometres between the geo-coordinates of the two plants. The histogram presents results from 560 switches (pairs of plants).


Fact 3. Local labour supply: export shocks increase hiring at the plant; nearby export shocks increase churn at the plant.

We turn our attention to hiring dynamics in the sector, from the point of view of employers – the garment plants – and establish that (i) when a plant produces high export volumes, its hiring goes up, and (ii) when nearby plants increase their exports, the plant sees increases in the churning of its workers.

The recruitment process, described in the overview of the research, links the hiring of workers to exports. On the hiring side, the HR data shows that factories hire new workers almost every month, with an average of around 65 new recruits a month. Exits are also common, with an average of 85 workers leaving their plant in any given month. Taken jointly, there is a significant amount of churn in a month, with an entry rate of approximately 4% and an exit rate of just over 6% in factories that have, on average, over one thousand employees.

On the exports side, there is significant month-to-month variability, with export volumes doubling on high-season months, relative to the median month. Our data shows that when exports observe such increases (i.e. double), hiring rates increase by approximately 50%. Figure 4 shows the impact of changes in exports in nearby factories towards a plant's hiring and worker exits, conditional on the plant's own volumes. We aggregate the exports of all factories located between various distance bins. We find that export increases in nearby factories generate churning at the plant of reference, an effect that is highly localised. These nearby exports spill over onto the plant, increasing both exits and hiring, and leaving overall employment approximately unchanged. Crucially, this churning effect vanishes rapidly with distance: while some estimates are noisy, we find no systematic responses to export increases in plants that are more than one kilometre away from the factory.

FIGURE 4: Spillover of Export Shocks



Notes: The figure presents coefficients and standard errors for two regressions at the factory-month level, using 69 plants from the HR records, where the outcomes are the (log) of the total number of newly employed workers at the plant on that month, and the total number of workers who exit the plant on that month. The data were kindly facilitated by Andreas Menzel and Christopher Woodruff (Menzel and Woodruff, 2021). Standard errors are clustered at the level of the geographic zone (Level 4 administrative unit) and month. All specifications include factory and month fixed effects, and the log export volume for the factory-month combination (q_{im}). The bars for 0-1 km represents the coefficient for the log total export volume for all other plants within 0 to 1 km of a given plant, for the given month. Coefficients are re-scaled to be comparable with the coefficient q_{im} , which takes value 0.264 for New Hires, and -0.006 for Worker Exits.

This suggests that plants ‘poach’ workers from nearby plants, rather than facing an infinite pool of workers that they can attract at a market going rate. Additional exercises show that workers hired at times of production bottlenecks (when exports are high), make higher daily and monthly earnings. This follows from these workers serving a higher number of overtime hours. Incumbent workers also see increases in their overtime, both within and above the limits of legal overtime. Altogether, the patterns presented here indicate that plants compete for workers at a small spatial scale.

Fact 4. Value of Travel Time: workers report being willing to commute further away for good job opportunities.

The earlier exposition implies that home-work commutes and factory switches are very local, that plants compete for workers locally, and that the information workers have on job opportunities decays sharply with distance. A potential explanation for these patterns is that commuting is too costly for workers, so faraway job opportunities are irrelevant to workers.

However, when workers are asked to rank the top job attributes according to their preferences, only 1.2% report “transportation” as the top factor, and 5.5% report “easy commute”. The incidence of these responses is negligible in comparison with pay-related attributes such as “timely payment” (66.3%) or the opportunity to work for paid overtime (43.5%).


Beyond these descriptives, to study the role of commuting costs on workers’ job location decisions, we perform a stated preferences elicitation exercise. We estimate a value of travel time (VOTT) for garment workers of 82.2 BDT per hour, which is close to twice the wage in this setting. This estimate implies that, for example, a worker is willing to commute 10 minutes farther for a pay raise of 3.5 BDT per hour, equivalent to an 8.5% raise above the average hourly wage – well within the wage dispersion we observe in the LFS. The linear estimate masks considerable heterogeneity, and a richer specification finds that the VOTT is not statistically significantly different from zero for short distances between 10 to 20 minutes, and it is higher for longer travel times. Thus, high costs of commute (or high disutility of commute) cannot account fully for the gravity in commutes documented in Fact 1.

Fact 5. Gravity in information: workers are uninformed about job opportunities far away.

For six months, we time a set of special surveys – vacancy ledgers – to line up with the hiring cycle. In every round, workers report the total number of job vacancies they have heard about, and offer details about the most recent of these. We find that information about job opportunities is in general sparse, for both job seekers and non-job seekers. In each survey round, roughly 75% of workers did not hear of any job vacancies.

When available, the information about job opportunities is highly local. We document this in two ways. First, workers report that the median one-way walking commute time from workers’ residence to the hiring plant whose vacancy they heard about is 15 minutes. Second, we geo-locate the vacancies reported by workers and calculate the distance to the worker’s current plant. Figure 5 shows that the median hiring the worker hears is about is within 800 metres from their residence, the 75th percentile being 2.5 kilometres. This gravity pattern is common across both job-seekers and non-seekers.

FIGURE 5: Gravity in Information



Notes: Once every month in February, March, April and May 2022, workers reported information about job vacancies they had heard about. Over the four rounds, we obtain information on 1,166 plants that workers have heard about (685 distinct workers). The histogram presents the distribution of the distance between the worker's residential geo-coordinates and those of the plant whose vacancy the worker heard about.

Hence, many workers do not hear about vacancies and, when they do, these vacancies are spatially concentrated around the workers' current residence. The former may appear hard to reconcile with the fact that workers are based in areas with high factory density. The latter may lead us to believe that an abundance of heterogeneous jobs nearby makes further away opportunities unnecessary. We address these two points turn.

First, we study the profile of plants whose vacancies are known to workers. The plants that workers hear about are significantly larger, of which there are fewer. When we account for the higher probability of vacancies in large plants being heard about, we find 2–6 factories that hire between 0 and 1 kilometre from a worker's home. Crucially, the number of factories rises steeply with distance from home, yet workers very rarely hear about these farther away plants.

Second, we consider the possibility that local labour markets are saturated with sufficiently heterogeneous job opportunities, making a small radius of search efficient for the worker. We offer two pieces of evidence against this interpretation. First, we remind the reader that, on any given hiring cycle, a large share of the workers, including job seekers and the unemployed, do not hear about any vacancy. Contrary to the hypothesis of saturation, this pattern is consistent with local scarcity. Second, we find that plant wages (net of non-

wage attributes) are correlated in nearby locations but not with faraway locations. This suggests that the *effective* number of opportunities available in their local labour markets may be lower than it appears. Instead, accessing vacancies from more distant locations increases the options workers have. An information-elicitation experiment where workers are presented with specific plants at different distances, instead of vacancy ledgers, supports these conclusions. Workers are significantly less likely to have ever heard about faraway plants, relative to closer plants. Similarly, they are less likely to have information about the plant's hiring, even when both plants are known to the worker.

Policy implications

The facts discussed above suggest that distant-dependent information frictions play an important role at explaining the small spatial scale of urban labour markets in the context that we study. Leaning on model of workers' location choices, we establish that our main finding, namely that information frictions explain the spatial distribution of workers, have important implications for workers' welfare.

Disentangling the frictions (commuting costs or information) that lie behind a given pattern of worker mobility in the urban space is critical for policy. First, the decomposition allows policy-makers to deploy the 'right' intervention. For instance, in a setting (like ours) where the cause behind the spatial concentration of commuting is information, the most cost-effective tool for expanding the spatial extent of the labour market may *not* be a large infrastructure intervention. Second, for a given intervention that reduces effective distance, the expected welfare gain will vary with the extent of the distance-dependent information friction – producing different results relative to a perfect information benchmark.

In practice, in our setting, a sharp decay in workers' information about job opportunities as distance increases explains, largely, the hyper-local labour markets that we document in this paper. As such, this mechanism explains why most workers work close to where they live, why job switches among workers are rare and local despite high willingness to commute farther afield for extra pay and, ultimately, while inequality in labour market conditions persists, at close spatial proximity.

References

- Kreindler, Gabriel and Yuhei Miyauchi, "Measuring Commuting and Economic Activity inside Cities with Cell Phone Records," *The Review of Economics and Statistics*, 2023.
- Lall, Somik V., Vernon Marquette Henderson Jr, and Anthony J. Venables, *Africa's cities : Opening doors to the world*, African Regional Studies Program Washington, D.C. : World Bank Group., 2017.
- Menzel, Andreas and Christopher Woodruff, "Gender wage gaps and worker mobility: Evidence from the garment sector in Bangladesh," *Labour Economics*, 2021, 71, 102000.
- Tsivandis, Nick, "Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio," *mimeo* 2022.