

Integration Costs and Missing Women in Firms around the World

By CONRAD MILLER, JENNIFER PECK AND MEHMET SEFLEK*

The female share of the labor force varies dramatically across regions, ranging from 46% in sub-Saharan Africa to 20% and 24% in the Middle East and North Africa (MENA) and South Asia.¹ Prior research argues that this variation can be explained in part by social norms regarding gender roles and their effects on labor supply decisions (e.g., Fernandez and Fogli, 2009; Alesina, Giuliano and Nunn, 2013). Recent work shows that these norms also constrain *firm* behavior and can have persistent effects on aggregate female labor *demand* (Miller, Peck and Seflek, 2022; Jayachandran, 2021).

In particular, where there are social norms that favor gender segregation, firms face additional costs to *integration*, defined here as employing both women and men. Workers, customers, and regulators may expect firms to establish gender-segregated facilities, including restrooms and workspaces. Firms may also segregate tasks or teams to limit interactions between men and women. For male-dominated firms, hiring women may necessitate changes in their workplace culture. In World Bank Enterprise Surveys from 2013 and 2014, 29% of South Asian firms claim that hiring women “could cause disruption in the working environment” and cite this as a constraint to hiring women. Firms may also need to restructure tasks and occupations to enable female employment or to reorganize their production process to comply with restrictions on women’s work-

ing hours.² These types of restrictions all constrain the production process for integrated firms. A distinguishing feature of the costs associated with accommodating social norms is that they often have a significant *fixed* component. We study the consequences of these integration costs for women’s employment.

The distribution of female employment across firms provides prima facie evidence that integration is costly for firms. In most of the world, all-male firms are rare—in sub-Saharan Africa, East Asia and the Pacific, Eastern and Central Europe, and Latin America and the Caribbean, only 2%–11% of medium (20–99 employee) manufacturing firms are all-male, and 1%–2% of large (100+ employee) firms are all-male.³ This is unsurprising given that women make up 27%–41% of employees in these firms overall. However, all-male firms are dramatically more common in MENA and South Asia: 48% and 50% of medium firms and 23% and 29% of large firms are all-male. While women make up a smaller share of employees in surveyed firms in these regions (17% and 15%), the share of firms that are all-male is substantially larger than what one would expect by chance.⁴ We show that this cannot be explained by differences in the female share of workers across regions or occupations. Instead, the mass of firms with zero female employees in MENA and South Asia

* Miller: Haas School of Business, UC Berkeley, Berkeley, CA 94720-1900, ccmiller@berkeley.edu. Peck: Swarthmore College, 500 College Avenue, Swarthmore PA 19081, jpeck1@swarthmore.edu. Seflek: Haas School of Business, UC Berkeley, Berkeley, CA 94720-1900, mseflek@berkeley.edu.

¹As measured by the World Bank and International Labour Organization in 2018. According to the World Bank, unpaid workers, family workers, and students are often omitted from these figures.

²The majority of countries in MENA and South Asia prohibit women from working at night, with such regulations typically justified by concerns over safety (International Finance Corporation, 2013; World Bank, 2018).

³See calculations in Miller, Peck and Seflek (2022). All-female firms are even more rare: less than 1% of medium and large firms are all-female.

⁴For example, if the gender of each employee were an independent draw from a binomial distribution where the probability an employee is female is 0.169 (the female share of manufacturing employment in MENA), the probability that a firm with 50 employees is all-male is $(1 - 0.169)^{50} = 0.0001$.

strongly suggests that firms face an extensive margin decision of whether to integrate their workforce. We deploy methodology developed in Miller, Peck and Seflek (2022) that uses the distribution of female employment across firms to estimate the share of firms with binding integration costs and counterfactual female employment at all-male firms. We find evidence of significant costs in MENA and South Asia, consistent with regional variation in social preferences for gender segregation. Integration costs are more binding in countries with stronger gender-segregation preferences, and our measures of these costs are negatively correlated with female labor force participation (LFP).

I. Data

We use cross-country data from the World Bank Enterprise Surveys, which survey private sector companies across 139 countries between 2005 and 2018. We limit our sample to manufacturing firms, where surveys record the gender composition of employees by occupation (production or nonproduction). We drop surveys where information on gender composition is missing for more than 20% of firms, firms with missing data on gender composition or with fewer than five employees, and then drop surveys with fewer than 100 remaining firms. This leaves 105 surveys in 65 countries.⁵ Survey participants are sampled by stratified random sampling with firm size categories, sectors, and within-country geographic regions as the strata. We use the provided survey weights so that firms are representative of the manufacturing sector in a given country and year.

We link our measure of integration costs to social preferences using data from the Arab Barometer, a cultural, religious, and political opinion survey that is run periodically across several MENA countries. The survey is designed to cover a representative sample of adults within each country and records gender segregation preferences, with all four survey groups asking about

agreement or disagreement with a statement regarding gender mixing in university classes.⁶ We use these responses to measure the share of respondents in each country that approve of gender mixing and match this information to the World Bank data.

II. Empirical Strategy and Results

If the costs of integration are largely fixed, firms will integrate only if their expected number of female employees under integration exceeds some threshold. Firms that have not paid these fixed integration costs are *ex-ante segregated*, i.e. face some binding fixed cost of integration. Our analysis follows the methodology developed in Miller, Peck and Seflek (2022) to test for integration costs and calculate ex-ante integration rates in Saudi firms. The empirical strategy is based on a partial equilibrium model of firm hiring in which (1) firms face an extensive margin integration decision (i.e. decide ex-ante whether to incur the integration costs associated with hiring both men and women) and (2) integration costs are largely fixed, so firms integrate only if they anticipate employing enough women to justify the costs. The probability that the top candidate for position i is female is modeled by a function $\theta(\cdot)$ of observable job characteristics, X_i . To test the null that all firms are ex-ante integrated, we estimate $\theta(X_i)$ using data on employees at all firms, simulate the distribution of female employment across firms using this estimate, and compare the simulated and observed distributions. Under the null, we would expect some firms to have zero female employees by chance alone. However, if integration costs bind for *some* firms, we show that we should see an excess mass or “bunching” of firms with zero female employees. We then re-estimate $\theta(X_i)$ and use this estimate to measure ex-ante integration rates as a function of expected female employment under integration.

⁶The statement is “It is acceptable in Islam for male and female university students to attend classes together” in waves I (2006-2009) and IV (2016-2017), and “Gender-mixed education should be allowed in universities” in II (20102011) and III (20122014).

⁵Online Appendix Table A2 lists the surveys we include.

A. Testing the Null of No Binding Integration Costs

We begin by testing the null that integration costs are nonbinding for all firms separately by region. Following the procedure developed in Miller, Peck and Seflek (2022), we calculate a separate $\theta(X_i)$ function for each survey (country-year pair), estimating a logistic regression model of the form

$$P(\text{Worker } i \text{ is female}) = \Lambda(X_i\beta),$$

and include location and occupation type (production or nonproduction) in X_i . We then simulate the distribution of female employment across firms under the null hypothesis given the estimated $\hat{\theta}^0(X_i)$ and compare that to the observed distribution.

Figure 1 compares the distributions for Ethiopia and India.⁷ In Ethiopia our simulation matches the actual distribution reasonably well. In particular, we do not substantively underpredict the number of all-male firms. By contrast, the pattern in Egypt is similar to the results from Miller, Peck and Seflek (2022) in Saudi Arabia: we substantially underpredict the number of firms with zero female employees. In Egypt, we predict that 22% of firms will be all male; in fact, 63% are. We also overpredict the number of firms with few female employees.

One concern with this comparison across countries is that we may see more bunching at zero in India and Egypt simply due to *censoring* at zero. To address this concern, we examine the same distributions for larger firms, where almost no firms would have zero female employees under the null. In Appendix Figure ??, we plot the same distributions for firms with at least 50 employees. Across all countries, virtually no firms have zero female employees in our simulations, and in Ethiopia, China, Russia, and Brazil, virtually none of these firms are all male in practice. Yet, about 30% of such firms in India and Egypt are all male.

If some firms are ex-ante segregated, then

$\hat{\theta}^0(X_i)$ will underestimate $\theta(X_i)$, since some of these segregated firms are included in its estimation. To correctly estimate $\theta(X_i)$ we therefore repeat the exercise limiting the data to integrated firms. This yields two functions that approximate $\theta(X_i)$ for ex-ante integrated firms: one that uses data on *ex-post* integrated firms ($\hat{\theta}^{EP}(X_i)$) and another that uses a structural model to jointly estimate $\theta(X_i)$ and the probability that a firm is ex-ante integrated as a function of firm characteristics, ($\hat{\theta}^S(X_i)$).⁸ We use these approximations of $\theta(X_i)$ in the remainder of the analysis.

B. Ex-Ante Integration Rates

We next calculate ex-ante integration rates as a function of $\bar{\theta}_j n_j$, a firm's expected number of female employees if ex-ante integrated. We use our estimates of $\theta(X_i)$ to estimate counterfactual female employment at non-integrated firms:

$$\sum_{i \in \text{firm } j} \theta(X_i) n_{ij} = \bar{\theta}_j n_j,$$


where n_{ij} is the number of type i jobs at firm j , n_j is the number of jobs at firm j , and $\bar{\theta}_j$ is average value of $\theta(X_i)$ at firm j given its job composition. We estimate $\theta(X_i)$ and $\bar{\theta}_j$ using ex-post integrated firms. In Figure 2, we plot ex-ante integration rates for our six most populous countries in the World Bank data. In Ethiopia, China, Russia, and Brazil, ex-ante integration rates are close to 100% for all values of $\bar{\theta}_j^{EP} n_j$. Ex-ante integration rates are uniformly lower in Egypt and India. As shown below, this pattern corresponds with more general regional differences in integration.

We construct three measures to summarize integration costs by country. We calculate average ex-ante integration rates in two ways: first using $\hat{\theta}^{EP}(X_i)$ based on data from ex-post integrated firms and second using $\hat{\theta}^S(X_i)$ from overall ex-ante integration rates estimated using the structural approach described in Miller, Peck and Seflek

⁷Appendix Figure A1 compares the distributions for the largest countries in each region represented in our data: Ethiopia, China, Russia, Brazil, India, and Egypt.

⁸See Miller, Peck and Seflek (2022) for details on how this structural model is specified and estimated.

Figure 1. : Distribution of Female Employment across Firms: Ethiopia and Egypt



Note: These figures compare observed and simulated distributions of female employment across firms in Ethiopia and India. The simulated distributions are simulated under the null that no firms face binding integration costs.

(2022). For our third measure, we calculate the implied ex-ante integration rate for a “representative” firm with $\hat{\theta}_j^{EP} n_j = 10$.

Social preferences are correlated with our measures of ex-ante integration rates within the MENA region. Figure 3 plots ex-ante integration rates by country (derived using ex-post integrated firms) against support for gender mixing from the Arab Barometer surveys. Countries with less support for gender mixing have lower ex-ante integration rates. This is consistent with our motivating hypothesis that integration is costly for firms where social norms for gender segregation are strong.

Ex-ante integration costs are also correlated with women’s labor market outcomes across countries. Figure 4 reports OLS estimates of a regression of LFP measures on all three measures of integration costs and plots country-level female LFP rates against overall ex-ante integration rates (derived using ex-post integrated firms). There are two main points to note. First, ex-ante integration rates are lowest in MENA and South Asia. Ten of the fourteen MENA and South Asian countries in our data have average ex-ante integration rates below 60%, and no other countries fall in this range. Second, ex-ante integration rates are positively correlated with female LFP rate, both in absolute terms (columns (1) and (2)) and relative to male LFP (columns (3) and (4)).⁹ The posi-

tive relationship between ex-ante integration rates and women’s economic engagement also holds within regions (see columns (2) and (4), specifications that include region fixed effects).

III. Conclusion

In countries where there are social norms for gender segregation, firms face costs to employing both men and women that are largely fixed. We use a joint test for whether integration costs bind for any firm and a methodology for evaluating the firm-level consequences of those costs using cross-sectional World Bank survey data covering 65 countries. We find that a large fraction of firms in MENA and South Asia employ only men due to integration costs, but integration costs do not constrain firms in other regions. This is consistent with regional variation in social preferences for gender segregation. Within MENA, we find that variation in integration costs across countries is consistent with local preferences. Our results suggest that integration costs prevent some firms from hiring superior female candidates, with significant consequences for women’s labor market outcomes. Overcoming these barriers may be key to increasing women’s employment and labor force participation in areas where these constraints bind.

ployment rates rather than LFP rates (see Appendix Table ??).

⁹The relationship is similar when we examine em-

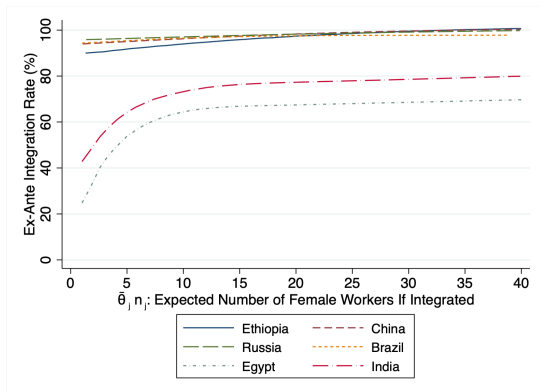


Figure 2. : Integration Rates by $\bar{\theta}_j n_j$

Note: This figure depicts the relationship between ex-ante integration rates and $\bar{\theta}_j n_j$, a firm’s expected number of female employees if ex-ante integrated.




Figure 3. : Integration Rates and Support for Gender Mixing in MENA

Note: This figure plots ex-ante integration rates ($\hat{\theta}^{EP}(X)$) against local support for mixed-gender university classes across MENA countries.




Figure 4. : Female Labor Force Participation and Integration Rates Across Countries

Note: This figure plots female LFP rates against ex-ante integration rates for 65 countries. Ex-ante integration rates are derived using country-specific $\hat{\theta}^{EP}(X_i)$, with estimates of $\theta(X_i)$ derived using ex-post integrated firms. Table reports OLS estimates of regressions of female LFP measures on ex-ante integration rates. LFP_F is the average percentage of women age 15+ in the labor force, and $LFP_F - LFP_M$ is the difference between female and male LFP. Robust standard errors are in parentheses. \sim p<0.1; * p<0.05; ** p <0.01.

	LFP_F		$LFP_F - LFP_M$	
Ex-ante integra- tion rate:	(1)	(2)	(3)	(4)
Overall ($\hat{\theta}^{EP}$)	0.425** (0.079)	0.278** (0.089)	0.453** (0.067)	0.298** (0.075)
Overall ($\hat{\theta}^S$)	0.615** (0.133)	0.287** (0.142)	0.638** (0.117)	0.292* (0.123)
Rep. firm ($\hat{\theta}_j^{EP} \cdot n_j = 10$)	0.489** (0.081)	0.337** (0.093)	0.496** (0.070)	0.355** (0.077)
Region FEs		✓		✓
Observations	65	65	65	65

REFERENCES

Alesina, Alberto, Paola Giuliano, and Nathan Nunn. 2013. “On the Origins of Gender Roles: Women and the Plough.” *Quarterly Journal of Economics*, 128(2): 469–530.

Fernandez, Raquel, and Alessandra Fogli. 2009. “Culture: An Empirical Investigation of Beliefs, Work, and Fertility.” *American Economic Journal: Macroeconomics*, 1(1): 146–177.

International Finance Corporation. 2013. *Investing in Women’s Employ-*

ment: Good for Business, Good for Development. World Bank.

Jayachandran, Seema. 2021. “Social Norms as a Barrier to Womens Employment in Developing Countries.” *IMF Economic Review*.

Miller, Conrad, Jennifer Peck, and Mehmet Seflek. 2022. “Missing Women, Integration Costs, and Big Push Policies in the Saudi Labor Market.” *American Economic Journal: Applied Economics*, 14(2).

World Bank. 2018. *Women, Business and the Law 2018.* The World Bank.