

NBER WORKING PAPER SERIES

INTEGRATION COSTS AND MISSING WOMEN IN FIRMS

Conrad Miller
Jennifer Peck
Mehmet Seflek

Working Paper 26271
<http://www.nber.org/papers/w26271>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2019

We thank Hadar Avivi for excellent research assistance. We thank Peter Blair, Patricia Cortes, Alessandra Gonzalez, Supreet Kaur, Pat Kline, Peter Kuhn, David Levine, Jonathan Leonard, Jeremy Magruder, Brendan Price, Joe Shapiro, Isaac Sorkin, Melanie Wasserman, and seminar participants at Berkeley (Haas and Economics), Brown, Dartmouth, Harvard (WiGi and Growth Lab), IMF, Michigan, NBER Labor Studies, Princeton (AALIMS), Swarthmore, and Stanford (SIEPR) for comments. This research was sponsored by a grant from the Evidence for Policy Design program at the Harvard Kennedy School and the Human Resources Development Fund of the Kingdom of Saudi Arabia. This paper subsumes and extends a paper previously circulated as “Big Push Policies and Firm-Level Barriers to Employing Women: Evidence from Saudi Arabia.” The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w26271.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Conrad Miller, Jennifer Peck, and Mehmet Seflek. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Integration Costs and Missing Women in Firms
Conrad Miller, Jennifer Peck, and Mehmet Seflek
NBER Working Paper No. 26271
September 2019
JEL No. J16,J23,J71,O53

ABSTRACT

Where social norms favor gender segregation, firms may find it costly to employ both men and women. If the costs of integration are largely fixed, firms will integrate only if their expected number of female employees under integration exceeds some threshold. Motivated by a simple model of firm hiring, we develop a methodology 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 validate our approach using administrative data and unique policy variation from Saudi Arabia. We provide suggestive evidence that integration costs reduce aggregate female employment. Using survey data on manufacturing firms in 65 countries, we find significant integration costs in the Middle East, North Africa, and South Asia but not in other regions.

Conrad Miller
Haas School of Business
University of California, Berkeley
Berkeley, CA 94720
and NBER
ccmiller@berkeley.edu

Mehmet Seflek
Haas School of Business
University of California, Berkeley
Berkeley, CA 94720
mseflek@berkeley.edu

Jennifer Peck
Department of Economics
Swarthmore College
500 College Avenue
Swarthmore, PA 19081
United States
jpeck1@gmail.com

I Introduction

The female share of the labor force varies dramatically across regions, ranging from 46.4% in sub-Saharan Africa to 20.3% 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 et al., 2013). In MENA and South Asia, strong preferences for gender segregation are common (Jayachandran, 2015). Women or their families may prefer that they not work in environments that require interaction with unrelated men. In this paper, we show that social norms constrain *firm* behavior too.

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, entrances, and workspaces. Firms may also segregate tasks to limit interactions between male and female employees. For male-dominated firms, hiring women may necessitate changes in their workplace culture. In World Bank 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.² The majority of countries in MENA and South Asia prohibit women from working at night; such regulations are typically justified by concerns over safety (International Finance Corporation, 2013; World Bank, 2018). These restrictions constrain the production process for integrated firms. Moreover, workers may prefer to work with same-sex coworkers (Gorman, 2005). We study the consequences of these integration costs for female employment.

The distribution of female employment across firms provides *prima facie* evidence that integration is costly. Table I presents the fraction of medium (20–99 employees) and large (100+ employees) manufacturing firms that employ only men in a sample of 65 countries, grouped by region.³ In most of the world, all-male firms of this type are rare—in sub-Saharan Africa, East Asia and the Pacific, Eastern and Central Europe, and Latin America and the Caribbean, only 1.8%–10.5% of medium firms are all male, and 0.5%–2.3% of large firms are all male.⁴ This is unsurprising given that women make up 27.0%–41.2% of employees in these firms overall. However, all-male firms are dramatically more common in MENA and South Asia: 48.1% and 49.9% of medium firms and 22.7% and 28.6% of large firms are all male. While women make up a smaller share of employees in surveyed firms in these regions (16.9% and 14.5% in MENA and South Asia), the share of firms that

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.

²These data are from World Bank Enterprise Surveys in Afghanistan (2014), Bangladesh (2013), India (2014), Nepal (2014), and Pakistan (2013). In addition, 29% of firms cite “challenges of hiring women given government regulations such as working hours and maternity leave,” and 32% of firms cite “required benefits and other expenses such as providing separate workplace facilities for women make them more expensive employees” as obstacles to hiring women.

³We describe these data in more detail in Section VI.

⁴All-female firms are even more rare: between 0.1%–1.3% and 0.1–0.4% of medium and large firms are all female. In MENA and South Asia, 0.1%–0.2% of medium and large manufacturing firms are all female.

are all male is substantially larger than what one would expect by chance.⁵ As we show below, this cannot be explained by differences in the female share of workers across regions, occupations, or industries. Instead, the mass of firms with zero female employees in MENA and South Asia strongly suggests that firms face an extensive margin decision of whether to integrate their workforce.

[Table 1 about here.]

Motivated by a simple model of firm hiring, we develop a methodology that uses the distribution of female employment across firms to assess whether and how integration costs constrain female employment at firms. We apply and validate the methodology using administrative data and unique policy variation from Saudi Arabia, a country that strictly limits between-gender interactions in the workplace over our sample period. We find that the majority of Saudi firms employ only men because they face binding integration costs. We document suggestive evidence that integration costs reduce aggregate female employment. We then apply our methodology to survey data on manufacturing firms in 65 countries and find significant integration costs in MENA and South Asia but not in other regions. Within MENA, we find that variation in integration costs across countries correlates with local preferences for gender segregation.

We first build a simple, partial equilibrium model of firm hiring based on Kuhn and Shen (2013). A firm posts an exogenously determined number of vacancies and receives a random draw of candidates for each vacancy. The firm hires their most preferred candidate for each vacancy from the candidate pools: they can choose to hire from a pool of only male candidates or pay a fixed integration cost and hire from both pools for all vacancies. This framework generates a threshold rule: firms pay for the ability to hire both male and female candidates (“ex-ante integrate”) if their expected number of female hires under integration is sufficiently large.⁶

Guided by the model, we develop a joint test for whether all firms are ex-ante integrated (i.e., whether integration costs affect hiring at any firm). We assume that, for ex-ante integrated firms, the probability of a female hire for a vacancy i is a function, $\theta(\cdot)$, of observable job characteristics, X_i . In our data, X_i includes occupation, industry, and the location of the job.⁷ Under the null hypothesis that integration costs do not bind at any firm, all firms are ex-ante integrated, so $\theta(X_i)$ can be estimated using data on employees at all firms. To test the null hypothesis, we simulate the distribution of female employment across firms using this estimate for $\theta(X_i)$ and compare that to the observed distribution. Under the null, we would expect some firms to have zero female employees by chance alone.⁸ However, if integration costs bind for *some* firms, we show that we

⁵For example, if the gender of each employee were independent draws from a binomial distribution where the probability an employee is female is 0.169, the female share of the manufacturing sector in MENA, then the probability that a firm with 50 employees is all male is $(1 - 0.169)^{50} = 0.0001$.

⁶In principle, some firms are at the margin of hiring from only a pool of female candidates and an integrated pool. Symmetrically, these firms ex-ante integrate only if their expected number of *male* hires under integration is sufficiently large. In practice, in the settings where integration costs are relevant, all-female firms are rare, suggesting few firms are on this margin. In the context of our model, this can be rationalized by the fact that women make up a small share of the workforce in these markets.

⁷We test and find support for this assumption in several ways described below.

⁸By chance, all the top candidates may be male for some firms. We would also expect this if $\theta(X_i)$ were zero for some vacancies.

should see an excess mass or “bunching” of firms with zero female employees.

In Saudi Arabia, we find exactly this pattern and reject the null hypothesis of no binding integration costs. In January 2009, our first month of data, 8% of Saudis in the private sector are female. If θ were constant across jobs, we simulate that 34% of firms in our sample would employ only men. If we allow X_i to include controls for location, occupation, and industry, this increases to 43%. In practice, 73% of firms in our sample employ only men. The relatively limited role for observable job characteristics under the null suggests unobservable job characteristics are an unlikely explanation for the larger-than-predicted share of all-male firms. By contrast, in a placebo test, we cannot reject the null hypothesis that there are no costs associated with employing both married and unmarried men.

Equipped with $\theta(X_i)$, we can estimate counterfactual female employment at each segregated all-male firm. However, if some firms face binding integration costs, the above approach will underestimate $\theta(X_i)$ by including firms that are not (ex-ante) integrated in its calculation. Instead, to correctly estimate $\theta(X_i)$, we must limit the data to ex-ante integrated firms. Accomplishing this is complicated by the fact that we only observe whether firms are “ex-post” integrated—whether they employ both men and women in practice. We take two approaches. First, we ignore the distinction and limit the data to ex-post integrated firms. This will lead to a slight upward bias in our estimate for $\theta(X_i)$. Second, we estimate a more parametric model that accommodates ex-ante integration as a potentially unobserved firm state. The two approaches yield similar estimates.

For each segregated firm, we use the firm’s job mix and our estimate for $\theta(X_i)$ to predict what its female employment would be if it were to integrate, holding the behavior of other firms fixed. Our estimates imply that about 65% of Saudi firms face binding integration costs. We find that ex-ante integration rates are increasing in a firm’s expected number of female employees if integrated, consistent with largely fixed integration costs. We match the observed distribution of female employment using our estimate of $\theta(X_i)$ alone to simulate each firm’s integration status and female employment.

While our ability to match the observed distribution of female employment is reassuring, it remains possible that our estimate of $\theta(X_i)$ based on integrated firms does not generate accurate estimates for counterfactual female employment at segregated firms. In fact, it is possible that with sufficiently heterogeneous preferences or candidate pools across firms, the observed distribution of female employment can be rationalized in the absence of any integration costs at all. Moreover, while a firm’s chances of integrating are increasing in its expected number of female employees under integration, it is not clear that this relationship is *causal*, as our model implies.

To further validate the model, we use two features of the Saudi data and context. The first feature is the panel structure of the data. We conduct two tests that exploit this feature. In the first test, we examine firm transitions from segregated to integrated using an event study framework. Our theory predicts “lumpy” transition dynamics; once a firm integrates, it hires women at rates similar to that of incumbent integrated firms with a similar mix of jobs. This prediction is borne out; six months after a newly integrated firm’s first female hire, 26% of their hires are female. This

matches the female share of hires for similar incumbent integrated firms. Moreover, our estimate of $\theta(X_i)$, derived from incumbent integrated firms, predicts the female share of hires across newly integrated firms with little bias.

We also test for *state dependence*, comparing hiring behavior at previously segregated and previously integrated firms. We expect previously integrated firms to tend to remain integrated, either because employee turnover is low, integration costs are sunk, or firm conditions that make integration appealing in the first place are persistent. Hence, we should see that hiring patterns in those firms exhibit little bunching and are consistent with no binding integration costs. Consistent with state dependence, we find strong evidence of bunching at previously segregated firms but not at previously integrated firms.

The second feature we use is a unique policy that generates exogenous variation in Saudi employment across firms. Like many Gulf countries, the Saudi labor market is characterized by high youth unemployment and low participation of nationals in the private sector. In 2011, the Saudi government launched *Nitaqat*, an ambitious gender-neutral nationalization quota policy designed to increase private sector employment of Saudi nationals. As shown in Peck (2017), firms generally respond to Nitaqat quotas by employing more Saudis. Firms vary in their distance from their quota at baseline, generating exogenous variation in Saudi employment growth across firms. We use this variation to further test the model. We find that firms that are above and below their Nitaqat quotas at the time the policy is implemented have similar observable characteristics and are on similar pre-trends. Following implementation, firms that are below their quotas experience a larger increase in Saudi employment and hiring. Consistent with the model, among previously all-male firms, we also find that below quota firms integrate at higher rates and have a larger female share of hires. Moreover, the magnitude of the increase in female share of hires is in line with what we would predict based on our estimate of $\theta(X_i)$ derived from incumbent integrated firms. We conclude that integration costs are an important driver of firm behavior in Saudi Arabia and our stylized model fits the data well.

While integration costs reduce female employment at individual firms, it is not clear what implications integration costs have for female labor market outcomes in the aggregate. As in Becker (1957), integrated firms may be sufficiently numerous or large to absorb female labor so that the existence of constrained male-only firms has no bearing on female wages and employment. On the other hand, in the presence of search frictions or insufficient entry or growth of integrated firms, integration costs will reduce aggregate demand for female labor.

We document suggestive evidence that integration costs reduced aggregate female employment in Saudi Arabia. In particular, we find that the Nitaqat quota policy nearly tripled the female share of Saudis working in the private sector within four years, from 10% in 2011 to 27% in 2015. This increase is concentrated at firms that were previously all male and were induced to integrate by the policy. This occurs despite a *decrease* in the gender wage gap over this period, which is in part driven by the introduction of a de facto minimum wage for Saudi workers in the private sector. Together, these findings suggests that Nitaqat increased relative aggregate demand for women by

inducing more firms to integrate.

We then apply our empirical strategy to World Bank survey data covering manufacturing firms in 65 countries. While Saudi Arabia has regulations that explicitly introduce integration costs, there may be significant integration costs in other countries with less strict regulations but where there are strong social norms for gender segregation. Consistent with Table I, we find significant integration costs in MENA and South Asia but not in other regions. For example, we estimate that 50% of Indian manufacturing firms are ex-ante segregated. This pattern is consistent with regional variation in social preferences for gender segregation (Jayachandran, 2015). Using survey data on social preferences across countries within MENA, we find that integration costs are more binding in countries where inhabitants prefer gender-segregated classes in local universities. Across countries, we also find a negative relationship between the female labor force participation rates and the share of firms that are ex-ante segregated. This relationship holds both between and within regions.

The notion that gender integration involves substantial, largely fixed costs has important implications for policy. In particular, our results suggest that “big push” demand-side policies that incentivize firms to integrate can substantially change firm hiring preferences at the margin.⁹ These policies can also have the potential for feedback effects by attracting more women to the labor market, which could in turn induce more firms to integrate. Though we cannot test this directly here, our results also suggest that one-time incentives to integrate may have long-lasting effects on female employment. This is because the types of costs we believe are associated with gender integration in this context—physical investment in new or restructured workspaces and facilities, change in organizational structure or culture—have a significant sunk component.¹⁰

We contribute to a large literature on how social and cultural norms affect women’s labor market outcomes. This literature primarily focuses on how social norms influence labor supply decisions (e.g., Fernandez, 2013). Most closely related is Bursztyjn et al. (2018), who study social norms over women’s labor supply in Saudi Arabia. They show that husbands underestimate the share of their peers who support wives participating in the labor market, and they provide evidence that correcting those misperceptions increases husbands’ willingness to support their wives joining the labor force. By contrast, we focus on how norms constrain labor demand and how firms respond to those constraints. Our paper also studies how policy interacts with cultural norms, as do Ashraf et al. (2019), who study how bride price traditions mediate educational investments in daughters’ education and household responsiveness to education policy.

⁹In the literature on racial discrimination in the US, there is evidence that affirmative action policies can have long-term effects on minority employment even after the policies end. Miller and Segal (2012) study affirmative action quotas imposed by federal courts on municipal police departments and other law enforcement agencies in the 1970s and find that black share gains due to the quotas do not erode after the policy ends. Miller (2017) studies federal affirmative action regulation of private employers and finds that the black share of employees continues to grow even after firms are no longer regulated. He argues that this persistence is in part driven by persistent changes in screening methods for potential hires. Relatedly, Whatley (1990) argues that the introduction of black workers into previously all-white Cincinnati firms during World War I made those firms more likely to hire black workers in the near future and employer learning is a key driver of this pattern.

¹⁰Policies that increase exposure to integrated workplaces may also change gender attitudes, which could in turn affect integration costs (Dahl et al., 2018).

Methodologically, our approach to inferring integration costs is similar in spirit to bunching estimators (Kleven 2016). Canonical bunching estimators exploit bunching in observed income distributions around discontinuities in tax rates to measure behavioral responses to taxes and transfers. In an application more closely related to ours, Garicano et al. (2016) and Gourio and Roys (2014) examine bunching in the firm size distribution to study the costs of labor regulations that apply to firms above a known size threshold. By contrast, we infer the existence of integration costs—a discontinuity at zero in the incentive to employ women at the margin—based on observed bunching at zero in the distribution of female employment across firms. To test for binding integration costs, we compare the observed distribution of female employment across firms to what we would predict if the gender of each employee were a random draw, conditional on observable job characteristics. The spirit of the exercise is similar to Ellison and Glaeser (1997) and Augereau et al. (2006).

This paper also relates to the literature on workplace segregation and its implications for labor market inequality. This literature has primarily focused on the United States. Tomaskovic-Devey et al. (2006) show that US men and women are segregated across firms, though this segregation has declined since 1960. As some firms pay more than others, this segregation can have important implications for gender earnings inequality (Groshen, 1991; Bayard et al., 2003; Card et al., 2016). While prior research has shown that skill differences and occupational preferences can explain between-establishment segregation to some degree, at least along the intensive margin, there is little research explaining why some firms employ no women at all.¹¹ Our findings also contrast with Pan (2015), who documents that the gender composition of occupations have historically displayed “tipping points,” where occupations like bank tellers and typesetters quickly transitioned from predominantly male to predominantly female. This pattern is consistent with men preferring to work in occupations with a higher male share of coworkers.¹² Instead, in Saudi Arabia, some firms employ only men, but when they begin to hire women, they quickly transition to look like incumbent integrated firms, which employ a mix of men and women.¹³

Finally, we build on a literature that studies dynamics and adjustment costs in firm-level labor demand, primarily as an input for understanding macroeconomic fluctuations. Firms are modeled as facing costs when making net or gross adjustments to their level of employment. A series of papers document that firms tend to change employment in a manner consistent with nonconvex

¹¹Using data on US manufacturing workers in 1890, Goldin (1986) argues that occupations within firms are segregated in part because female workers have higher turnover rates and are therefore better suited for piece rate compensation when monitoring is costly.

¹²In the Goldin (2014) “pollution theory” of discrimination, men may prefer to have more male workers if it increases the prestige associated with their work.

¹³Though integration costs do not appear to substantively constrain firm behavior outside of MENA and South Asia today, they likely played a more significant role in the past. Humphries (1987) describes social anxiety regarding integrated workplaces in 19th-century England, which contemporaries worried would lead to “immoral” interactions. She argues that these attitudes increased industrial segregation by sex. Breckinridge (1906) describes contemporary attitudes in the US: “it is well known that the unregulated mingling of men and women under conditions of darkness, fatigue, or the excitement due to the constant apprehension of danger may give rise to immoral intercourse. On this account we find women generally prohibited from working in mines, and . . . other forms of employment at night.” At the turn of the century, state laws restricted the hours that women could work in factories and required separate restrooms and dressing rooms (Mies, 1906).

adjustment costs: adjustment tends to be lumpy, with extended periods of inactivity and sharp, large changes (see, e.g., Varejão and Portugal, 2007; Hamermesh and Pfann, 1996). We study a different type of adjustment, moving from an all-male to an integrated workforce, and document that the pattern of adjustment within and across firms is consistent with largely fixed, potentially one-time adjustment costs.

The remainder of the paper is organized as follows. Section II introduces a model of firm hiring in the presence of integration costs. Section III describes the Saudi Arabia context and data. In Section IV, we use the model to develop a methodology for assessing whether and how integration costs constrain female employment at firms. We apply the methodology and test additional model predictions using Saudi data. In Section V, we discuss the aggregate consequences of integration costs for female labor market outcomes, using the introduction of Nitaqat employment quotas as a case study. Section VI assesses integration costs in 65 other countries using World Bank survey data. Section VII concludes.

II A Model of Firm Hiring

In this section, we develop a simple, partial equilibrium model of an individual firm’s hiring strategy to study the implications of integration costs for female employment. Our model is a modified version of Kuhn and Shen (2013). We assume wages and a firm’s candidate pool of potential hires are fixed. The firm must decide which pool of candidates to hire from.

A firm must fill n vacancies, where n is set exogenously. For each vacancy, the firm receives k applications from two types of candidates: type F and type M . Share δ of candidates are type F , and share $(1 - \delta)$ are type M .

Let the net value to the firm of an individual candidate, j , be

$$U_j = v^G + \epsilon_j, \quad G \in (M, F),$$

where ϵ_j is an independent draw from a distribution with cumulative distribution function (CDF) $F(\epsilon_j)$. The difference $v^F - v^M$ embodies between-group differences in expected revenue productivity, wage costs, and turnover. This difference may also reflect employer tastes.

The firm will choose the best worker among candidates it can hire. The question is, which candidate pool will it hire from? At no additional cost, the firm can hire from either the type F or type M pool, but not both. To hire from *both* pools for all vacancies, the firm must pay fixed integration cost c . We assume the firm must choose their hiring strategy prior to observing their candidates. Formally, the firm’s problem is to choose D^F and D^M to maximize

$$nE[\max(U_j; D^F, D^M)] - cD^F D^M,$$

where D^F and D^M are indicators for the firm considering type F and M candidates and $E[\max(U_j; D^F, D^M)]$ gives the expected value of the maximum U_j drawn from the sample of candidates the firm can

hire defined by D^M and D^F . At cost c , integration yields hires of (weakly) higher quality.

To get a simple, closed-form solution, we assume that $F(\epsilon_j)$ is a type I extreme value distribution with scale parameter β .¹⁴ β indexes how much candidate quality varies within group.

The expected value of the highest U_j in a sample of size $s \in \{\delta k, (1 - \delta)k\}$ drawn from a single group, F or M , is

$$U_*^G = \mu^G + \beta \log(s), \quad G \in (M, F),$$

where $\mu^G \equiv v^G + \beta\gamma$ is the expected net value of a single candidate from group G .

The expected value of the highest U_j drawn from a combined sample of all candidates is

$$U_*^I = \beta \log \left[\delta \exp \left(\frac{\mu^F}{\beta} \right) + (1 - \delta) \exp \left(\frac{\mu^M}{\beta} \right) \right] + \beta \log k.$$

The firm's problem of choosing what pools to hire from is equivalent to choosing the maximum of nU_*^F (only type F), nU_*^M (only type M), and $nU_*^I - c$ (both types).

We first consider the choice between hiring only type M candidates and hiring from both types. The firm will pay the fixed integration cost and hire from both types if

$$U_*^I - U_*^M > \frac{c}{n}. \quad (1)$$

The left-hand side of this expression can be expressed as

$$U_*^I - U_*^M = \beta \log \left[\frac{\delta}{1 - \delta} \exp \left(\frac{v^F - v^M}{\beta} \right) + 1 \right]. \quad (2)$$

Let θ denote the probability that the firm's preferred candidate from the combined pool is type F , where

$$\theta = \frac{\delta \exp \left(\frac{v^F}{\beta} \right)}{\delta \exp \left(\frac{v^F}{\beta} \right) + (1 - \delta) \exp \left(\frac{v^M}{\beta} \right)}.$$

Rearranging, we get

$$\frac{1}{1 - \theta} = \frac{\delta}{1 - \delta} \exp \left(\frac{v^F - v^M}{\beta} \right) + 1. \quad (3)$$

Combining (2) and (3), we have

$$\begin{aligned} U_*^I - U_*^M &= \beta \log \left[\frac{\delta}{1 - \delta} \exp \left(\frac{v^F - v^M}{\beta} \right) + 1 \right] \\ &= -\beta \log[1 - \theta] \\ &\approx \beta\theta. \end{aligned}$$

¹⁴The CDF is $F(\epsilon_j) = \exp(\exp(-\epsilon_j/\beta))$. It follows that $\text{Var}(\epsilon_j) = \frac{\beta^2\pi^2}{6}$ and $E(\epsilon_j) = \beta\gamma$, where γ is Euler's constant.

Combining the expression above with (1), an *approximate* condition for the firm to pay the fixed integration cost and hire from the combined pool is

$$n\theta > \frac{c}{\beta}. \quad (4)$$

The left-hand side of (4) is the firm’s expected number of type F hires if it were to integrate. Hence, the firm’s integration decision follows a threshold rule. If $n\theta$ exceeds integration costs (rescaled by β), the firm integrates and hires from both pools. Intuitively, θ is increasing in female labor supply (δ) and $v^F - v^M$, which embodies net productivity of, and employer tastes for, women relative to men.

Next we consider the choice of hiring only from the type F candidate pool. Symmetrically, we have $U_*^I - U_*^F = \beta \log[\theta]$. Hence, if $\theta < \frac{1}{2}$, then $U_*^F < U_*^M$ and no firm will hire from only the type F candidate pool. We find below that in contexts where integration costs are relevant, θ is generally below $\frac{1}{2}$. This is consistent with the fact that all-female firms are rare.

The model is one period but can be readily extended to multiple periods. In a dynamic setting, where n or θ is varying over time, we must distinguish between ongoing and one-time sunk integration costs. Integration decisions now depend on the future path of n and θ and whether integration costs are ongoing or one-time sunk costs. For example, if θ is increasing over time, firms have more incentive to wait to integrate if integration costs are on-going rather than one-time costs.

While framed in terms of hiring, the model also has straightforward implications for female employment. If turnover rates are similar for men and women, then a firm’s female share of hires will equal its female share of employees. Otherwise, the female share of employees will equal the duration-weighted female share of hires.

III Saudi Arabia Context and Data

We apply the model to data to (1) develop a joint test for whether all firms are ex-ante integrated and (2) to estimate counterfactual female employment at segregated firms. We first apply our methodology using administrative data from Saudi Arabia. Saudi Arabia is a useful starting point for two reasons. First, there are strict regulatory and social constraints on between-gender interactions in the workplace, so there are clear reasons to think integration costs may bind in this context. Second, rich administrative data and unique policy variation, which we describe below, provide opportunities to validate our approach. In this section, we describe the Saudi context and data.

III.A Women in the Saudi Workforce

There are several reasons to think that integration costs may be particularly important for Saudi firms. First, Saudi Arabia has extremely low female employment rates by international standards but also has high female unemployment rates. In 2008, before the start of our sample period,

the employment rate for women was 8.4%, and for men it was 56.8% (World Bank, 2016); official unemployment rates were 26.9% for women and 6.8% for men (GaStat, 2011). These patterns are even more pronounced in the private sector, as Saudi women have typically relied on the public sector for work.¹⁵ These disparities are not driven by differences in skill: education levels are also comparable for men and women, and women are more educated among private sector workers and the unemployed.¹⁶ Wages also tend to be significantly lower for Saudi women as compared to Saudi men: in January 2009, our first month of administrative data, the average monthly full-time wage for women at baseline is about half the wage for men (see Appendix Table A3). Even when controls are added for education, location, and occupation, women earn about 40% less than men in January 2009.¹⁷

The Saudi private sector is composed primarily of male, non-Saudi expatriate workers (see Appendix Table A2 for details). Among Saudis, women form a small share of private sector employment, with Saudi women making up 8.5% of Saudi employees in January 2009. Saudi men and women work in similar occupations, though women are more likely than men to work in sales, clerical, and technician roles, and are less likely to work in service and engineering occupations. Occupations with the lowest proportions of Saudi women tend to be those where expatriate workers are concentrated. Among expatriate workers in the private sector, 1%–2% are female. This excludes expatriate domestic workers, of whom 30% are female.

Low female employment in the private sector is likely attributable to a variety of factors on both sides of the market. Female employment in the public sector in part likely reflects women’s work preferences: jobs in education are widely seen as culturally appropriate for women, and completely segregated gender environments are also seen as highly desirable (Evidence for Policy Design, 2015). As we will argue, low female employment in the private sector also reflects significant additional firm-level costs to employing women. At the same time, female employment has become a priority for the Saudi government. The Kingdom’s Vision 2030 economic strategy has an explicit goal of increasing women’s labor force participation to 30% by 2030.

III.B Firm-Level Costs of Employing Women in Saudi Arabia

There are a variety of features of the Saudi labor market that may create additional costs for firms as they begin to hire women. Some of these are specific to Saudi Arabia’s legal requirements around women’s employment: at the time, Saudi Arabia was the only country with legally mandated gender segregation, even in private workplaces. Other costs to firms apply more broadly, including costs associated with accessing a new labor market, attracting female workers, complying with rules around childcare and maternity leave, and complying with other gender-based legal restrictions

¹⁵Even by 2014, women overwhelmingly worked in the public sector, with 74% of employed women working in girls’ schools in 2014 (Evidence for Policy Design, 2015).

¹⁶Unemployed women with college degrees outnumbered men by almost four to one in 2008.

¹⁷A Mincer regression of the log of private sector wages at baseline on employee characteristics indicates that Saudi women earn 40% less than men within occupations after controlling for educational attainment, years of potential experience, and location (all with indicator variables).

around women’s work. Many of these costs are fixed in the sense that they do not depend on the number of female workers that firms employ. These include one-time switching costs as well as ongoing costs that apply to integrated firms. Firms may also face differential per worker, or variable, costs in employing women instead of men.¹⁸

III.B.1 Fixed Costs

The fixed integration costs of employing women are especially striking in the Saudi context. In particular, it may have been costly for firms to comply with government regulations regarding gender segregation in the workplace. During the study period, the government required that a firm employing women provide them separate workstations, a private space to pray and take breaks, convenient restroom access, and a separate entrance to the building or workplace. Meeting rooms also had to be adjusted to accommodate mixed-gender meetings: firms were initially required to hold these meetings only in private and later to make them fully visible to the rest of the office. Employing women exposes firms to inspections and potential fines through the Ministry of Labor and Social Development (MLSD)¹⁹ and the Ministry of Municipal and Rural Affairs (Khoja and Thomas, 2018).²⁰ Even when not legally required, Saudi firms may experience additional costs in providing a workplace environment that is acceptable to female workers.²¹ In addition to the explicit integration costs associated with making a workplace compliant with segregation regulations, the cost of learning how to comply with these rules may also present a barrier to hiring women.

Low historical female employment may also lead to high search costs on both sides of the market: firms may have limited access to hiring and referral networks with female employees, and women may have little information about opportunities for private sector employment. Furthermore, Saudi firms must also develop a strategy for navigating the relationship with male guardians: this is no longer explicitly required by the government, but many firms do ask for guardian permission when recruiting female workers.²² More broadly, firms may also need to develop different types of HR policies to attract and to retain female employees, such as offering parental leave, facilitating childcare, and addressing workplace harassment.²³ Addressing these HR issues involves learning

¹⁸These integration costs are sometimes explicitly cited when discussing obstacles to female employment. One business owner told the *New York Times*, “If they hire women to work, they need another office, with electricity, a dedicated security guard, computers... This is a major cost, especially for small, local companies.” (*New York Times*, 2012) Lubna Olayan, a female Saudi CEO, describes integration obstacles, such as difficulties navigating labor law and social customs, when providing the required segregation for her company’s male and female employees (*Fortune*, 2015).

¹⁹The Ministry of Labor (MoL) was reorganized into the Ministry of Labor and Social Development in 2015.

²⁰In addition to fines for not providing gender-segregated workplaces, a new fine was introduced in October 2015 to penalize female employees individually for not wearing a headscarf in the workplace (Khoja, 2016).

²¹For example, female workers may prefer to have female coworkers, so that female employees require a compensating differential that is declining in the firm’s number of female employees. Though not a fixed cost, this type of integration cost would generate similar firm behavior.

²²The guardianship requirement was lifted by the Ministry of Labor in 2008. There are still 18 countries where women must have a (male) guardian’s permission to get a job (World Bank, 2018).

²³Some of these adjustments are mandated for firms above a particular size: Saudi’s labor law requires firms that employ more than 50 women with at least ten children under age six must provide childcare access, and firms with more than 100 women must provide a childcare center.

by doing, and these costs will be higher for firms that have never recruited women than for firms that already have female employees.

Firms may also need to restructure their task allocations or working hours to accommodate female employees. This type of reassessment can similarly present a one-time hurdle to overcome before hiring women. For example, firms may have a narrow view of the qualifications they require (e.g., certain types of degrees)²⁴ or years of experience, which disqualify many female applicants. Overcoming these barriers may require firms to think flexibly about how they structure their tasks across occupations within the firm. This might include restructuring shifts and working hours, as Saudi Arabia is among the 44 countries that restrict the working hours of women. Firms may also face costs due to the lower mobility of female employees.²⁵ Some firms address this by providing group transportation for their employees, a lumpy, ongoing cost.

III.C Nitaqat Nationalization Quotas

We further test our approach using exogenous variation in Saudi employment across firms generated by the Nitaqat quota policy. The Nitaqat program is an ongoing gender-neutral nationalization quota policy first instituted in 2011.^{26,27} The policy was designed to address growing national unemployment, which in 2011 had reached 40% for Saudis in the 20–25 age group, in the context of the low participation of nationals in the private sector. At the time, foreign guest workers made up 90% of non-oil private sector employment, with the majority of Saudis employed in the public sector. Under Nitaqat, the Saudi government began requiring private sector firms to attain set nationalization quotas for their employees. The Ministry of Labor (MoL) first announced plans for Nitaqat in early 2011, with detailed information about the program structure, targets, and penalties released to firms in June 2011. Sanctions for noncompliance were phased in, starting just three months later in September 2011.

The program first classified firms according to industry and size. Firms were assigned to one of 41 initial industries and to one of five size groups: tiny (< 10 employees), small (10–49 employees), medium (50–499 employees), large (500–2,999 employees), and giant (3,000+ employees). Industry classifications were made using economic activities registered with the Ministry of Commerce; size group classifications were made using the total number of Saudi nationals registered as employees with the General Organization for Social Insurance (GOSI) and foreign nationals with visas sponsored by the company who were registered with the National Information Center.

Within each cell of this industry by size classification, the MoL then defined four different color bands corresponding with a company’s Saudi employee percentage. Firms with fewer than ten employees were not subject to Nitaqat regulations over this period and were classified with a White color band, though they were added in to the regulation after 2012. A large-sized manufacturing firm, for example, faced the following cutoffs for the four color bands:

²⁴Engineering, for example, was not offered to Saudi women as an undergraduate degree program until 2005.

²⁵Women were not permitted to drive in Saudi Arabia until June 2018.

²⁶See Peck (2017) for a more detailed description of the Nitaqat program and its effects.

²⁷Other contemporary labor policies are described in Appendix A.

Red:	0%–7%
Yellow:	7%–19%
Green:	19%–34%
Platinum:	35+%

These cutoffs were set based on pre-Nitaqat Saudization rates so that slightly less than half of firms in each cell would be classified as Green or Platinum, with the intention that the Yellow/Green quota cutoff be attainable for most firms in each cell. The Red/Yellow and Green/Platinum cutoffs were set at the discretion of MoL staff, and the MoL used its visa issuance and foreign recruitment services to enforce the program. These services were tied directly into the monitoring system and were implemented automatically on firms that failed to meet their nationalization quotas.²⁸ Firms in the Green and Platinum bands were given access to a streamlined visa renewal service, while firms in the Red and Yellow bands faced restrictions on their ability to renew existing visas, obtain new visas, and access the MoL’s foreign recruitment services. Sanctions against Red firms phased in more quickly and were slightly more strict than those placed on Yellow firms, but both types of enforcement were disruptive for firms with large numbers of expatriate employees. All sanctions were enforced on both categories of firms by the end of the first year of the program. Platinum firms were given some additional benefits in terms of the ease of their visa renewal but were mostly treated the same as Green firms.

The program dramatically increased the number of Saudis in the private sector (Peck, 2017), with firms complying with the program by increasing their Saudi employment. For our purposes, we categorize firms by their Nitaqat status at the start of the program: “Above Quota” firms are those in the Green and Platinum bands in July 2011, and “Below Quota” are those in the Red and Yellow bands. Nitaqat quotas also served as a way for the government to introduce a de facto minimum wage for Saudis. In September 2012 the government announced that only Saudis paid at least 3,000 SAR per month would count as a full Saudi employee; those paid 1,500 SAR would count as half an employee for Nitaqat purposes, and those between 1,500 and 3,000 SAR would be linearly prorated. This restriction was applied to firms beginning in February 2013.²⁹

Overall, Nitaqat quotas were effective at increasing Saudi employment in the private sector, though it did so at significant cost in terms of firm exit and expatriate employment. Firm-level increases in Saudi employment were closely related to distance from the quota, with below quota firms significantly increasing their Saudi hiring (Peck, 2017).

III.D Data

We test for integration costs in the Saudi context using administrative social security data from GOSI. These data (hereafter referred to as the GOSI data) contain information on all Saudis employed in the private sector between January 2009 and June 2015.³⁰ The data set is used to track

²⁸Color band status in the system was based on a 13-week moving average of the number of Saudi workers registered with GOSI divided by the total number of workers.

²⁹*Saudi Gazette*, September 9, 2012, “Nitaqat percentage linked to salary.”

³⁰Our data only go back to 2009 due to a change in how the data were stored and collected by GOSI. Unfortunately, we are not able to obtain information prior to this year.

Saudi employees for social security eligibility and withdrawal purposes and contains information on worker characteristics such as gender, age, education level, and marital status; job characteristics such as occupation, work location, full-time status, and wages; and firm information such as their administrative identifiers and industries. While we cannot identify establishments in the data, the definition of the firm we use in this paper can be thought to be a legal commercial organization within a particular province or major city. Our definition of firms is described in more detail in Appendix B. In total, the GOSI data set contains information on approximately 2.8 million unique individuals and 430,000 firms.

The data set is structured such that each observation represents an individual’s employment at a firm for a particular wage and occupation. In other words, the observations are at the individual-occupation-wage-firm level, and they contain date ranges that specify when the individual was working in a particular firm in the noted occupation and was paid the corresponding wage. We transform these data into an unbalanced monthly panel for each Saudi employee. We drop entries for part-time work, which only affects about 47,000 of the 2.8 million employees in the data. If an individual has more than one full-time job in a given month, we keep only the observation for the job with the highest wage.

To standardize the occupations in the data (which are based on categorizations by GOSI), and make them more comparable to international classifications, we create a crosswalk between the occupations and the International Labour Organization’s (ILO) 2008 International Standard Classification of Occupations (ISCO-08). We classify each occupation to the two-digit ISCO-08 group, reducing the number of occupations from 2,151 to 40. This significant drop in occupations is primarily due to inconsistent naming, misspellings, and changes to the GOSI classification scheme over time. Appendix Table A4 lists the top ten most common ISCO-08 coded occupations in June 2011.

There are 37 work locations provided in the data. We limit our analysis to locations with at least 50 firms with five or more Saudi employees in January 2009. This leaves us with 17 locations that account for 95% of firms and 98% of workers. In January 2009, 83% of workers are located in four cities: Riyadh, Jeddah, Damman, and Khobar.

We also test our model using firm responses to Nitaqat Saudi employment quotas. We use the Nitaqat data to obtain a list of firms and their quota compliance status for the second week of June 2011, when the program began assessing quotas and began reporting status to firms. This gives us a sample of approximately 1.07 million firms at our baseline, over 990,000 of which were originally exempt from the program for having fewer than ten employees. Approximately 113,000 of these firms appear in the GOSI data.³¹ The details of merging the two data sets are described in more detail in Appendix B.

³¹The big drop in the number of baseline firms between the two data sets is primarily due to the fact that many firms in the white color band do not need to hire any Saudi employees, and therefore they do not appear in the GOSI data since it only contains information on firms that have hired at least one Saudi between 2009 to 2015. Additionally, some firms exit the market before hiring any Saudis, as Peck (2017) documents, so they again would not appear in our GOSI data.

There is a potential concern that GOSI records may not accurately reflect real employment if firms falsify their employee records with GOSI to meet their Nitaqat quotas. This may be a particular concern for female employment if firms are more likely to fraudulently register women’s ID numbers. We discuss this possibility in Appendix D by examining the share of workers in the GOSI data with “active” subsequent career trajectories by month of hire. We find that women hired after Nitaqat are no less likely to have active careers than those hired in the pre-period, particularly when compared to men and when controlling for observable worker characteristics.

Unfortunately, the GOSI data do not include information non-Saudi workers. Our references to the composition of workers throughout the paper refer only to Saudi employees. For example, firms that we identify as “all-male” may in fact employ non-Saudi women. Due to this data limitation, the integration costs we identify are those for hiring *Saudi* women, rather than women more generally. In practice, the integration costs associated with either Saudi or non-Saudi women likely have substantial overlap. Consistent with this, firms that employ non-Saudi women tend to employ Saudi women as well.³²

IV Empirical Strategy and Results

The central ideas of the model are (1) firms face an extensive margin integration decision and (2) integration costs are largely fixed, so firms integrate only if they anticipate employing enough women to justify the costs. To take the model to the data, the central assumption we make is that the probability that the top candidate for position i is female is a function, $\theta(\cdot)$, of *observable* job characteristics, X_i . We assume other factors that determine this probability are *uncorrelated* with the identity of the firm. More formally, we assume that for position i ,

$$P(\text{top candidate for position } i \text{ is female} | X_i) = \theta(X_i) + \epsilon_i,$$

where $E[\epsilon_i | i \in \text{firm } j] = 0$. For ex-ante integrated firms, this characterizes the probability that the hire for a position is female. For ex-ante segregated firms, this characterizes the same *counterfactual* probability if an individual firm were to integrate, holding the behavior of other firms fixed.

Building on this assumption, we develop a joint test for whether integration costs are nonbinding at all firms so that all firms are ex-ante integrated. We then show how to estimate $\theta(X_i)$ when some firms are ex-ante segregated and use this estimate to construct counterfactual female employment at each segregated all-male firm. We also test our central assumption in several ways described below.

³²Among all firms that employed both non-Saudi women and Saudis in 2012 or 2013, 75% employed Saudi women. Among firms that employed both non-Saudis (women or otherwise) and Saudis, only 39% employed Saudi women.

IV.A Testing the Null of No Binding Integration Costs

We first test the null hypothesis that no firm faces binding integration costs and all firms are ex-ante integrated.³³ The distribution of female hires across firms should be consistent with $\theta(X_i)$, the probability that the top candidate for position i is female given job characteristics X_i . In other words, conditional on job characteristics, different firms should hire women at similar rates, and any variation across firms is due to chance alone. Our procedure for testing the null hypothesis is as follows: (1) estimate $\theta(X_i)$, (2) simulate the implied distribution of female hires across firms, and (3) compare that to the distribution we observe in practice. We describe each step in more detail below.

By contrast, if some firms *do* face binding integration costs, we show in Appendix C that the simulation will generally *underpredict* the number of firms with zero female hires. The intuition is that when some firms are in fact ex-ante segregated, female hires are more concentrated across firms than the simulation predicts.

While the model is framed in terms of firm hiring, the test we first develop here uses cross-sectional, firm-level data on female *employment*. We do this in part because firm-level data on current employees are more readily available than data on worker flows.³⁴ If turnover rates are similar for men and women, then a firm’s female share of hires will equal its female share of employees. (In Saudi Arabia, turnover rates are similar for Saudi men and women.³⁵) Otherwise, the female share of employees will equal the duration-weighted female share of hires. We conduct a similar test in Section IV.C.2 that examines hiring rather than employment.

IV.A.1 Estimating $\theta(X_i)$

We estimate $\theta(X_i)$ with a job-level regression model using jobs at all firms meeting our sample criteria. We use cross-sectional data from January 2009, the first month of our data. We limit to firms with at least five Saudi employees to reduce the degree of “chance” segregation.³⁶ While firms with fewer than five Saudi employees account for the majority of firms, they account for less than 10% of Saudi private sector employment. We show in Table II that the industry composition of all firms, and those with at least five Saudi employees, are comparable.

[Table 2 about here.]

We estimate a logistic regression model of the form

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

³³A special case would be that integration costs are zero (i.e., do not exist).

³⁴For example, the World Bank data we use in Section VI to identify integration costs outside of Saudi Arabia includes information on employment but not hiring.

³⁵The monthly turnover rates in the GOSI data are 3.5 and 4.2 percentage points for men and women. Adjusting for job characteristics (occupation, industry, and location) and month, turnover rates are 5% lower for women.

³⁶In other words, this restriction reduces the number of firms that are potentially ex-ante integrated but ex-post segregated.

where X_i includes fixed effects for job location, two-digit occupation, and one-digit industry. We label the function we estimate as $\hat{\theta}^0(X_i)$.

In Table III we summarize $\hat{\theta}^0(X_i)$ across all jobs and the explanatory power of location, occupation, and industry fixed effects for these estimates. The mean is 0.08, the median is 0.027, and the standard deviation is 0.155. Across one-digit occupations, $\hat{\theta}^0(X_i)$ is largest among professionals at 0.23 and lowest among plant and machine operators at 0.007. Across industries, $\hat{\theta}^0(X_i)$ is largest in community and social services at 0.43 and lowest in electricity, gas, and water at 0.008. In separate linear regression models, occupation and industry explain 73% and 62% of the variance in $\hat{\theta}^0(X_i)$, while location explains only 6%. $\hat{\theta}^0(X_i)$ explains 31% of variation in worker sex across positions.

[Table 3 about here.]

IV.A.2 Simulation Results

Next, we simulate the distribution of female employment across firms using our estimate, $\hat{\theta}^0(X_i)$, and compare the result to the distribution we observe. In each simulation, we take a random draw from a uniform distribution for each position i . If that draw is below $\hat{\theta}^0(X_i)$, the worker in that position is labeled as female; if not, the worker is labeled as male. We then sum up to the firm level to get the simulated total of female employees at each firm. We repeat this procedure 1,000 times.

We plot the simulated and observed distributions of female employment in Figure I. We plot the share of firms with zero female employees separately due to the difference in scale. We also plot the share of firms with 1, 2, 3, 4, 5, 6–10, 11–25, and > 25 female employees. The error bars represent the 5th and 95th percentiles across simulations for the share of firms with a given number of female employees.

[Figure 1 about here.]

We substantially underpredict the number of firms with zero female employees. While we predict that 43% of firms will have zero female employees, on average, across simulations, in fact, 73% of firms have zero female employees. We also overpredict the number of firms with few female employees, particularly in the one to four range. For all simulations, we reject equality of the distributions in a Kolmogorov-Smirnov test at $\alpha = 0.01$. Overall, the pattern is consistent with binding integration costs at many firms.

IV.A.3 What if $\theta(\cdot)$ Is Misspecified?

However, our simulated distribution may also fail to match the observed distribution because we have misspecified $\theta(X_i)$. This would be true if there are job characteristics that are not included in X_i that (1) help to explain the probability that the top candidate for a position is female and (2) vary systematically across firms, conditional on X_i .

One concern is that the occupation and industry classifications in our data may be too coarse, as there may be systematic variation in sex composition between subcategories of jobs. For example, for the same occupation, commerce firms that sell men’s clothing may skew male compared to commerce firms that sell women’s clothing. Under this misspecification, our simulation may underpredict the number of firms with zero female employees, not because some firms have not ex-ante integrated but because some firms in fact have smaller θ values than we estimated. In other words, we may underestimate the number of firms that are all male simply because those firms employ workers in job types that few women work in.

There is reason to think misspecification is not a first-order issue. First, generating the number of all-male firms we observe would require a substantial role for unobservable job characteristics relative to observable characteristics in determining θ . To see this, note that if we had estimated θ using no covariates so that $\hat{\theta}^0 = 0.08$ for all jobs, we would simulate that 34% of firms would be all male. Adding our observable job characteristics increases this value to 43%, but brings the value nowhere close to the observed value, 73%. Unobservable job characteristics would need to be very predictive relative to observable job characteristics to match the distribution of female employment in the data.³⁷

Second, the job characteristics we use—location, two-digit occupation, and one-digit industry—explain much of the systematic variation in sex composition across jobs. For comparison, we use the 2015 American Community Survey to measure the additional predictive power of job characteristics included in those data (e.g., four-digit industry and occupation codes) for predicting the sex of a worker. In those data, job characteristics analogous to what we use in the Saudi data explain 28% of variation in sex across workers. Location and four-digit occupation and industry explain 36% of the variation, only a modest increase.

Still, we test our specification of $\theta(X_i)$ in several ways below, including a placebo test described in the next section and more direct tests described in Sections IV.C and IV.D.

IV.A.4 Placebo Test: Married versus Unmarried Men

As a placebo test for the simulation procedure, we adopt the same approach but test a different null hypothesis: no firm faces binding integration costs, where integration refers to employing both *married* and *unmarried* men. We have no reason to think such an integration cost exists. Hence, restricting to male hires, if the probability of a married hire for vacancy i is only a function of observable job characteristics, we should be able to match the observed distribution of married male employment across firms.

Across jobs, we find that the estimated probability that a worker is married has a mean of 0.246 and varies significantly, with a standard deviation of 0.093. Appendix Figure A1 plots the simulated and observed distributions of married male employment. Reassuringly, the simulation matches the observed distribution well. For all simulations, we fail to reject equality of the distributions in a Kolmogorov-Smirnov test at $\alpha = 0.01$. In a context where integration costs do not exist, our

³⁷The reasoning behind this argument is similar to that of Altonji et al. (2005) and Oster (2019).

methodology can match the observed distribution of employment.

IV.B Estimating $\theta(X_i)$ When Integration Costs Bind

If some firms are ex-ante segregated, $\hat{\theta}^0(X_i)$ will underestimate $\theta(X_i)$ because we include these firms in its estimation. To correctly estimate $\theta(X_i)$, we must limit the data to ex-ante integrated firms.

A key problem with executing this is that we *do not observe* whether firms are *ex-ante* integrated. Instead, we observe whether they are “*ex-post*” integrated—whether they employ both men and women in practice. We take two approaches to address this issue. First, we ignore the distinction and limit the data to ex-post integrated firms. This will lead to an upward bias in our estimate for $\theta(X_i)$.³⁸ Second, we estimate a more parametric model that accommodates ex-ante integration as a potentially unobserved firm state.

Equipped with an estimate of $\theta(X_i)$, we can use the following to estimate counterfactual female employment for firms that did not integrate:

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

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. In other words, once we know the probability that the top candidate for a given job is female, we can predict female employment for each firm under integration given its job mix.³⁹ We can also test whether firm ex-ante integration rates are increasing in $\bar{\theta}_j n_j$, as our assumption that integration costs are largely fixed would suggest.

IV.B.1 Using Ex-Post Integrated Firms to Estimate $\theta(X_i)$

We first limit the data to ex-post integrated firms when estimating $\theta(X_i)$. This approach is straightforward and transparent, though it will generate an upward bias because we exclude ex-ante integrated firms that, by chance alone, do not hire any women.

Table II summarizes the characteristics of ex-post integrated firms, using the same sample restrictions described in Section IV.A.1. Overall, the female share of employment is 8.2%, while the female share of employment at integrated firms is 12.5%. We label our function estimated using only ex-post integrated as $\hat{\theta}^{EP}(X_i)$.

Column (2) of Table III summarizes $\hat{\theta}^{EP}(X_i)$ for *all* jobs, not just those at ex-post integrated firms. In separate linear regression models, occupation and industry explain 69% and 60% of the variance in $\theta(X_i)$, while location explains only 9%.

³⁸Intuitively, the bias is small because the distinction between ex-ante and ex-post integrated firms is only relevant for small firms, which account for a small share of employment.

³⁹For this construction, we assume that a firm’s job mix does not depend on its integration status.

IV.B.2 Modeling Firm Integration States

The second approach we take is to directly model the distinction between ex-ante and ex-post integrated firms and to structurally estimate $\theta(X_i)$.

Let j index firms, and let N_j denote the number of positions at firm j . Let y_{ij} be an indicator that equals one if position i in firm j is filled by a female employee. Denote $K_j = \sum_i^{N_j} y_{ij}$ as the number of female employees at firm j .

Let π_j denote the probability that firm j has *not* paid its integration cost and so is not able to employ women. Hence, with probability $1 - \pi_j$, the firm is ex-ante integrated. We will model π_j as a function of observable firm characteristics. Finally, among ex-ante integrated firms, denote the probability that position i is filled by a female employee as θ_{ij} . As above, we model θ_{ij} as a function of observable job characteristics, X_{ij} .⁴⁰

With these terms defined, we can define the likelihood function for each firm. Without loss of generality, we order each firm's workers such that the first K_j workers are female and the remaining $N_j - K_j$ are male. Denote $Y_j = (Y_{1j}, \dots, Y_{N_jj})$ as the firm-specific vector of outcomes. The likelihood function for firm j is

$$P(Y_j = Y) = \begin{cases} \pi_j + (1 - \pi_j) \prod_{i=1}^{N_j} (1 - \theta_{ij}) & \text{if } K_j = 0 \\ (1 - \pi_j) \prod_{i=1}^{K_j} \theta_{ij} \prod_{i=K_j+1}^{N_j} (1 - \theta_{ij}) & \text{if } 0 < K_j < N_j \\ (1 - \pi_j) \prod_{i=1}^{N_j} \theta_{ij} & \text{if } K_j = N_j. \end{cases}$$

We model both θ_{ij} and π_j in logistic regression models with explanatory variables X_{ij} and Z_j , respectively:

$$\begin{aligned} \theta_{ij} &= \Lambda(X_{ij}\beta) \\ \pi_j &= \Lambda(Z_j\gamma) \end{aligned}$$

where Λ is the logistic function. In the vector of firm characteristics, Z_j , we include fixed effects for location and industry and a cubic in log firm size.⁴¹ For the vector of hire characteristics, X_{ij} , we include fixed effects for two-digit occupation codes, location, and one-digit industry. We estimate the model using an expectation-maximization (EM) algorithm. Estimation details are provided in Appendix C. We label these structural estimates for $\theta(X_i)$ as $\hat{\theta}^S(X_i)$.

Column (3) of Table III summarizes the estimates and how they vary across jobs. The average value of $\hat{\theta}^S(X_i)$ is 0.123. These estimates are similar to those from Section IV.B.1 using only ex-post integrated firms; the correlation between $\hat{\theta}^S(X_i)$ and $\hat{\theta}^{EP}(X_i)$ is 0.82. The average value of

⁴⁰For ease of notation, in this section we index positions separately by firm.

⁴¹We measure firm size here using the firm's number of Saudi employees.

π_j is 0.65, indicating 65% of firms are ex-ante segregated.

IV.B.3 Applying Estimates of $\theta(X_i)$

We next examine how integration rates relate to $\bar{\theta}_j n_j$, a firm's expected number of female employees if ex-ante integrated. Panel A of Figure II plots ex-post integration rates as a function of $\bar{\theta}_j n_j$. It also plots simulated ex-post integration rates under the counterfactual that all firms are ex-ante integrated. The distinction between ex-ante and ex-post integration is that firms that pay their integration costs may still hire only men by chance alone. The simulated ex-post integration rate in Panel A of Figure II estimates this chance factor.

[Figure 2 about here.]

We use the relationships illustrated in Panel A of Figure II to estimate *ex-ante* integration rates as a function of $\bar{\theta}_j n_j$. To see how, let I_j be an indicator for whether firm j has ex-ante integrated. Then

$$P(K_j > 0) = P(K_j > 0 | I_j = 1) \times P(I_j = 1),$$

where $P(K_j > 0)$ is the probability that firm j is *ex-post* integrated. Grouping firms by their value of $\bar{\theta}_j n_j$, we get

$$\begin{aligned} P(K_j > 0 | \bar{\theta}_j n_j) &= E[P(K_j > 0 | I_j = 1) \times P(I_j = 1) | \bar{\theta}_j n_j] \\ &\approx P(K_j > 0 | I_j = 1, \bar{\theta}_j n_j) \times P(I_j = 1 | \bar{\theta}_j n_j), \end{aligned}$$

where the approximation holds because conditional on $\bar{\theta}_j n_j$, $P(K_j > 0 | I_j = 1)$ varies little across firms.

The two relationships depicted in Panel A of Figure II correspond to $P(K_j > 0 | \bar{\theta}_j^{EP} n_j)$ and $P(K_j > 0 | I_j = 1, \bar{\theta}_j^{EP} n_j)$, where $\bar{\theta}_j^{EP}$ is the natural estimate for $\bar{\theta}_j$ constructed using $\hat{\theta}^{EP}(X_i)$. Hence, the ratio of the actual and simulated ex-post integration rates provides an estimate of the ex-ante integration rate, $P(I_j = 1)$, as a function of $\bar{\theta}_j n_j$. Panel B of Figure II plots this estimate of ex-ante integration rates as a function of $\bar{\theta}_j^{EP} n_j$. We also plot our structural estimates of ex-ante integration rates, π_j , as a function of $\bar{\theta}_j^S n_j$, where $\bar{\theta}_j^S$ is the estimate for $\bar{\theta}_j$ constructed using $\hat{\theta}^S(X_i)$.

For both estimates, we find that ex-ante integration rates are increasing in $\bar{\theta}_j n_j$. This pattern is consistent with firms facing an integration threshold rule with respect to $\bar{\theta}_j n_j$.

IV.B.4 Can $\theta(X_i)$ Match the Distribution of Female Employment?

As an additional test for whether $\theta(X_i)$ is well specified, we evaluate whether a simulation of the distribution of female employment across firms that allows for integration rates to vary by $\bar{\theta}_j n_j$ fits the observed distribution. For each firm, we take a uniform random draw and label the firm as integrated if the draw is below the corresponding values in Panel B of Figure II given the firm's value of $\bar{\theta}_j^{EP} n_j$. If the firm is not labeled as integrated, we assign it a value of zero for its female

employment. For firms labeled as integrated, we simulate a value of female employment as above, this time using $\hat{\theta}^{EP}(X_i)$ to assign the gender for the employee in each position.

Figure III compares the simulated distribution of female employment to the observed distribution. While, by construction, we will match the share of firms with zero female employees, the simulation is not guaranteed to match other parts of the distribution. In fact, it does match quite well. For all simulations, we cannot reject equality of the simulated and observed distributions in a Kolmogorov-Smirnov test. This suggests that we have included the most relevant job characteristics in X_i (or that other relevant characteristics are not concentrated within firms) and have mapped them appropriately to hiring probabilities, at least among ex-ante integrated firms.

[Figure 3 about here.]

IV.C Using Panel Structure to Validate the Model

So far we have only used the January 2009 cross-section to test for binding integration costs and to measure counterfactual female employment at ex-ante segregated firms. The patterns we document imply that a majority of Saudi firms face binding integration costs. We will apply the approach outlined here to cross-sectional data on manufacturing firms in 65 countries in Section VI.

Given that cross-sectional firm data are typically more available than panel firm data, it is useful that we can apply the model using cross-sectional data. However, the panel structure of the Saudi data is valuable because it permits additional tests of the model. In particular, panel data allow us to further probe our key assumption that $\theta(X_i)$ dictates counterfactual female employment for ex-ante segregated firms.

We use the panel structure to conduct two tests. First, we test whether our estimate of $\theta(X_i)$ provides unbiased predictions for the female share of hires at *newly integrated* firms. With panel data, we can observe firm transitions from segregated to integrated. The model predicts an *extensive margin* adjustment: abrupt changes occur in the gender composition of hires for these firms as they move from hiring no women to hiring women at a rate dictated by their job composition. By contrast, if the bunching at zero we observe in Figure I is driven by unobserved heterogeneity in job characteristics, we would expect these transitions to reflect *intensive margin* changes in θ or chance variation in the candidate pool. In this case, we expect transitions to be smooth and the female share of hires at newly integrated firms to be low relative to observably comparable incumbent integrated firms.

Second, we test for *state dependence*, comparing hiring behavior at previously segregated and previously integrated firms. If previously integrated firms tend to maintain their integrated status, we should see that hiring patterns in those firms are consistent with no binding integration costs.

IV.C.1 Checking Predictions for Counterfactual Female Employment

First, we test whether our estimate of $\theta(X_i)$ provides unbiased predictions for the female share of hires at newly integrated firms. This is a powerful out-of-sample test for whether our estimate of

$\theta(X_i)$ predicts counterfactual female employment at segregated firms because we do not use this set of firms to estimate $\theta(X_i)$. We also examine the transition dynamics of these newly integrated firms.

We first examine hiring at newly integrated firms in an event study. We plot the female share of hires at integrating firms in the months following a firm’s first observed female hire.⁴² We limit to firms with at least five Saudi employees in the month prior to integration. We observe 8,307 transitioning firms meeting this size threshold. Prior to integrating, we observe transitioning firms in the GOSI data for an average of 28 months. At each firm, we calculate the female share of hires in six-month increments before and after a firm’s first female hire. We then take the average across all firms meeting the sample restrictions and exclude firms that do not make a hire in a given six-month increment from the calculation for that period.

The event study is shown in Panel A of Figure IV. By construction, among hires made prior to integration, there are no women. Among firms that we observe integrating, we observe an average of 25 male hires made over 24 months prior to a firm’s first female hire. We see that the female share of hires changes abruptly at newly integrated firms, consistent with an extensive margin response.⁴³ Among hires made in the six months following integration, including the first female hire, 55% are female. This drops to about 26% in the following six-month period and remains relatively steady thereafter. By contrast, if we thought the excess mass of firms with zero female workers we observed in Figure I was driven by unobserved heterogeneity in job characteristics, we would expect a gradual and potentially short-lived increase in female hiring rather than the discrete and sustained increase we observe.

[Figure 4 about here.]

We compare the observed increase to what we would predict using an estimate of $\theta(X_i)$ derived from hires at incumbent integrated firms. As in Section IV.B.1, we construct our predictions by estimating a logistic regression of the form:

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

where i indexes the position for each hire. As above, X_i includes fixed effects for job location, two-digit occupation, and one-digit industry. In addition, we allow predictions to vary over time by including in X_i fixed effects for each half-year and interactions between the location, occupation, and industry controls with an indicator for hires made after June 2011, the month Nitaqat is implemented. We limit estimation to all firms that are ex-post integrated in January 2009, regardless of whether any of their subsequent hires are female. These firms should provide a valid estimate for $\theta(X_i)$ if integrated firms remain ex-ante integrated moving forward, an assumption we verify in the next section. We simply label this estimate $\hat{\theta}(X_i)$.

⁴²In this exercise, we exclude firms that have female employees when they are first observed in the GOSI data.

⁴³Note that the period labeled as “0 to 5” months includes the first female hire herself.

We include the $\hat{\theta}(X_i)$ -based prediction for the female share of hires in Panel A of Figure IV. We find that the magnitude of this change matches what we would predict using $\hat{\theta}(X_i)$, at least on average. Next, we check how well firm-specific predictions for the female share of hires of newly integrated firms matches the realized female share of hires. In Panel B of Figure IV, we group newly integrated firms into deciles based on the predicted female share of hires and plot bin averages against their observed female share of hires 12 or more months following their first female hire. If the predictions are *unbiased*, the binned averages will fall on the 45-degree line. This is similar to the pattern we observe, though the observed female share of hires is slightly below the 45-degree line, with the gap increasing in the predicted female share of hires.

IV.C.2 State Dependence

An immediate implication of the model is state dependence: the hiring behavior of firms that have already paid their integration costs will differ from the behavior of firms that have not. In particular, we should *not* observe bunching for the former set of firms. While we cannot observe each firm’s current state, we proxy for their current state using their baseline ex-post segregation status. This proxy should closely correlate with a firm’s current state if integration costs are sunk or if the conditions that led the firm to integrate are highly persistent over time. We test the null hypothesis of no binding integration costs but conduct separate tests for firms that are ex-post integrated and ex-post segregated as of January 2009.

We conduct a similar test to that described in Section IV.A, except we pool *hires* between February 2009 and June 2015. We limit to firms that have at least five Saudi hires over this period. To classify firms as ex-post integrated or segregated in January 2009, we also limit to firms that had Saudi employees in January 2009. We estimate $\theta(X_i)$ separately by baseline integration status and include the same job characteristics we use in Section IV.C.1: fixed effects for each half-year and fixed effects for job location, two-digit occupation, and one-digit industry, all interacted with an indicator for hires made after June 2011.

Table IV compares the two sets of firms. There are 2,796 firms meeting the sample criteria that were ex-post integrated in January 2009 (“baseline integrated”) and 12,617 firms that were ex-post segregated (“baseline all male”). Baseline integrated firms are larger, pay higher wages, and concentrated in community and Social services. For baseline all-male and integrated firms, the female share of recent hires is 19.2% and 48.4%. Figure V and VI plot the simulated and observed distribution of female employment for baseline all-male and integrated firms.

[Table 4 about here.]

[Figure 5 about here.]

For baseline all-male firms, the pattern is similar to that observed in Figure I. The simulations underpredict the number of firms that employ zero female workers (16.1% versus 34.2%) and overpredict the number of firms that employ fewer than ten (68.1% versus 51.2%). For all simulations,

a Kolmogorov-Smirnov test rejects equality of the observed and simulated distributions at the 0.01 significance level.

[Figure 6 about here.]

By contrast, the simulated distribution for baseline integrated firms matches the observed distribution relatively well. For all simulations, a Kolmogorov-Smirnov test fails to reject equality of the observed and simulated distributions at the 0.01 significance level.

Consistent with our interpretation of bunching as evidence for the presence of ex-ante segregated firms, there is little evidence of bunching at firms that are likely ex-ante integrated.

IV.D Using Policy Variation to Verify Threshold Rule

We next test whether the Nitaqat employment quotas induce firm integration and increase female hiring in a manner consistent with the model. In particular, Nitaqat provides a direct test for the model prediction that firm integration decisions follow a threshold rule in $\bar{\theta}_j n_j$. As we will demonstrate, Nitaqat incentivizes some firms to increase their number of Saudi hires (n) and incentivizes larger increases at some firms than others. An implication of the model is that by increasing n , Nitaqat will induce some firms to integrate and will increase their female share of hires by a magnitude predicted by $\bar{\theta}_j$.

We study the causal effects of Nitaqat using a difference-in-difference research design, comparing Above and Below Quota firms before and after the policy is implemented. We first show that Nitaqat increases Saudi employment at private sector firms, with larger increases at firms that needed to increase their Saudi share of employees to satisfy their quota.

[Table 5 about here.]

Figure VII compares the growth in Saudi employment at Below Quota and Above Quota firms. Saudi employment is expressed as a percentage change relative to May 2011, the month before the implementation of Nitaqat. Prior to Nitaqat, Saudi employment at Below and Above Quota firms move in tandem. Following the implementation of Nitaqat, hiring at Below Quota firms increases sharply relative to Above Quota firms. By June 2015, the last month of our data, Saudi employment has increased by about 50% at Above Quota firms while doubling at Below Quota firms. We will use this policy-induced variation in firm size to test the predictions of the model.

[Figure 7 about here.]

In particular, we test whether Nitaqat increases the following at Below Quota firms relative to Above Quota firms: (1) Saudi hiring, (2) integration rates, and (3) female share of hires. We also evaluate whether the female share of hires increases by a degree consistent with $\theta(X_i)$. For the analysis below, we limit to Above and Below firms that were ex-post segregated and employed at

least five Saudis in January 2009. Each plot as an average across firms, where each firm that is present in that period is weighted equally.

Table V compares descriptive statistics for these two sets of firms as of June 2011. There are 2,224 Below Quota firms and 1,559 Above Quota firms satisfying our sample criteria. The firms are generally similar except, as expected, Above Quota firms have more Saudi employees (an average of 45.5) than Below Quota firms (33.7). Above Quota firms also pay higher average wages.

We first look at Saudi hiring. Figure VIII plots the average number of hires each half-year, separately for Below and Above Quota firms. Consistent with Figure VII, we see that a gap emerges between Below and Above Quota firms at the implementation of Nitaqat. While the number of hires is comparable at the two sets of firms prior to Nitaqat, a gap of five Saudi hires per half-year emerges in the second half of 2011, which drops to about two hires by the second half of 2012 and stagnates thereafter.

[Figure 8 about here.]

Second, we look at integration rates. Our model predicts that by increasing n at Below Quota firms relative to Above Quota firms, Nitaqat will increase relative integration rates at Below Quota firms too, as more firms will cross their integration threshold. We plot integration rates in Panel A of Figure IX. As the model predicts, the share of Below Quota firms that are integrated increases relative to the same share of Above Quota firms following the implementation of Nitaqat. An immediate difference of about 8 percentage points emerges by the second half of 2011. The gap fluctuates between 8 and 14 percentage points thereafter.

[Figure 9 about here.]

Third, we look at the female share of hires. With constant underlying rates of female hiring, Nitaqat could increase integration rates at Below Quota versus Above Quota firms by chance alone. By contrast, the model predicts an *increase* in the female share of hires at Below Quota versus Above Quota firms. We plot the female share of hires in Panel B of Figure IX. As the model predicts, the female share of hires at Below Quota firms increases relative to Above Quota firms following the implementation of Nitaqat. The magnitude of this relative increase—2–4 percentage points—is in line with what we would predict given the differential in integration rates (about 11 percentage points) and our $\hat{\theta}(X_i)$ estimates from Section IV.C.1. Averaging across post-Nitaqat hires within firms, the average estimated value of $\bar{\theta}_j$ is 0.25.

Finally, we present corresponding difference-in-difference estimates in table form. We estimate models of the form:

$$Y_{jt} = \alpha_j + \tau_t + \beta \text{Post}_t \times \text{Below}_j + \epsilon_{jt}, \quad (5)$$

where α_j are firm fixed effects, τ_t are half-year fixed effects, Post_t is an indicator for post-Nitaqat implementation, and Below_j is an indicator for Below Quota firm. The coefficient β is the post-Nitaqat differential change in the outcome for Below Quota firms relative to Above Quota firms.

We estimate equation (5) for the same three outcomes: total hires, integration status, and female share of hires.

The estimates are presented in Table VI. Over the full period, Saudi hires increase at Below Quota firms relative to Above Quota firms by about three per half-year. The integration rate increases by about 11 percentage points, and the female share of hires increases by 2.26 percentage points.

[Table 6 about here.]

V Aggregate Effects of Integration Costs: Evidence from Nitaqat

We find that integration costs bind for a substantial share of firms in Saudi Arabia. Individually, these firms would employ more women if they were to integrate, as they would do if their integration costs were eliminated. However, what would happen in the aggregate if integration costs were eliminated? Would aggregate demand for female labor increase? How would the female share of the workforce and gender differences in wages change? As in Becker (1957), integrated firms may be sufficiently numerous or large to absorb female labor so that the existence of constrained male-only firms has no bearing on female wages and employment. On the other hand, in the presence of search frictions or insufficient entry or growth of integrated firms, integration costs will reduce aggregate demand for female labor.

To assess the aggregate consequences of integration costs, we would ideally work with exogenous variation in integration costs across labor markets. Lacking such variation, we take a different approach. We examine the labor market response to a policy that reduces the share of firms that face binding integration costs and assess the effects of the policy on female employment and the gender wage gap. If the policy increases female employment or relative wages, this would suggest that the presence of integration costs depresses those outcomes. The logic of our approach is to essentially use the policy as an instrument for (binding) integration costs. The exclusion restriction implicit in our argument is that the policy only affects our outcomes of interest by reducing the set of firms with binding integration costs. We discuss this exclusion restriction below.

In particular, we investigate the aggregate effects of Nitaqat. As discussed in Section IV.D, Nitaqat induced many firms to integrate and hire women for the first time. We show this directly in Figure X, which plots the share of firms that employ both Saudi men and women over time, restricting to firms with at least five Saudi employees. There is a clear trend break that begins just as Nitaqat is implemented, which coincides with the first vertical marker. The share of firms that are ex-post integrated increases from about 30% in the first quarter of 2011 to about 67% in the third quarter of 2013 and remains roughly flat thereafter. This flattening occurs soon after a dramatic increase in the effective minimum wage for Saudis in the private sector, which is implemented in February 2013. We discuss the effects of this minimum wage increase in more detail below.

Next, we explore how the female share of the workforce evolves in response to Nitaqat. In Section IV.D we document that among firms that are all male in January 2009, Below Quota firms

increase their female share of hires relative to Above Quota firms. Figure X plots the female share of Saudis in the private sector over time, pooling employment at all firms. The overall pattern matches that of integration rates. Nitaqat led to a dramatic increase in the female share of Saudis in the private sector, from 10% in 2011 to 27% in 2015. This increase occurs primarily within sectors, as measured by industry and occupation. This increase in the female share of Saudis in the private sector is not offset by a decrease in the public sector; in fact, the female share of Saudis working in the public sector also increases over this period, from 33% in 2011 to 40% in 2015.⁴⁴

[Figure 10 about here.]

While this increase is striking, it does not necessarily indicate an important role for integration costs. Nitaqat may also increase the female share of employment through a price effect. If male labor is relatively inelastic, then Nitaqat may bid up men's relative wages, increasing relative relative demand for women. This would be a violation of our exclusion restriction: a path through which Nitaqat increases the female share of employment that is unrelated to integration costs per se.

However, the evidence suggests that the scarcity of male labor is not the driving force behind the dramatic increase in the female share of the workforce. In fact, the gender wage gap *decreases* following Nitaqat. Moreover, after the effective minimum wage reduces the wage gap even further, the female share of the workforce remains elevated. This is illustrated in Figure XI, which plots the female-male wage gap over time. The figure includes two measures of the gender wage gap: (1) the raw difference in average log wages for women and men and (2) the gap within labor market entry cohorts.

Prior to Nitaqat, the wage gap is relatively flat. The raw wage gap is 60 log points; within cohorts, the gap is about 35 log points. Following Nitaqat, but prior to the minimum wage increase, the wage gap decreases by about 10 log points. The 2013 minimum wage increase leads to a substantial reduction in the gender wage gap. Following its introduction, about 65% of women and 40% of men earn the new minimum wage. The raw wage gap drops to about 30 log points. Within cohorts, the wage gap drops to 4–9 log points. Yet, from Figure X, we can see that the female share of the private sector workforce is *increasing* over this period. This share stagnates beginning in 2013 but remains elevated thereafter. The fact that both female relative wages and employment increase is difficult to reconcile with a price-based explanation. Instead, the evidence is consistent with Nitaqat increasing relative demand for female labor by increasing the set of firms that integrate.

[Figure 11 about here.]

Finally, we explore the possibility that Nitaqat led to a shift in women's labor *supply*. Female labor force participation in Saudi Arabia is among the lowest in the world, at 17.8% in 2011 (GaStat, 2011). While this low rate is likely driven by multiple factors, one may be that households perceive

⁴⁴Source: Ministry of Civil Service. This figure does not include employees in the security and military sectors.

that few firms are willing to hire women in the first place. Nitaqat may cause an outward shift in women’s labor supply by increasing the set of firms that are ex-ante integrated. In fact, integration costs as a barrier to women’s employment may generate feedback effects: women may only enter the labor market if enough firms have integrated, while firms only integrate if they can anticipate employing enough women to justify the costs of integration.

Unfortunately, we do not have data on labor supply decisions; in particular, we do not have data on anyone that is not employed in the private sector. Instead, we look at the response to Nitaqat for firms that had integrated prior to the policy’s implementation. While Figure X shows that the female share of the workforce is increasing, we expect this increase to be concentrated at firms induced to integrate by the policy. Firms that integrated prior to Nitaqat are already employing a mix of men and women and face an increase in the relative price of women. In the absence of a supply response, we would expect to see the female share of employment at these firms weakly decreasing.

Figure XII plots the female share of employment over time in firms that employed Saudis in January 2009, split by the firm’s ex-post integration status in that month. For *both* sets of firms, there is a marked increase in the female share of employment beginning with Nitaqat’s integration. As expected, the increase is larger for baseline-segregated firms. But for baseline integrated firms, the increase is also substantial: from 14.5% in Q1 of 2011 to 20.6% of Q1 2015.

[Figure 12 about here.]

VI Assessing Integration Costs Across Countries

We now apply our methodology to a broader range of countries using survey data on manufacturing firms from the World Bank. Although Saudi Arabia has regulations that explicitly impose integration costs on firms, strong social norms in other countries may implicitly impose similar constraints. In particular, social norms for gender segregation are common in MENA and South Asia (Jayachandran, 2015), and firms may need to restructure the workplace to accommodate social preferences. For example, in interviews with Pakistani men conducted by the World Bank, most expressed reservations about the appropriateness of women interacting with men outside the family, with one business owner quoted as saying, “free mixing of the sexes is encouraged in modern workplaces, but in our culture, this may lead to social problems” (Amir and Durrani, 2019).

As in Saudi Arabia, firms in markets with low baseline female employment may also face particular challenges in recruiting and retaining female employees. Firms may also need to restructure tasks and occupations to facilitate female employment. The International Finance Corporation (IFC) cites the case of an automotive parts plant in Thailand that reduced the weight of boxes in one assembly line to make the work more physically manageable for their female workers (International Finance Corporation, 2013). This type of reorganization can also address restrictions on work shifts: 44 countries, including Saudi Arabia, restrict the working hours of women. In India, for example, women are prohibited from working after 7:00 p.m. The IFC also describes the

case of a chemicals company that identified positions that could be redesigned to fit within these scheduling limitations to accommodate female workers (International Finance Corporation, 2013).

Although particularly acute in Saudi Arabia, the issue of mobility is a common one for women globally. There are 16 other countries where women are legally less able to travel outside the home than men (World Bank, 2018). Even in countries where women are not legally restricted in their mobility, in practice, transportation is a serious barrier to employment for women around the world. Cultural norms and security concerns greatly limit women’s mobility in Pakistan (Field and Vyborny, 2016) and Indonesia (Schaner and Das, 2016), for example. These restrictions can prevent women from taking jobs, lower their after-transport wages relative to men, or make them less reliable employees. Many firms address these problems by offering private transport services for their employees (International Finance Corporation, 2013).

VI.A Data

In this section we review the two data sets we use to assess integration costs and measure preferences for gender segregation across countries: World Bank Enterprise Surveys and Arab Barometer.

VI.A.1 World Bank Enterprise Surveys

To assess integration costs outside of Saudi Arabia, we use publicly available cross-country survey data from the World Bank Enterprise Surveys, which represent 139 countries from various points between 2005 and 2018.⁴⁵ The enterprise surveys are conducted on a representative sample of each country’s private sector, covering businesses’ experiences with access to finance, access to government services, their performance, and other stated constraints. Responses are gathered from business owners and managers, while accountants and HR staff are called in to discuss sales and labor issues. While surveys are conducted for each country, responses are standardized to allow for cross-country comparisons.

The surveys are primarily focused on the manufacturing and services sectors with only private companies in the sample. We limit our samples to manufacturing firms, where surveys include questions on the gender composition of employees by occupation. Unfortunately, the occupation specificity in these surveys are limited to whether workers are employed in production or nonproduction positions. Next, we drop surveys where information on gender composition is missing for more than 20% of firms. In remaining surveys, we drop firms with missing data on gender composition or with fewer than five employees. We then drop surveys with fewer than 100 remaining firms. This leaves us with 105 surveys in 65 countries. Appendix Table A1 lists the surveys (country by year pairs) we include in our analysis.

Survey participants are sampled by stratified random sampling with firm size categories, sectors, and within-country geographic regions as the strata. Larger companies are oversampled given the relatively few large companies in the sampling frame. We use the provided survey weights so that

⁴⁵The data are accessible through <https://www.enterprisesurveys.org>.

firms are representative of the manufacturing sector in a given country and year.

VI.A.2 Arab Barometer

We link our measure of integration costs to social preferences using data from the Arab Barometer survey. The Arab Barometer is a cultural, religious, and political opinion survey that is run periodically across several MENA countries, with four waves of survey data beginning in 2006. The survey is designed to cover a representative sample of adults within each country, with at least 1,000 responses in each of waves II–IV for each country. The Arab Barometer is unique in recording gender segregation preferences, with all four survey groups asking about agreement or disagreement with a statement regarding gender mixing in university classes. In waves I (2006–2009) and IV (2016–2017), the statement is “It is acceptable in Islam for male and female university students to attend classes together,” and in waves II (2010–2011) and III (2012–2014), the statement is “Gender-mixed education should be allowed in universities.” We use responses to these statements to measure the share of respondents in each country that approve of gender mixing. We match this information to data on surveyed countries in the World Bank Enterprise Survey. The eight matched observations are Egypt (waves II–IV), Iraq (II and III), Jordan (I–IV), Lebanon (I–IV), Morocco (I, III, IV), the West Bank and Gaza (I–IV), Tunisia (II–IV), and Yemen (I–III).

VI.B Results

VI.B.1 Testing the Null of No Binding Integration Costs

We begin by testing the null hypothesis that integration costs are nonbinding for all firms, separately by region. We follow the procedure described in Section IV.A. We calculate a separate $\theta(X_i)$ function for each survey (country by year pair) 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 $\hat{\theta}^0(X_i)$ and compare that to the observed distribution. Note that because our occupational measure is quite coarse, it is more likely that $\theta(X_i)$ is misspecified in an economically significant, so even under the null hypothesis, we may have trouble matching the observed distribution of female employment.

Figure XIII compares the distributions for the largest countries in each region that are represented in our data: Ethiopia, China, Russia, Brazil, India, and Egypt. In Ethiopia, China, Russia, and Brazil, our simulations match the actual distributions reasonably well. In particular, we do not substantively underpredict the number of all-male firms.

By contrast, the patterns in India and Egypt are similar to what we found in Saudi Arabia. We substantially underpredict the number of firms with zero female employees. In India and Egypt, we predict that 20.7% and 21.7% of firms will be all male; in fact, 49.9% and 62.5% are. In addition, we 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 A2, we plot the same distributions but limit to 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.

[Figure 13 about here.]

VI.B.2 Ex-Ante Integration Rates

We next calculate ex-ante integration rates as a function of $\bar{\theta}_j n_j$ as in Section IV.B.3. We estimate $\theta(X_i)$ and $\bar{\theta}_j$ using ex-post integrated firms. In Figure XIV, we plot ex-ante integration rates for our six populous countries. 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. In fact, the relationship between ex-ante integration rates and $\bar{\theta}_j^{EP} n_j$ in Egypt and India is similar to what we measure in Saudi Arabia, as depicted in Figure II. As will be shown below, this pattern reflects more general regional differences in integration.

We also calculate ex-ante integration rates for each of the 65 countries represented in the World Bank data. We do this for two exercises. First, we use the Arab Barometer survey to test whether integration costs are correlated with social preferences for gender segregation. Second, we examine how integration costs relate to women’s labor force participation rates and the female share of the workforce.

We construct three measures to summarize integration costs by country. We calculate average ex-ante integration rates in two ways. First, we estimate $\theta(X_i)$ using ex-post integrated firms and construct overall ex-ante integration rates as in Section IV.B.3. Second, we estimate overall ex-ante integration rates using the structural approach described in Section IV.B.2. While overall ex-ante integration rates are straightforward to interpret, they also reflect variation in the distribution of $\bar{\theta}_j n_j$ across countries. For our third measure, we calculate the implied ex-ante integration rate for a “representative” firm with $\bar{\theta}_j^{EP} n_j = 10$.

[Figure 14 about here.]

Social preferences are correlated with ex-ante integration rates within the MENA region. Figure XV plots ex-ante integration rates by country (derived using ex-post integrated firms) against support for gender mixing from the Arab Barometer surveys. There is a strong positive correlation between ex-ante integration costs and local preferences for gender segregation within MENA. This is consistent with our motivating hypothesis that integration is costly for firms where social norms for gender segregation are strong.

[Figure 15 about here.]

Ex-ante integration costs are also correlated with women’s labor market outcomes across countries. Figure XVI plots country-level female labor force participation rates against overall ex-ante integration rates (derived using ex-post integrated firms). Table VII reports ordinary least squares (OLS) estimates of a regression of labor force participation measures on all three measures of integration costs. There are two main points to note. First, ex-ante integration rates are lowest in MENA and South Asia. Ten of the 14 MENA and South Asian countries in our data have average ex-ante integration rates below 60%, and no other countries fall in this range.⁴⁶ Second ex-ante integration rates are positively correlated with female labor force participation rate, both in absolute terms (columns (1) and (2)) and relative to male labor force participation (columns (3) and (4)). As we show in Appendix Table A5, the relationship is similar when we examine employment rates rather than labor force participation rates. The positive relationship between ex-ante integration rates and women’s economic engagement holds both between and within regions. This can be seen from that fact that when we add region fixed effects as controls (columns (2) and (4)), the relationship is muted but remains statistically and economically significant.

[Figure 16 about here.]

[Table 7 about here.]

VII Conclusion

We posit that where there are social norms for gender segregation, firms face costs to employing both men and women that are largely fixed. Motivated by a simple model of firm hiring, we develop a joint test for whether integration costs bind for any firm and a methodology for evaluating the firm-level consequences of those costs. We validate our approach using administrative employer-employee data and unique policy variation from Saudi Arabia, a country that strictly regulates between-gender interactions in the workplace during our period of study. We then apply our methodology to 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.

We find suggestive evidence that integration costs may depress aggregate demand for female labor. First, we document that Nitaqat—a gender-neutral policy in Saudi Arabia that had the unintended consequence of inducing many firms to integrate—increased female employment and wages. Second, across countries, integration costs are negatively correlated with the female share of the labor force. A 10 percentage point increase in ex-ante segregation rates is associated with a 4 to 6 percentage point decrease in female labor force participation rates.

⁴⁶The four MENA countries with ex-ante integration rates (derived using ex-post integrated firms) above 60% are Israel (89%), Lebanon (83%), Morocco (87%), and Tunisia (89%).

Integration costs also have the potential to generate a coordination problem: firms may not integrate unless enough women enter the labor market, and women may not enter the labor market unless enough firms have integrated. This interaction between the two sides of the market may generate a feedback loop: for example, a firm’s decision to integrate may increase the supply of women searching in the labor market, which in turn induces other firms to integrate. Unfortunately, we are unable to quantify the potential magnitude of these spillovers because we do not have data on labor supply and how female labor supply responds to the integration of local firms. However, such a coordination problem could be solved by policy.

Our results suggest that integration costs prevent some firms from hiring superior female candidates. Though beyond the scope of this paper, a natural question for further research is: how do integration costs affect productivity, both for firms and in the aggregate?

References

- Alesina, Alberto, Paola Giuliano, and Nathan Nunn**, “On the Origins of Gender Roles: Women and the Plough,” *Quarterly Journal of Economics*, May 2013, *128* (2), 469–530.
- Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber**, “An Evaluation of Instrumental Variable Strategies for Estimating the Effects of Catholic Schooling,” *Journal of Human Resources*, 2005, *40* (4), 791–821.
- Amir, Saman and Ahmad Shah Durrani**, “Should women get a job? “Yes...but” say Pakistani men,” Apr 2019.
- Arab News**, “Women’s rights gain focus in the Kingdom,” *Arab News*, 2010. June 15, 2010.
- , “Saudi Women No Longer Need Guardians’ Consent to Receive Services,” *Arab News*, 2017. May 5, 2017.
- Ashraf, Nava, Natalie Bau, Nathan Nunn, and Alessandra Voena**, “Bride Price and Female Education,” *Journal of Political Economy*, 2019.
- Augereau, Angelique, Shane Greenstein, and Marc Rysman**, “Coordination versus Differentiation in a Standards War: 56K Modems,” *RAND Journal of Economics*, Winter 2006, *37* (4), 887–909.
- Bayard, Kimberly, Judith Hellerstein, David Neumark, and Kenneth Troske**, “New Evidence on Sex Segregation and Sex Differences in Wages from Matched Employee-Employer Data,” *Journal of Labor Economics*, October 2003, *21* (4), 887–922.
- Becker, Gary**, *The Economics of Discrimination*, University of Chicago Press, 1957.
- Breckinridge, Sophonisba P.**, “Legislative Control of Women’s Work,” *Journal of Political Economy*, February 1906, *14* (2), 107–109.

- Bursztyn, Leonardo, Alessandra L. Gonzalez, and David Yanagizawa-Drott**, “Misperceived Social Norms: Female Labor Force Participation in Saudi Arabia,” June 2018.
- Card, David, Ana Rute Cardoso, and Patrick Kline**, “Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women,” *Quarterly Journal of Economics*, May 2016, 131 (2), 633–686.
- Dahl, Gordon B., Andreas Kotsadam, and Dan-Olof Rooth**, “Does Integration Change Gender Attitudes? The Effect of Randomly Assigning Women to Traditionally Male Teams,” February 2018. Unpublished manuscript.
- Ellison, Glenn and Edward L. Glaeser**, “Geographic Concentration in U.S. Manufacturing Industries: A Dartboard Approach,” *Journal of Political Economy*, 1997, 105 (5), 889–927.
- Evidence for Policy Design**, “Back to Work in a New Economy: Background Paper on the Saudi Labor Market,” Harvard Kennedy School, Harvard University April 2015.
- , “The JPCs seem to be working. Why not use them to increase female employment?,” EPoD Policy Insight, Harvard Kennedy School, Harvard University May 2017. <https://epod.cid.harvard.edu/sites/default/files/2018-06/JPC.pdf>.
- Fernandez, Raquel**, “Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century,” *American Economic Review*, 2013, 103 (1), 472–500.
- **and Alessandra Fogli**, “Culture: An Empirical Investigation of Beliefs, Work, and Fertility,” *American Economic Journal: Macroeconomics*, January 2009, 1 (1), 146–177.
- Field, Erica and Kate Vyborny**, “Female Labor Force Participation: Pakistan Country Study for the Asian Development Bank,” Working Paper April 2016.
- Fry, Erika**, “This CEO is helping Saudi women break a gender barrier,” *Fortune*, 2015. September 14, 2015.
- Garicano, Luis, Claire Lelarge, and John Van Reenan**, “Firm Size Distortions and the Productivity Distribution: Evidence from France,” *American Economic Review*, November 2016, 106 (11), 3439–79.
- GaStat**, *Labor Force Survey 2011*, General Authority for Statistics, Kingdom of Saudi Arabia, 2011.
- Goldin, Claudia**, “Monitoring Costs and Occupational Segregation by Sex: A Historical Analysis,” *Journal of Labor Economics*, January 1986, 4 (1), 1–27.
- , “A Pollution Theory of Discrimination: Male and Female Differences in Occupations and Earnings,” in Leah Platt Boustan, Carola Frydman, and Robert A. Margo, eds., *Human Capital in History: The American Record*, University of Chicago Press, 2014, pp. 313–348.

- Gorman, Elizabeth H.**, “Gender Stereotypes, Same-Gender Preferences, and Organizational Variation in the Hiring of Women: Evidence from Law Firms,” *American Sociological Review*, August 2005, 70 (4), 702–728.
- Gourio, François and Nicolas Roys**, “Size-dependent regulations, firm size distribution, and reallocation,” *Quantitative Economics*, July 2014, 5 (2), 377–416.
- Groshen, Erica L.**, “The Structure of the Female/Male Wage Differential: Is It Who You Are, What You Do, or Where You Work?,” *Journal of Human Resources*, Summer 1991, 26 (3), 457–472.
- Hamermesh, Daniel S. and Gerard A. Pfann**, “Adjustment Costs in Factor Demand,” *Journal of Economic Literature*, September 1996, 34 (3), 1264–1292.
- Humphries, Jane**, ““... The Most Free From Objection” The Sexual Division of Labor and Women’s Work in Nineteenth-Century England,” *Journal of Economic History*, December 1987, 47 (4), 929–949.
- International Finance Corporation**, *Investing in Women’s Employment : Good for Business, Good for Development*, World Bank, 2013. <https://openknowledge.worldbank.org/handle/10986/16257>.
- Jayachandran, Seema**, “The Roots of Gender Inequality in Developing Countries,” *Annual Review of Economics*, August 2015, 7, 63–88.
- Khoja, Sara**, “Considerations for Female Employees in KSA,” Clyde and Co.: Insight 2016. <https://www.clydeco.com/insight/article/considerations-for-female-employees-in-ksa>.
- **and Sarit Thomas**, “Labour and Employment Compliance in Saudi Arabia,” in del Rey Salvador and Robert J. Mignin, eds., *International Labour and Employment Compliance Handbook*, Kluwer Law International BV, Netherlands, 2018.
- Kleven, Henrik**, “Bunching,” *Annual Review of Economics*, October 2016, 8, 435–464.
- Kuhn, Peter and Kailing Shen**, “Gender Discrimination in Job Ads: Evidence from China,” *Quarterly Journal of Economics*, February 2013, 128 (1), 287–336.
- Mies, Frank P.**, “Statutory Regulation of Women’s Employment—Codification of Statutes,” *Journal of Political Economy*, February 1906, 14 (2), 109–118.
- Miller, Amalia R. and Carmit Segal**, “Does Temporary Affirmative Action Produce Persistent Effects? A Study of Black and Female Employment in Law Enforcement,” *Review of Economics and Statistics*, 2012, 94 (4), 1107–1125.
- Miller, Conrad**, “The Persistent Effect of Temporary Affirmative Action,” *American Economic Journal: Applied Economics*, July 2017, 9 (3), 152–90.

- Oster, Emily**, “Unobservable Selection and Coefficient Stability: Theory and Evidence,” *Journal of Business and Economic Statistics*, 2019, *37* (2), 187–204.
- Pan, Jessica**, “Gender Segregation in Occupations: The Role of Tipping and Social Interactions,” *Journal of Labor Economics*, April 2015, *33* (2), 365–408.
- Peck, Jennifer R.**, “Can Hiring Quotas Work? The Effect of the Nitaqat Program on the Saudi Private Sector,” *American Economic Journal: Economic Policy*, May 2017, *9* (2), 316–347.
- Sara Hamdan**, “Saudi Arabia Signals Openness to Women Seeking Work,” *New York Times*, 2012. September 5, 2012.
- Saudi Gazette**, “Nitaqat percentage linked to salary,” *Saudi Gazette*, 2012. September 9, 2012.
- Schaner, Simone and Smita Das**, “Female Labor Force Participation in Asia: Indonesia Country Study,” ADB Economics Working Paper 474, Asian Development Bank February 2016.
- Tomaskovic-Devey, Donald, Kevin Stainback, Tiffany Taylor, Catherine Zimmer, Corre Robinson, and Tricia McTague**, “Documenting Desegregation: Segregation in American Workplaces by Race, Ethnicity, and Sex, 1966-2003,” *American Sociological Review*, August 2006, *71* (4), 565–588.
- Varejão, José and Pedro Portugal**, “Employment Dynamics and the Structure of Labor Adjustment Costs,” *Journal of Labor Economics*, January 2007, *25* (1), 137–165.
- Whately, Warren C.**, “Getting a Foot in the Door: “Learning,” State Dependence, and the Racial Integration of Firms,” *Journal of Economic History*, March 1990, *50* (1), 43–66.
- World Bank**, *World Development Indicators 2016*, The World Bank, 2016.
- , *Women, Business and the Law 2018*, The World Bank, 2018.

A Saudi Female Employment Policies

In addition to Nitaqat, the Saudi government also pursued a slate of practical measures designed to increase women’s employment over the study period, including the Retail Employment Decree, the Hafiz program, and updates to the guardianship system. The King issued a royal decree in 2011 mandating that shops selling lingerie and cosmetics employ only Saudi women as salesclerks beginning in August 2012. The decree was expanded to also cover stores selling women’s clothing and accessories beginning in January 2014. There were recently plans to further expand the decree to cover all stores selling goods of primary interest to women, such as pharmacies with cosmetics sections and fabric stores (Evidence for Policy Design, 2015).

Though not gender-specific, the Hafiz unemployment assistance program has also drawn women into the workforce and supported their private sector job search. Hafiz provides a monthly financial stipend to unemployed Saudis who make weekly check-ins to a government-sponsored online job search portal (Taqat Online). More than 90% of Hafiz beneficiaries have been women (Evidence for Policy Design, 2017). The MoL removed regulations requiring women to obtain permission from a male guardian to apply for private sector jobs.⁴⁷ Many firms still require a guardian’s approval, though the Ministry recently forbade this practice among government employers.⁴⁸

B Appendix: Data

Administrative data from the Nitaqat program is used to identify the Nitaqat compliance status of firms. As described by Peck (2017), the Nitaqat database is used to track compliance with national quotas on Saudi employment in the private sector. The database collects information on whether a given firm was subject to quotas during a given week, and, if so, whether it met the quotas for that particular week. These data provide weekly quota compliance information from June, 2011 (the start of the Nitaqat program) until December, 2013.

Firms are defined differently between the Nitaqat and GOSI data sets. In the latter, firms are defined by their legal status as a commercial organization operating in potentially multiple industries. In the Nitaqat data, however, the operations of such firms are further classified into entities, which are subject to different quotas depending on the industry category each entity operates in and, as described in the main text, the size group based on the total number of employees. For example, a firm operating a bakery and a jewelry store would be considered two separate entities facing different quotas (and would therefore contain two entries in the data for each time period)⁴⁹. In the GOSI data, however, such a firm would be considered a single firm. Firms with multiple entities can also list as a single entity (in the “Multiple Economic Activities” industry) but would be subject to the most stringent quota they face based on the entities under their umbrella. To

⁴⁷Jafar AlShayeb, *Arab News* June 15, 2010 “Women’s rights gain focus in the Kingdom”

⁴⁸Lulwa Shalhoub, *Arab News* May 5, 2017 “Saudi women no longer need guardians’ consent to receive services” <http://www.arabnews.com/node/1094681/saudi-arabia>

⁴⁹An entity consisting of multiple branches (e.g., a national franchise) are counted as a single entity for each branch of the MLSL labor office they are linked to.

harmonize the definition of the firm between the two data sets, firms with multiple entities in the Nitaqat data were aggregated together by summing their employee counts, and assigning the color and size status by the most binding entity quota (as measured by the number of Saudis required to fulfill it) the firm faces. The number of Saudis the firm needs to hire, however, was summed across all entities to create a single metric for the distance of the firm to the quota. This transformation only affects 58,000 of the approximately 1.07 million firms in the Nitaqat data.

In addition to the distinction between entities and firms, it should be noted that the firm identifiers used by both GOSI and the Nitaqat data define firms with a national or multicity presence as separate commercial organizations depending on the geographic MLSD office they register with. For example, a firm with branches in Riyadh and Dammam would count as two firms, both of which are subject to separate quota calculations. The geographic scope of the MLSD offices is quite broad, and are typically at the provincial level. The definition of the firm we use in this paper therefore can be thought of as a legal commercial organization within a particular province.

To clean up potentially erroneous observations, we drop individuals with ages below 10 or above 100 in the GOSI data.

C Appendix: Testing and Estimation Details

C.1 Simulating Bunching at Zero When Integration Costs Bind

In this section we demonstrate the rationale for our test of the null hypothesis that no firms face binding integration costs, developed in Section IV.A. In this test, we simulate the distribution of female hires across firms under the null hypothesis, and compare this to the distribution we observe in practice. We show here that if some firms are in fact ex-ante segregated, the distribution we simulate will generally *underpredict* the number of firms with zero female hires. We demonstrate this point using simulation.

Our simulation exercise builds on the model above by positing that some exogeneously determined γ share of firms are integrated and that, under the null hypothesis, θ_0 is the probability that each hire is female. Under both hypotheses, θ_0 is the expected female share of employees pooled across all firms. Firms are characterized by their number of hires, n .

Under the null hypothesis (H_0), all firms are integrated ($\gamma = 1$). In this case, firms which do not hire any women do so by chance alone. Alternatively (H_a), if $\gamma < 1$, then some firms do not hire women because they are ex-ante segregated. We show via simulation that under H_a there are generally a greater share of firms with zero female employees.

We consider two scenarios: one where the probability of integration is constant across firms and a second where integration rates are increasing in firm size.

C.1.1 Constant Integration Probability

First, we assume that the probability of integration is constant across firms, and given by γ . In this case, under H_a , the probability that a hire is female at an ex-ante integrated firm is $\theta_a = \theta_0/\gamma$. Our simulation is structured as follows. We first set a value of γ , the share of integrated firms, and θ_a , the probability a hire is female in an ex-ante integrated firm under H_a . Then, for each run of the simulation, we:

1. sample firm sizes (ie. the total number of employees) from a log-normal distribution with mean 50 and standard deviation 500, approximately matching the distribution of firm sizes we observe in our Saudi employment data (see Table II);
2. determine whether a firm in our sample is integrated with probability $\gamma < 1$ for H_a ; all firms are considered integrated under H_0 ($\gamma = 1$);
3. for H_a determine the gender of each hire via a binomial draw with probability θ_a that each hire is female. Sum these hires to determine the count of female employees for each firm under H_a ;
4. set θ_0 , the probability of a hire being female under H_0 using the overall female share of employment simulated in the prior step⁵⁰, then similarly determine the gender of each hire and count the number of female employees for each firm under H_0

After running the above simulation 1,000 times, we calculate the share of simulations where the number of firms with zero female employees under H_a (Z_a) exceeds the same value under H_0 (Z_0). We show in Figure C1 what the distribution of female employee counts look like under both hypotheses for $\gamma = 0.7$ and $\theta_a = 0.5$.⁵¹ Each column represents the mean across simulations, whereas the error bars represent the 5th and 95th percentiles.

We then repeat this exercise by iterating over values of $\gamma \in (0, 1)$ and $\theta_a \in (0, 1)$. We plot the share of simulations where $Z_a > Z_0$ for each γ and θ_a value in figure C2 below.

Except the largest values of γ ($\gamma \geq 0.9$), $Z_a > Z_0$ for virtually all simulation draws. When γ is large, $Z_a > Z_0$ for the majority of simulation draws, but this share gets as low as the 0.6 – 0.7 range (when $\gamma = 0.95$ and $\theta_a < 0.075$).

C.1.2 Integration Rates Increasing in n

If integrated costs are largely fixed, firms which have to hire more employees may be more likely to integrate. In this case, integration rates are increasing in n . To account for this, we again draw n from log-normal distribution, and also generate firm specific integration likelihoods $\gamma_i \sim$

⁵⁰This allows us to have approximately equal numbers of female employees under both hypotheses.

⁵¹These values are chosen primarily for testing purposes. Repeating the exercise for different values results in similar patterns as shown below.

$Beta(\beta \frac{\gamma}{1-\gamma}, \beta)$ where $\gamma = \bar{\gamma}_i$.⁵² To introduce the correlation between these two marginal distributions, we conduct a Cholesky decomposition to create a joint distribution of n_i and γ_i across such that the correlation between n and γ is positive.

We then continue the simulations as above, but iterate over values of $\bar{\gamma}_i$ and θ_a , and determine whether a firm is integrated according to its specific γ_i integration probability. We show in figure C3 the distribution of female employment for $\gamma = 0.7$ and $\theta_a = 0.5$ as above. We similarly plot the share of simulations where $Z_a > Z_0$ for each γ and θ_a value in figure C4.

As above, except the largest values of γ ($\gamma \geq 0.9$), $Z_a > Z_0$ for virtually all simulation draws.

C.2 Structural Estimation of θ using Expectation-Maximization

In Section IV.B.2 we modeled the distinction between ex-ante and ex-post integrated firms to structurally estimate $\theta(X_i)$. We use an expectation-maximization algorithm to estimate these parameters. Continuing from Section IV.B.2, the likelihood function for firm j is

$$P(Y_j = Y) = \begin{cases} \pi_j + (1 - \pi_j) \prod_{i=1}^{N_j} (1 - \theta_{ij}) & \text{if } K_j = 0 \\ (1 - \pi_j) \prod_{i=1}^{K_j} \theta_{ij} \prod_{i=K_j+1}^{N_j} (1 - \theta_{ij}) & \text{if } 0 < K_j < N_j \\ (1 - \pi_j) \prod_{i=1}^{N_j} \theta_{ij} & \text{if } K_j = N_j. \end{cases}$$

We model both θ_{ij} and π_j in logistic regression models with explanatory variables X_{ij} and Z_j , respectively:

$$\begin{aligned} \theta_{ij} &= \Lambda(X_{ij}\beta) \\ \pi_j &= \Lambda(Z_j\gamma) \end{aligned}$$

where Λ is the logistic function.

The log-likelihood for firm j is

$$\log(f_j) = \log(P(Y_j = Y)) = \begin{cases} -\log(1 + e^{Z_j\gamma}) + \log\left(e^{Z_j\gamma} + \prod_{i=1}^{N_j} (1 + e^{X_{ij}\beta})^{-1}\right) & \text{if } K_j = 0 \\ -\log(1 + e^{Z_j\gamma}) - \sum_{i=1}^{N_j} \left(1 + e^{X_{ij}\beta}\right) + \sum_{i=1}^{K_j} X_{ij}\beta & \text{if } 0 < K_j < N_j \\ -\log(1 + e^{Z_j\gamma}) + \sum_{i=1}^{N_j} \left[X_{ij}\beta - \log(1 + e^{X_{ij}\beta})\right] & \text{if } K_j = N_j. \end{cases}$$

⁵²We pick this particular form of the *Beta* distribution as its mean is γ . In other words, for a given share of firms integrated, we can generate a distribution of integration likelihoods for each firm such that the mean is equal to the overall share of firms integrated. In this case β acts as a scaling parameter but does not affect the mean.

Combining each firm's log likelihood, we write our log-likelihood function as:

$$l(\beta, \gamma; Y_j, X_{ij}, Z_j) = \sum_{j=1}^J \log(f_j)$$

We obtain maximum likelihood estimates of γ and β using the expectation-maximization (EM) algorithm. The EM algorithm is an iterative method to find maximum likelihood estimates, where the model depends on unobserved latent variables. The EM algorithm alternates between an expectation (E) step, which creates a function for the expectation of the log-likelihood evaluated at the current estimate for the parameters, and a maximization (M) step, which computes parameters maximizing the expected log-likelihood found in the E step.

For each firm j , let the unobserved random variable I_j indicate whether a firm has ex-ante integrated. When $I_j = 0$, firm j is ex-ante segregated and Y_j is necessarily zero. When $I_j = 1$, firm j is ex-ante integrated. If we could observe I_j for every firm, then the log-likelihood for firm j given complete data (Y_j, I_j) would be

$$\log(f_j|I_j, X_{ij}, Z_j) = (1-I_j) \underbrace{(Z_j\gamma - \log(1 + e^{Z_j\gamma}))}_{\log_\gamma(f_j|I_j, X_{ij}, Z_j)} + I_j \underbrace{\left[-\sum_{i=1}^{N_j} \log(1 + e^{X_{ij}\beta}) + \mathbf{1}_{0 < K_j \leq N_j} \sum_{i=1}^{K_j} X_{ij}\beta \right]}_{\log_\beta(f_j|I_j, X_{ij}, Z_j)}.$$

Therefore the complete data log-likelihood function is

$$\begin{aligned} l_c(\beta, \gamma|Y_j, I_j, X_{ij}, Z_j) &= \sum_{j=1}^J \log(f_j|I_j, X_{ij}, Z_j) \\ &= \sum_{j=1}^J [\log_\gamma(f_j|I_j, X_{ij}, Z_j) + \log_\beta(f_j|I_j, X_{ij}, Z_j)] \\ &= l_c(\gamma|Y_j, I_j, X_{ij}, Z_j) + l_c(\beta|Y_j, I_j, X_{ij}, Z_j). \end{aligned}$$

The EM algorithm begins with starting values $\omega^{(0)} = (\gamma^{(0)}, \beta^{(0)})$. Our starting value for $\beta^{(0)}$ is derived from estimating the linear regression $Y_j = X_{ij}\beta$ and setting $\beta^{(0)} = \hat{\beta}$. For $\gamma^{(0)}$, we estimate the regression $\mathbf{1}_{K_j > 0} = Z_j\gamma$ and similarly set $\gamma^{(0)} = \hat{\gamma}$.

From these initial values, we proceed iteratively, with (r) indexing the iteration:

- **E Step:** estimate I_j by its conditional mean $I_j^{(r)}$ given $\omega^{(r)} = (\gamma^{(r)}, \beta^{(r)})$:

$$\begin{aligned}
\hat{I}_j^{(r)} &= E[I_j | Y_j, X_{ij}, Z_j, \gamma^{(r)}, \beta^{(r)}] \\
&= \frac{P(Y_j | I_j = 0)P(I_j = 0)}{P(Y_j | I_j = 0)P(I_j = 0) + P(Y_j | I_j = 1)P(I_j = 1)} \\
&= \begin{cases} \left[1 + e^{-G_j \gamma} \prod_{i=1}^{N_j} (1 + e^{X_{ij} \beta})^{-1} \right]^{-1} & \text{if } K_j = 0 \\ 0 & \text{if } K_j \neq 0 \end{cases}
\end{aligned}$$

- **M Step for γ :** we find $\gamma^{(r+1)}$ by maximizing $l_c(\gamma | Y_j, I_j, X_{ij}, Z_j)$. This can be accomplished by logistic regression of $I_j^{(r)}$ on Z_j . It is equivalent to solving the FOC of $l_c(\gamma | Y_j, I_j, X_{ij}, Z_j)$:

$$\sum_{j=1}^J \left(I_j^{(r)} - \frac{e^{Z_j \gamma}}{1 + e^{Z_j \gamma}} \right) Z_j = 0.$$

- **M Step for γ :** we find $\gamma^{(r+1)}$ by maximizing $l_c(\gamma | Y_j, I_j, X_{ij}, Z_j)$. This can be accomplished by logistic regression of $I_j^{(r)}$ on Z_j . It is equivalent to solving the FOC of $l_c(\gamma | Y_j, I_j, X_{ij}, Z_j)$:

$$\sum_{j=1}^J \left(I_j^{(r)} - \frac{e^{Z_j \gamma}}{1 + e^{Z_j \gamma}} \right) Z_j = 0.$$

From the above, we obtain estimates for β and γ for iteration (r) and repeat the exercise until $\|\beta^{(r+1)} - \beta^{(r)}\| + \|\gamma^{(r+1)} - \gamma^{(r)}\| < 0.0001$.

D Appendix: Ghost Employment

The main text mentions the concern that firms may falsify their employee records with GOSI to meet their quotas after Nitaqat, so reported employment numbers may not reflect real employment, particularly for women. Private sector firms are required to register their employees with GOSI and to pay a fraction of the reported wage into the employee's social security account. Nationals may not be registered as full-time employees for more than one firm at the same time. Workers have some incentive to make sure these records are filed accurately so that their eventual retirement payments are accurate. The Nitaqat enforcement system draws directly on these GOSI records to monitor the number of Saudi workers registered as employees at each firm. "Ghost employment" is used to refer to a variety of situations in which the worker is not doing the job as reported to GOSI. This can range from cases of outright fraud (e.g., where a worker's National ID Number is used without the worker's knowledge or permission) to cases where the worker draws the reported

salary but does not perform meaningful work at the firm.⁵³ This ghost employment would cause our analysis to overstate the degree to which firms hire Saudi women in response to employment quotas. In this analysis we investigate whether this phenomenon becomes more common after the start of Nitaqat and whether it appears to be more common for women than for men.

To do this, we examine the share of workers hired in each month who appear to have “active” career trajectories. We define a worker as being active if their job history shows that they switch firms, receive wage increases, change occupations, or make above minimum wage. We can be reasonably confident that workers that experience these events are “real” employees: firms have no incentive to report paying fake workers above minimum wage (as this simply increases their GOSI payments without providing Nitaqat benefits), and there is similarly no reason to promote them, give them raises, or move their IDs to other firms. We construct an indicator equal to 1 if the worker experiences any of these actions (change wage or occupation, switch firms, or make above minimum wage) within 24 months of their first appearance in the GOSI system.⁵⁴

In addition to capturing ghost employment, GOSI records may be inaccurate for several other reasons. First, firms may register artificially low wages in order to minimize their social security payments on behalf of their employees. This can in principle be checked by the worker, but there are some accounts of workers being surprised by their wage records upon retirement. Firms may also neglect to record promotions in the GOSI system, so recorded wages may lag actual wages. Movements across firms seem likely to be accurate, as a prior employer will not want to make payments for people who are no longer employees, and new firms will want to have the worker’s national ID number released so they can register a new hire. These will bias the measure toward under-counting active employees, so the count of “inactive” workers should be assumed to include not only ghost employees, but also employees whose records are not updated promptly as well as workers who simply do not experience job status changes over the period.⁵⁵

Figure E1 shows a plot of the share of workers hired in each month that experience at least one of these events within 24 months of being hired. The share of workers who change job status is relatively steady for both genders at about 58% for men and 47% for women. As discussed before, there are a variety of reasons (aside from ghost employment) why this might only apply to half of workers. First, workers may simply not be promoted within 24 months of their first entry into the private sector. Second, they may be promoted but not have the promotions recorded in GOSI. Although only about half of workers experience official status changes within two years of hire, the

⁵³There may also be cases in between, for example where workers collect a one-time payment or ongoing small payment from the firm to use their ID numbers.

⁵⁴One potential issue is the de facto increase in the minimum wage in 2013. GOSI had previously required firms to enter a minimum wage of 1500SAR per month. In January 2013 firms were only given pro-rated Nitaqat credit for Saudi employees paid less than 3000SAR a month (e.g., a worker being paid the previous minimum of 1500SAR would count as 0.5 Saudis for Nitaqat purposes). Because of this, we do not consider increases from 1500 to 3000SAR that occur after January 2013 to be wage increases.

⁵⁵Firms may also retain previous workers who have exited the labor market on their GOSI employment rolls. These workers will mistakenly appear to be active. Because we focus on workers hired between 2009 and 2013 we expect that this will comprise a only a very small part of the workforce, as these workers would need to enter the labor force after 2009, experience a change in wage, occupation, or firm, and then leave the private sector workforce without retiring and drawing their GOSI pension.

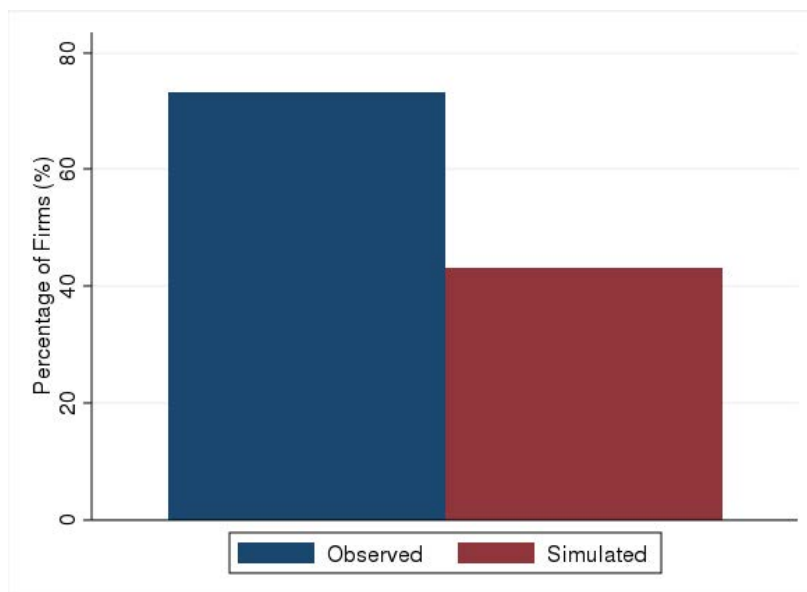
patterns are similar across genders and relatively stable over time. There is a slight decrease in the share of workers promoted for those hired after Nitaqat.

Within these series we may be concerned also about compositional changes in the types of workers that are being hired before and after Nitaqat as well as the types of firms that hire Saudis before and after the policy change. There is ample evidence that Saudis hired after Nitaqat are different from those hired before: more are women, more are hired with lower skill levels, and married women are more likely to join the labor force. Red and Yellow firms, which were most incentivized to increase Saudi hiring, were also potentially less desirable places for Saudis to work and may be less likely to keep their GOSI records up to date and to promote their employees over time. Figure E2 shows the plot of these shares controlling for some worker characteristics: age, education, and marital status of the new hires.


Women are more likely to be active workers when controlling for observable worker characteristics, and the likelihood of promotion appears to be steadily increasing over time for women. We therefore conclude that even if ghost employment is captured by the GOSI data it does not appear to worsen after Nitaqat, and does not worsen for women in particular.

FIGURE I
DISTRIBUTION OF FEMALE EMPLOYMENT ACROSS FIRMS

(a) Percentage of Firms with Zero Female Employees




(b) Percentage of Firms with > 0 Female Employees



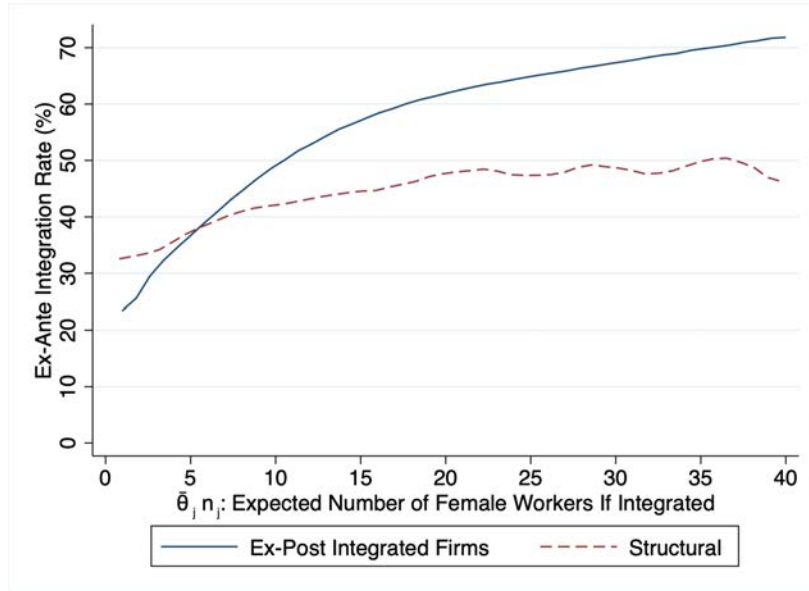
Note: This set of figures compares the observed distribution of female employment across firms in January 2009 to distributions simulated under the null hypothesis that no firm faces binding integration costs. Sample selection and simulation details are described in Sections IV.A.1 and IV.A.2. Panel A plots the share of firms with zero female employees in both the observed and simulation distributions. Panel B plots the share of firms with various nonzero totals of female in employees in both the observed and simulated distributions. For all simulations, a Kolmogorov-Smirnov test rejects equality of the observed and simulated distributions at the 0.01 significance level.

FIGURE II
INTEGRATION RATES BY $\bar{\theta}_j n_j$

(a) Integration Status




(b) Ex-Ante Integration Rates




Note: This set of figures depicts the relationship between ex-post and ex-ante integration rates and $\bar{\theta}_j n_j$, a firm’s expected number of female employees if ex-ante integrated. We construct $\bar{\theta}_j n_j$ for each firm j using an estimate of $\theta(X_i)$ —either $\hat{\theta}^{EP}(X_i)$ or $\hat{\theta}^S(X_i)$ —and the job composition of firm j . Panel A plots both the observed ex-post integration rate and the simulated ex-post integration rate, where the latter is simulated under the null hypothesis that all firms are ex-ante integrated. In Panel A, $\bar{\theta}_j$ is constructed using $\hat{\theta}^{EP}(X_i)$ and the job mix in firm j . Panel B plots ex-ante integration rates. “Ex-Post Integrated Firms” is constructed as described in Section IV.B.3. “Structural” plots the average estimated values of π_j (described in Section IV.B.2) as a function of $\bar{\theta}_j n_j$, where $\bar{\theta}_j$ is constructed using $\hat{\theta}^S(X_i)$ and the job mix in firm j .

FIGURE III
 DISTRIBUTION OF FEMALE EMPLOYMENT UNDER MODEL

(a) Percentage of Firms with Zero Female Employees





(b) Percentage of Firms with > 0 Female Employees



Note: This set of figures compares the observed distribution of female employment across firms in January 2009 to simulated distributions where we allow ex-ante integration rates to vary by $\bar{\theta}_j n_j$ as in Figure II. The simulation exercise is described in more detail in Section IV.B.3. Sample selection details are described in Section IV.A.1. Panel A plots the share of firms with zero female employees in both the observed and simulation distributions. Panel B plots the share of firms with various nonzero totals of female in employees in both the observed and simulated distributions. For all simulations, a Kolmogorov-Smirnov test fails to reject equality of the observed and simulated distributions at the 0.01 significance level.


FIGURE IV
 FEMALE SHARE OF HIRES AT NEWLY INTEGRATED FIRMS




Note: This set of figures describes the sex composition of hires at newly integrated firms. Panel A plots the female share of hires made at integrating firms in six-month increments relative to a firm’s first observed female hire, averaged across firms. We restrict to firms with at least five Saudi employees in the month prior to integration. Panel B compares the female share of hires at newly integrated firms to their $\theta(X_i)$ -based predicted values, where $\theta(X_i)$ is estimated using firms that are ex-post integrated in January 2009. $\theta(X_i)$ estimation details are provided in Section IV.C.1. The vertical axis depicts the female share of hires that are made 12 or more months following a firm’s first female hire. The horizontal axis depicts the $\theta(X_i)$ -based prediction for this value.

FIGURE V
 DISTRIBUTION OF FEMALE HIRING, BASELINE ALL-MALE FIRMS

(a) Percentage of Firms with Zero Female Hires




(b) Percentage of Firms with > 0 Female Hires




Note: This set of figures compares the observed and simulated distributions of female hires across firms that are ex-post segregated in January 2009 for hires made between February 2009 and June 2015. The simulated distributions are simulated under the null hypothesis that no firm faces binding integration costs over the hiring period. Sample selection and simulation details are described in Sections IV.C.2. Panel A plots the share of firms with zero female hires in both the observed and simulation distributions. Panel B plots the share of firms with various nonzero totals of female hires in both the observed and simulated distributions. For all simulations, a Kolmogorov-Smirnov test rejects equality of the observed and simulated distributions at the 0.01 significance level.

FIGURE VI
 DISTRIBUTION OF FEMALE HIRES, BASELINE INTEGRATED FIRMS

(a) Percentage of Firms with Zero Female Employees




(b) Percentage of Firms with > 0 Female Hires



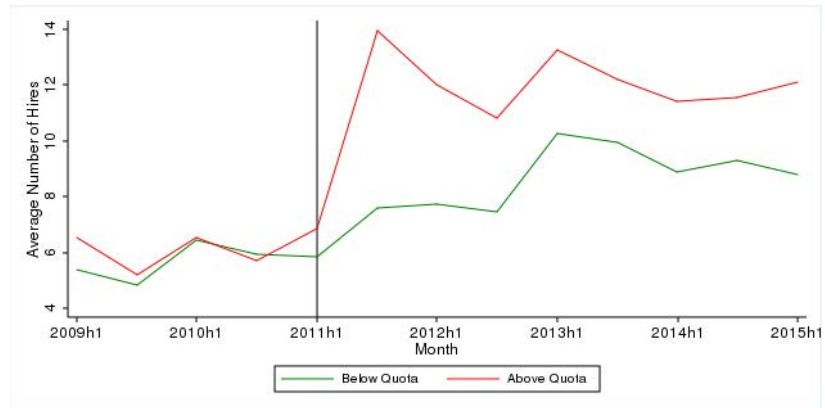
Note: This set of figures compares the observed and simulated distributions of female hires across firms that are ex-post integrated in January 2009 for hires made between February 2009 and June 2015. The simulated distributions are simulated under the null hypothesis that no firm faces binding integration costs over the hiring period. Sample selection and simulation details are described in Sections IV.C.2. Panel A plots the share of firms with zero female hires in both the observed and simulation distributions. Panel B plots the share of firms with various nonzero totals of female hires in both the observed and simulated distributions. For all simulations, a Kolmogorov-Smirnov test fails to reject equality of the observed and simulated distributions at the 0.01 significance level.

FIGURE VII
 SAUDI EMPLOYMENT AND BELOW VERSUS ABOVE QUOTA FIRMS



Note: This figure compares the percent change in Saudi employees relative to June 2011 for Above Quota and Below Quota. Above Quota firms are Green and Platinum firms, and Below Quota firms are Yellow and Red firms. Color refers to firm quota status in June 2011, and the vertical line marks June 2011.

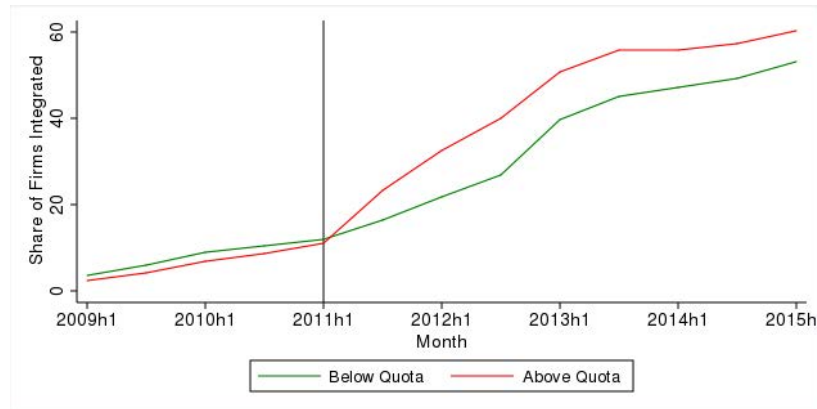
FIGURE VIII
 NUMBER OF HIRES PER HALF-YEAR AT BASELINE-SEGREGATED FIRMS



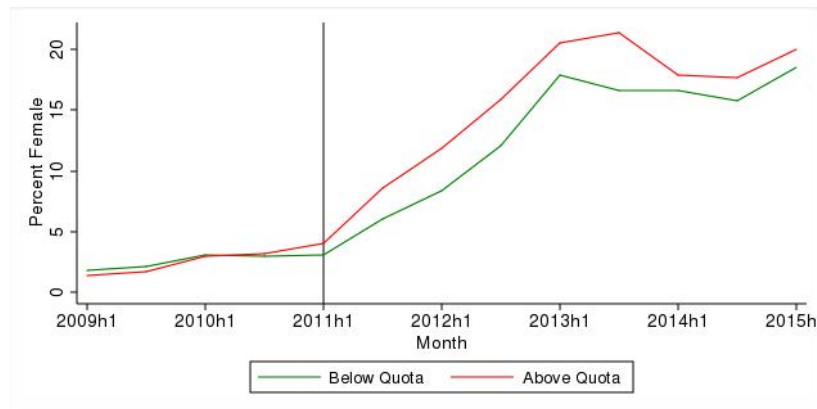
Note: This figure compares the number of hires per half-year for Above Quota and Below Quota firms that were ex-post segregated in January 2009. Above Quota firms are Green and Platinum firms, and Below Quota firms are Yellow and Red firms. Color refers to firm quota status in June 2011. We restrict to firms that had at least five Saudi employees in January 2009. The vertical line marks the first half of 2011. Nitaqat is implemented in June 2011.

FIGURE IX
 INTEGRATION RATES AND FEMALE SHARE OF HIRES OVER TIME

(a) Ex-Post Integration Rates




(b) Female Share of Hires



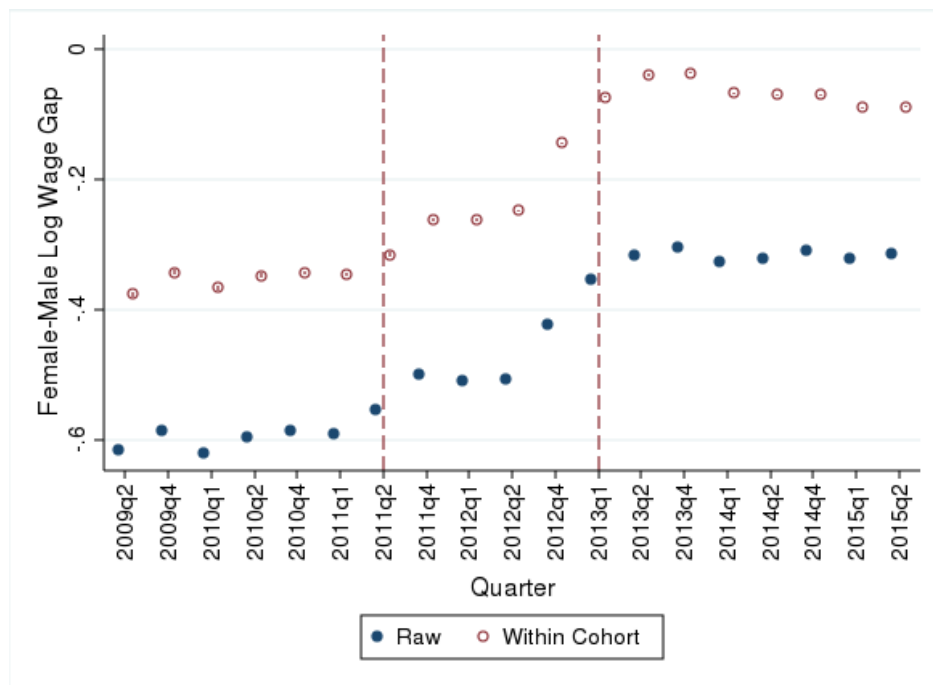
Note: This set of figures compares integration rates and the female share of hires for Above Quota and Below Quota firms that were ex-post segregated in January 2009. Above Quota firms are Green and Platinum firms, and Below Quota firms are Yellow and Red firms. Color refers to firm quota status in June 2011. We restrict to firms that had at least five Saudi employees in January 2009. The vertical line marks the first half of 2011. Nitaqat is implemented in June 2011. Panel A plots the share of firms that are ex-post integrated by half-year. Panel B plots the female share of hires at each firm, averaged across firms. Firms that do not make any hires in a given half-year are not included in the calculation for that period.

FIGURE X
 INTEGRATION RATES AND FEMALE SHARE OF WORKFORCE OVER TIME




Note: This figure plots the female share of full-time Saudi workers and the share of firms that employ both men and women, both on a quarterly basis. For the latter outcome, firms are restricted to those with at least five Saudi employees. The vertical lines correspond to the implementation of Nitaqat (in Q2 of 2011) and the increase in the de facto minimum wage (in Q1 of 2013).

FIGURE XI
GENDER WAGE GAP OVER TIME



Note: This figure plots the female-male log wage gap on a quarterly basis. It includes both the raw log wage gap and the log wage gap controlling for cohort fixed effects, where cohorts refer to the year of the earliest start date for a worker as recorded in the GOSI data. The vertical lines correspond to the implementation of Nitaqat (in Q2 of 2011) and the increase in the de facto minimum wage (in Q1 of 2013).


FIGURE XII
 FEMALE SHARE OF WORKFORCE BY BASELINE INTEGRATION STATUS



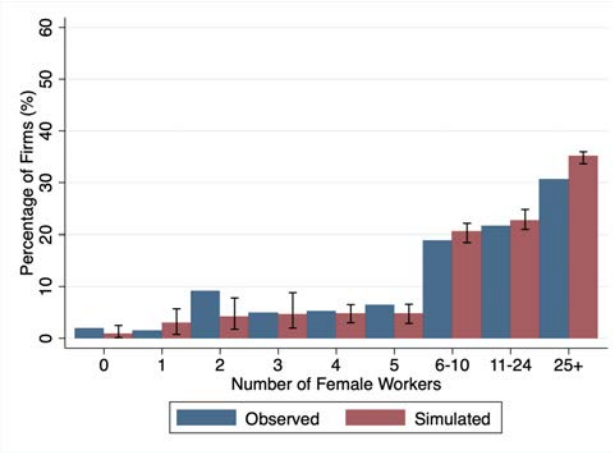
Note: This figure plots the female share of employment over time in firms that employed Saudis in January 2009, split by the firm's ex-post integration status in that month. These shares are measured on a quarterly basis. The vertical lines correspond to the implementation of Nitaqat (in Q2 of 2011) and the increase in the de facto minimum wage (in Q1 of 2013).

FIGURE XIII
 DISTRIBUTION OF FEMALE EMPLOYMENT ACROSS FIRMS, BY COUNTRY


(a) Ethiopia




(b) China




(c) Russia




(d) Brazil



(e) India

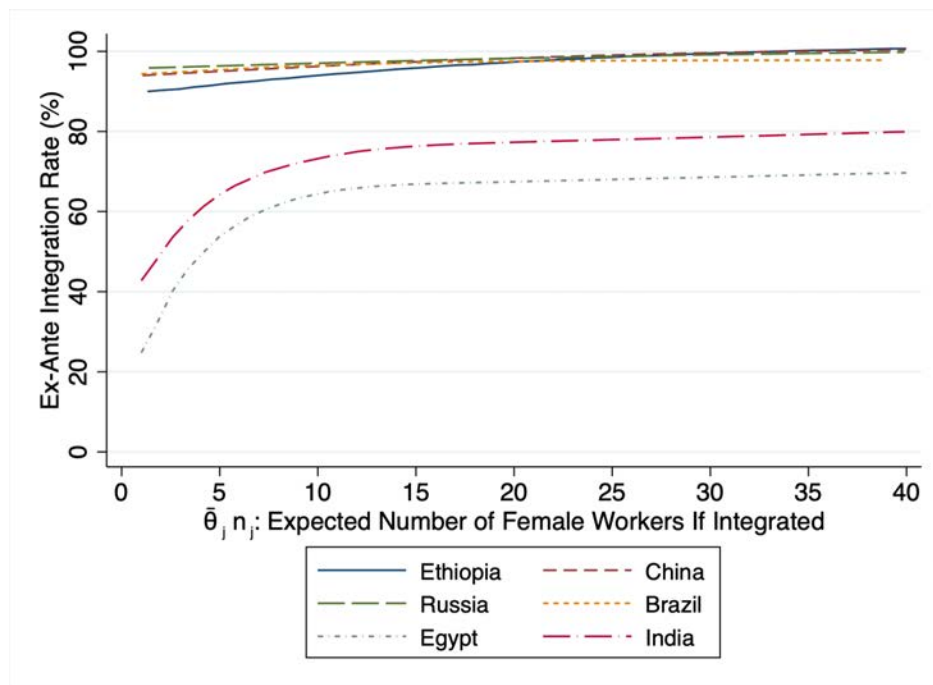


(f) Egypt



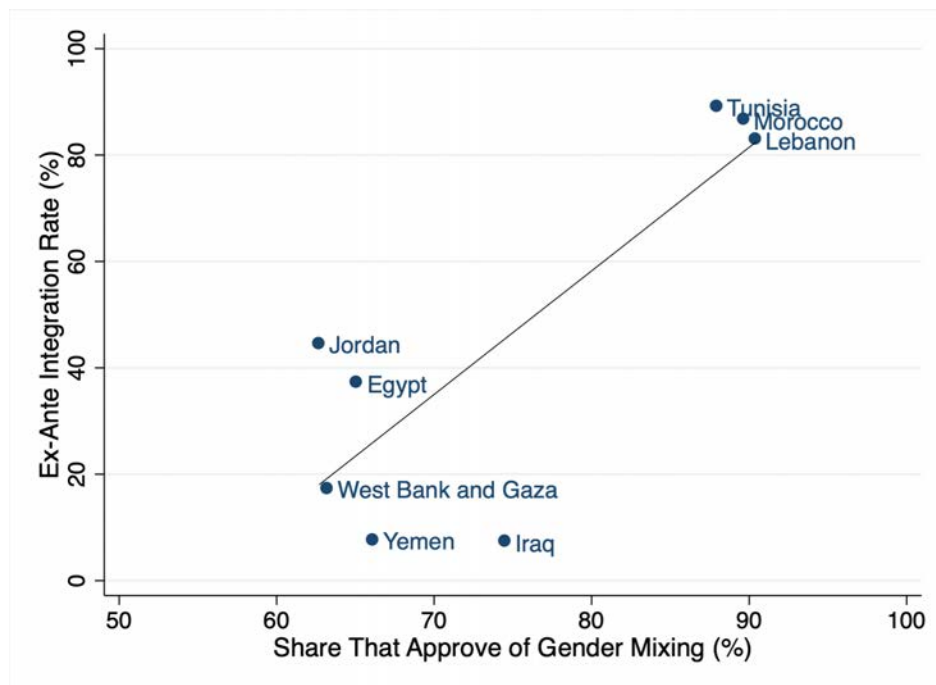
Note: This set of figures compares observed and simulated distributions of female employment across firms for six countries: Ethiopia, China, Russia, Brazil, Egypt, and India. The simulated distributions are simulated under the null hypothesis that no firm in that countries faces binding integration costs. Sample selection and simulation details are described in Sections VI.A.1 and VI.B.1.

FIGURE XIV
 INTEGRATION RATES BY $\bar{\theta}_j n_j$ AND COUNTRY




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 in six countries: Ethiopia, China, Russia, Brazil, Egypt, and India. We construct $\bar{\theta}_j$ for each firm j using $\hat{\theta}^{EP}(X_i)$ and the job mix of firm j . Ex-ante integration rates are constructed as described in Section IV.B.3.

FIGURE XV
 INTEGRATION RATES AND SUPPORT FOR GENDER MIXING IN MENA



Note: This figure plots ex-ante integration rates against local support for mixed-gender university classes across countries in the Middle East and North Africa (MENA). Ex-ante integration rates are derived using country-specific $\hat{\theta}^{EP}(X_i)$, estimates of $\theta(X_i)$ derived using ex-post integrated firms. We measure local support for gender mixing using the Arab Barometer survey. This survey is described in more detail in Section VI.A.2.

FIGURE XVI
 FEMALE LABOR FORCE PARTICIPATION AND EX-ANTE INTEGRATION RATES



Note: This figure plots female labor force participation rates against ex-ante integration rates for 65 countries. Ex-ante integration rates are derived using country-specific $\hat{\theta}^{EP}(X_i)$, estimates of $\theta(X_i)$ derived using ex-post integrated firms. Female labor force participation rates come from the World Bank Development Indicators.

TABLE I
MANUFACTURING FIRMS WITH ZERO FEMALE EMPLOYEES AND WORKFORCE COMPOSITION,
BY REGION

	All-male share of firms (%), by size		Female share (%)	
	Medium (20–99)	Large (100+)	Surveyed firms	Labor force
Sub-Saharan Africa	10.5	2.3	27.0	47.5
East Asia and Pacific	1.8	0.5	41.2	42.8
Eastern and Central Europe	2.5	0.7	38.4	43.9
Latin America and Caribbean	3.0	0.8	32.8	41.1
Middle East and North Africa	48.1	22.7	16.9	21.1
South Asia	49.9	28.6	14.5	23.5

Source: World Bank Enterprise Survey, 2006–2018. Survey data cover manufacturing firms in 65 countries. Appendix Table A1 lists the countries included in this tabulation. Statistics are calculated using survey weights within each country and year, then averaged across years within a country, then averaged across countries within a region, weighting by 2018 population. Female share of labor force is derived from 2018 World Bank Development Indicators for the same countries, is also a population-weighted average, and is not restricted to manufacturing.

TABLE II
FIRMS WITH SAUDI EMPLOYEES IN JANUARY 2009

	All	≥ 5 Saudi employees	
		All	Ex-post integrated
# of firms	27,294	7,943	2,123
Number of Saudi employees			
Mean	16.3	52.0	118.2
Median	2	12	20
SD	353	654	1254.3
Female share of employees (%)			
Mean	3.6	9.1	34.0
Median	0	0	25
SD	13.7	21.1	28.8
Avg. monthly wage (Riyals)	3,058	3,971	4,112
<i>Industry (%)</i> :			
Agriculture and fishing	0.9	1.0	0.4
Commerce	32.5	28.1	23.1
Community/social services	9.4	13.9	35.0
Construction	28.5	21.6	12.2
Electricity, gas, and water	0.7	1.4	1.0
FIRE	10.8	11.1	14.7
Manufacturing	13.3	17.1	10.7
Mining	1.0	1.6	0.8
Telecommunications	3.1	4.2	2.1

Notes: This table presents descriptive statistics for firms with any Saudi employee in January 2009. The second column limits to firms with at least five Saudi employees. We limit the analysis to firms with at least five Saudi employees throughout Section IV.A. The third column further limits to firms that are employment both men and women. The average wage at a firm is measured in nominal Saudi Riyals in January 2009.

TABLE III
SUMMARY OF θ ESTIMATES, JANUARY 2009

	Naive ($\hat{\theta}^0$)	Ex-post integrated ($\hat{\theta}^{EP}$)	Structural ($\hat{\theta}^S$)
Mean	0.082	0.125	0.123
Median	0.027	0.063	0.061
Standard deviation	0.155	0.180	0.177
<i>Pairwise R²:</i>			
Location	0.06	0.09	0.09
Occupation	0.73	0.69	0.43
Industry	0.62	0.60	0.62

Notes: This table summarizes three estimates for $\theta(X_i)$: (1) the “naive” estimate ($\hat{\theta}^0$), described in Section IV.A.1, which estimated using data from all firms; (2) the estimate using ex-post integrated firms ($\hat{\theta}^{EP}$); and (3) the structural estimates ($\hat{\theta}^S$), where the model and estimation are described in Section IV.B.2. Each estimated function is applied to all jobs in firms meeting the sample criteria described in Section IV.A. The *Pairwise R²* values are the R^2 values from separate linear regressions of the $\theta(X_i)$ estimates on location fixed effects, two-digit occupation fixed effects, and one-digit industry fixed effects.

TABLE IV
FIRM DESCRIPTIVE STATISTICS, BY BASELINE FIRM STATUS

	Baseline all male	Baseline integrated
# of firms	12,617	2,796
Number of Saudi hires		
Mean	61.9	219.7
Median	18	47
SD	194.6	930.9
Female share of hires (%)		
Mean	19.2	48.4
Median	10.0	46.3
SD	23.4	33.4
Avg. monthly wage (Riyals)	3,238	3,709
<i>Industry (%)</i> :		
Agriculture and fishing	1.0	0.3
Commerce	29.7	20.6
Community/social services	6.9	41.6
Construction	30.0	11.4
Electricity, gas, and water	0.8	0.9
FIRE	9.9	12.5
Manufacturing	16.3	10.4
Mining	1.4	0.6
Telecommunications	3.9	1.7

Notes: This table presents descriptive statistics for firms with any Saudi employee in January 2009 that hire at least five Saudis between February 2009 and June 2015. The first column includes firms that are all-male in January 2009. The second column includes firms that are ex-post integrated in January 2009. The average wage at a firm is measured in nominal Saudi Riyals at the time of hiring.

TABLE V
 BASELINE-SEGREGATED FIRMS IN JUNE 2011

	Below Quota	Above Quota
# of firms	2,224	1,559
Number of Saudi employees		
Mean	33.7	45.5
Median	12	16
SD	101	130
Female share of employees (%)		
Mean	2.2	1.7
Median	0	0
SD	8.6	7.0
Avg. monthly wage (Riyals)	3,680	4898
<i>Industry (%)</i> :		
Agriculture and fishing	1.4	0.8
Commerce	25.4	25.1
Community/social services	7.0	4.6
Construction	30.3	26.2
Electricity, gas, and water	1.2	1.6
FIRE	8.0	10.3
Manufacturing	20.1	22.4
Mining	1.7	3.7
Telecommunications	5.0	5.4

Notes: This table presents descriptive statistics as measured in June 2011 for firms with at least five Saudi employees in January 2009. The first column limits to Below Quota firms, those with Yellow and Red color statuses in June 2011. The second column limits to Above Quota firms, those with Green and Yellow color statuses in June 2011.

TABLE VI
 ABOVE VERSUS BELOW QUOTA
 DIFFERENCE-IN-DIFFERENCE ESTIMATES

Outcome:	Total hires	Integrated ($\times 100$)	Female share of hires (%)
Post \times Below Quota	2.88** (0.65)	10.68** (1.23)	2.26** (0.72)
Half-year FEs	✓	✓	✓
Firm FEs	✓	✓	✓
Observations	48,963	48,963	32,997

Notes: This table reports OLS coefficient estimates for Equation (5), a firm-level difference-in-difference model for number of hires, integration rates, and the female share of hires. Each observation reflects a firm by half-year pair. Column (3) excludes periods where a firm makes no hires. Firms are limited to those with at least five Saudi employees. *Post* is an indicator for half-years after H1 of 2011. Robust standard errors are clustered at the firm level.
 ~ significant at 10% level; * significant at 5% level; ** significant at 1% level.

TABLE VII
FEMALE LABOR FORCE PARTICIPATION AND INTEGRATION RATES ACROSS COUNTRIES

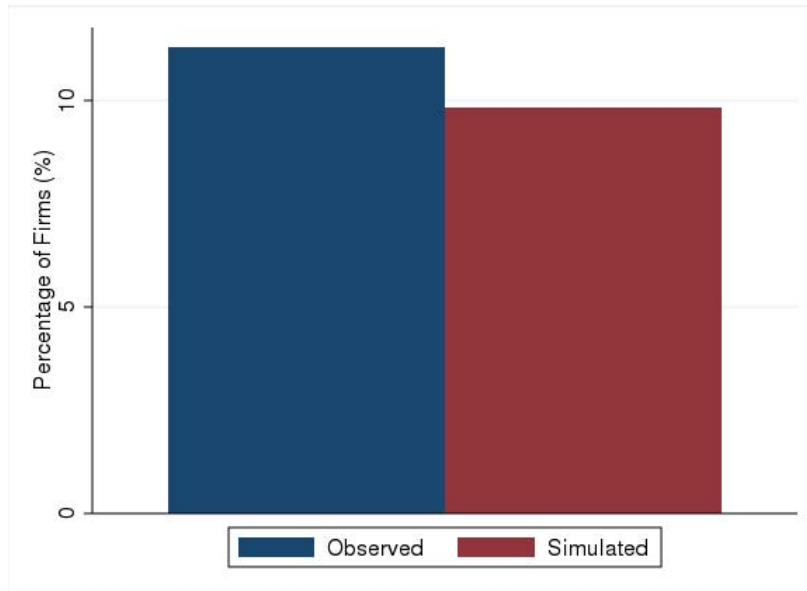
	LFP_F		$LFP_F - LFP_M$	
	(1)	(2)	(3)	(4)
<i>Ex-ante integration rate:</i>				
Overall ($\hat{\theta}^{EP}$)	0.425** (0.079)	0.278** (0.089)	0.453** (0.067)	0.298** (0.075)
Overall ($\hat{\theta}^S$)	0.615** (0.133)	0.287** (0.142)	0.638** (0.117)	0.292* (0.123)
Representative firm ($\bar{\theta}_j^{EP}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

Notes: This table reports the OLS estimates of a regression of labor force participation measures on ex-ante integration rates. LFP_F is the percentage of women over the age of 15 in the labor force in the years the manufacturing surveys were conducted, averaged across years. $LFP_F - LFP_M$ is the difference between female and male labor force participation in the same years. The top panel reports the results from a regression of labor force participation rates on average ex-ante integration rates in surveyed firms, where ex-ante integration rates are derived using ex-post integrated firms in Section IV.B.3. The middle panel uses the structural estimates of overall ex-ante integration rates as in Section IV.B.2. The bottom panel reports the results from a regression on predicted ex-ante integration rates in each country for a “representative” firm with $\bar{\theta}_j \times n_j = 10$ (i.e., a firm with ten expected female employees if integrated), where $\bar{\theta}_j$ is derived using $\hat{\theta}^{EP}(X_i)$. Columns (2) and (4) report results for regressions that include region fixed effects. Robust standard errors are in parentheses.


~ significant at 10% level; * significant at 5% level; ** significant at 1% level.

FIGURE A1
 DISTRIBUTION OF MARRIED MALE EMPLOYMENT ACROSS FIRMS

(a) Percentage of Firms with Zero Married Male Employees

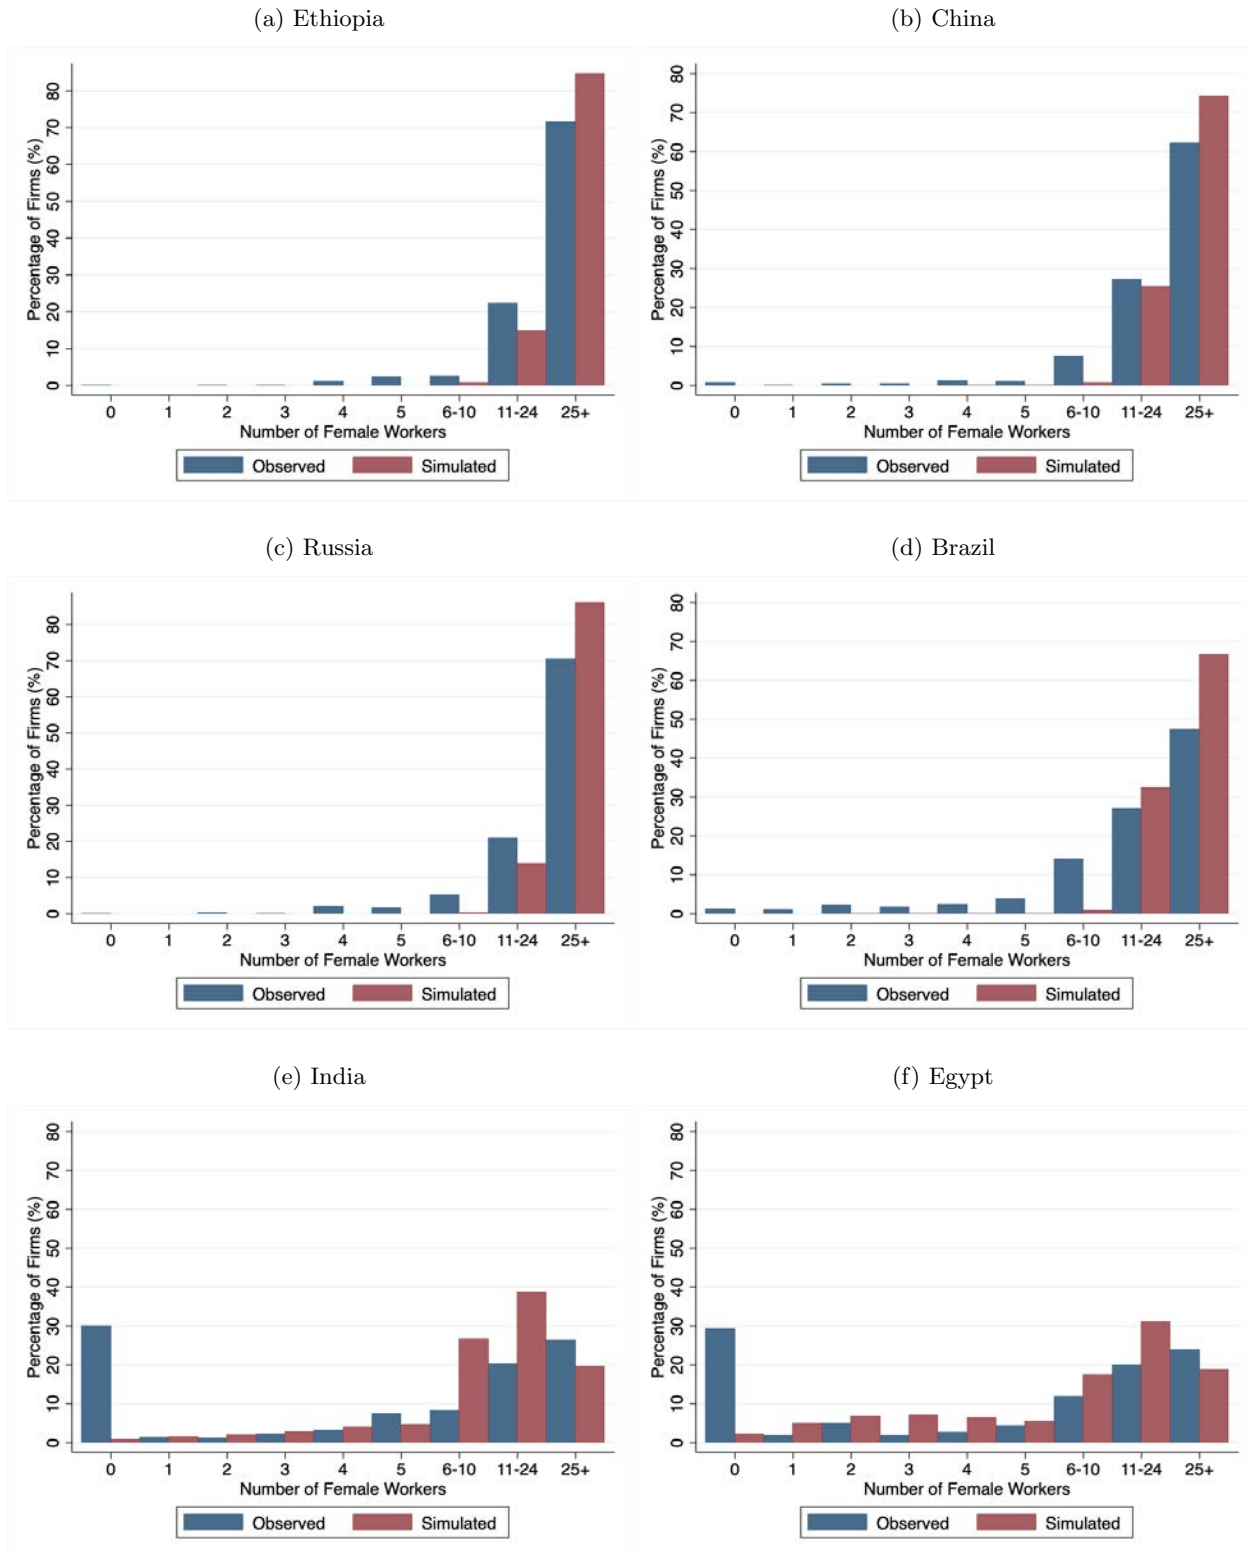


(b) Percentage of Firms with > 0 Married Male Employees




Note: This set of figures compares the observed distribution of married male employment across firms in January 2009 to distributions simulated under the null hypothesis that no firm faces binding integration costs. Sample selection and simulation details are described in Sections IV.A.1 and IV.A.2. Panel A plots the share of firms with zero married male employees in both the observed and simulation distributions. Panel B plots the share of firms with various nonzero totals of married male in employees in both the observed and simulated distributions. For all simulations, a Kolmogorov-Smirnov test fails to reject equality of the observed and simulated distributions at the 0.01 significance level.

FIGURE A2
 DISTRIBUTION OF FEMALE EMPLOYMENT ACROSS *Large* FIRMS, BY COUNTRY




Note: This set of figures compares observed and simulated distributions of female employment across firms for six countries: Ethiopia, China, Russia, Brazil, Egypt, and India. We limit to firms with at least 50 employees. The simulated distributions are simulated under the null hypothesis that no firm in that countries faces binding integration costs. Sample selection and simulation details are described in Sections VI.A.1 and VI.B.1.

FIGURE C1
 SIMULATED DISTRIBUTION OF FEMALE EMPLOYMENT FOR $\gamma = 0.7$ AND $\theta_a = 0.5$




Note: This figure plots the distribution of the count of female employees across firms based on 1,000 simulations of firm sizes, integration probabilities (γ) and the share of female labor in the workforce (θ_a). The H_0 category supposes that all firms are integrated ($\gamma = 1$), and the H_a category supposes that some firms are ex-ante segregated ($\gamma < 1$).

FIGURE C2
 SHARE OF SIMULATIONS WITH $Z_a > Z_0$ BY θ_a AND γ




Note: This heatmap plots the share of simulations with $Z_a > Z_0$, or the share of simulations where there are more firms with no female employees under H_0 vs. H_a while varying values of θ and γ .

FIGURE C3
 SIMULATED DISTRIBUTION OF FEMALE EMPLOYMENT FOR $\bar{\gamma}_i = 0.7$ AND $\theta_a = 0.5$ –
 INTEGRATION RATES INCREASING IN n




Note: This figure plots the distribution of the count of female employees across firms based on 1,000 simulations of firm sizes, integration probabilities (γ) and the share of female labor in the workforce (θ_a) when firm integration probabilities correlate positively with firm size. The H_0 category supposes that all firms are integrated ($\gamma = 1$), and the H_a category supposes that some firms may still be segregated ($\gamma < 1$).

FIGURE C4
 SHARE OF SIMULATIONS WITH $Z_a > Z_0$ BY θ_a AND γ – INTEGRATION RATES INCREASING IN n



Note: This heatmap plots the share of simulations with $Z_a > Z_0$, or the share of simulations where there are more firms with no female employees under H_0 vs. H_a while varying values of θ and γ and when firm integration probabilities correlate positively with firm size.

FIGURE E1
SHARE HIRED IN MONTH WHO CHANGE STATUS



Note: This figure plots the share of employees who are first hired in each month who change wage or occupation, switch firms, or earn above minimum wage within two years of hire. Dashed lines show the 95% confidence interval for month indicator variables.

FIGURE E2
 SHARE HIRED IN MONTH WHO CHANGE STATUS (WITH WORKER-LEVEL CONTROLS)



Note: This figure plots the share of employees who are first hired in each month who change wage or occupation, switch firms, or earn above minimum wage within two years of hire when controlling for employee characteristics. Indicator variables are used to flexibly control for age, education, and marital status of new hires. Dashed lines show the 95% confidence interval for month indicator variables.

TABLE A1
LIST OF WORLD BANK ENTERPRISE SURVEYS

Country	Region	Year	# of Firms	% Female	Country	Region	Year	# of Firms	% Female
Argentina	LAC	2006	559	23.8	Lao	EAP	2018	119	34.5
Argentina	LAC	2010	703	18.9	Lebanon	MNA	2013	174	21.0
Argentina	LAC	2017	571	20.8	Madagascar	AFR	2009	185	50.5
Armenia	ECA	2009	108	35.3	Mexico	LAC	2006	1060	35.9
Azerbaijan	ECA	2009	118	39.2	Mexico	LAC	2010	1065	31.5
Azerbaijan	ECA	2013	107	31.8	Mongolia	EAP	2009	126	56.7
Bangladesh	SAR	2007	1160	46.1	Mongolia	EAP	2013	106	51.7
Bangladesh	SAR	2013	1073	46.1	Morocco	MNA	2013	120	45.8
Belarus	ECA	2013	110	44.0	Myanmar	EAP	2014	314	58.2
Bosnia-Herzegovina	ECA	2009	112	36.3	Nepal	SAR	2009	122	12.7
Bosnia-Herzegovina	ECA	2013	103	37.3	Nepal	SAR	2013	231	14.9
Bolivia	LAC	2006	340	28.0	Nicaragua	LAC	2006	292	19.0
Bolivia	LAC	2010	106	20.2	Pakistan	SAR	2007	640	1.5
Botswana	AFR	2006	101	45.3	Panama	LAC	2006	223	27.3
Brazil	LAC	2009	1205	34.3	Panama	LAC	2010	105	33.0
Bulgaria	ECA	2007	501	52.2	Paraguay	LAC	2006	351	28.3
Chile	LAC	2006	602	25.7	Paraguay	LAC	2010	107	33.8
Chile	LAC	2010	755	19.0	Peru	LAC	2006	337	45.0
China	EAP	2012	1597	39.8	Peru	LAC	2010	715	25.4
Colombia	LAC	2006	588	53.0	Peru	LAC	2017	508	29.9
Colombia	LAC	2010	665	50.4	Philippines	EAP	2009	846	44.3
Colombia	LAC	2017	481	43.3	Poland	ECA	2009	108	43.7
Costa Rica	LAC	2010	285	22.5	Poland	ECA	2013	123	24.7
Croatia	ECA	2007	303	40.5	Romania	ECA	2009	135	44.7
Croatia	ECA	2013	109	42.0	Romania	ECA	2013	157	24.0
Dominican Republic	LAC	2010	109	27.2	Russia	ECA	2009	540	45.3
DRC	AFR	2006	128	10.1	Russia	ECA	2012	1106	42.6
DRC	AFR	2010	100	11.0	Serbia	ECA	2009	119	33.7
Ecuador	LAC	2006	336	24.2	Serbia	ECA	2013	105	40.8
Ecuador	LAC	2010	114	25.8	South Africa	AFR	2007	619	30.1
Egypt	MNA	2013	1535	11.4	Sri Lanka	SAR	2011	345	42.5
Egypt	MNA	2016	1063	13.0	Sweden	ECA	2014	277	22.8
El Salvador	LAC	2006	384	48.2	Tajikistan	ECA	2008	102	39.3
El Salvador	LAC	2010	121	44.0	Trinidad and Tobago	LAC	2010	110	29.1
El Salvador	LAC	2016	336	39.1	Tunisia	MNA	2013	280	43.1
Ethiopia	AFR	2011	218	44.6	Turkey	ECA	2008	699	27.7
Ethiopia	AFR	2015	340	37.0	Turkey	ECA	2013	872	23.0
Georgia	ECA	2008	104	38.9	Uganda	AFR	2006	254	20.7
Guatemala	LAC	2006	266	32.1	Uganda	AFR	2013	267	24.1
Guatemala	LAC	2010	326	30.3	Ukraine	ECA	2008	381	47.0
Guatemala	LAC	2017	118	33.3	Ukraine	ECA	2013	537	43.4
Honduras	LAC	2006	221	29.0	Uruguay	LAC	2006	315	38.9
Honduras	LAC	2010	111	28.4	Uruguay	LAC	2010	303	31.4
India	SAR	2014	6282	11.6	Uzbekistan	ECA	2008	116	37.8
Indonesia	EAP	2015	978	39.4	Uzbekistan	ECA	2013	116	32.8
Iraq	MNA	2011	377	1.6	Vietnam	EAP	2009	716	47.5
Israel	MNA	2013	170	29.0	West Bank and Gaza	MNA	2013	123	4.1
Jordan	MNA	2013	238	13.2	Yemen	MNA	2010	191	3.4
Kazakhstan	ECA	2009	147	40.8	Zambia	AFR	2007	276	12.2
Kazakhstan	ECA	2013	153	28.5	Zambia	AFR	2013	283	13.9
Kenya	AFR	2007	373	15.5	Zimbabwe	AFR	2011	332	16.2
Kenya	AFR	2013	338	19.0	Zimbabwe	AFR	2016	262	21.1
Kenya	AFR	2018	269	15.5					

Note: This table lists the World Bank Enterprise Surveys that we include in our analysis. We limit our samples to manufacturing firms, where surveys include questions on the gender composition of employees by occupation. Next, we drop surveys where information on gender composition is missing for more than 20% of firms. In remaining surveys, we drop firms with missing data on gender composition or fewer than 5 employees. We then drop surveys with fewer than 100 remaining firms. This leaves us with 105 surveys in 65 countries. The abbreviations used are: Latin America and the Caribbean (LAC), Europe and Central Asia (ECA), East Asia and the Pacific (EAP), Middle East and North Africa (MNA), Sub-Saharan Africa (SSA), and South Asia (SAR).

TABLE A2
COMPOSITION OF PRIVATE SECTOR

	Saudi		Non-Saudi	
	Male	Female	Male	Female
<i>Share of workforce</i>				
2005	11.0%	0.6%	86.9%	1.5%
2010	9.6%	0.8%	88.4%	1.3%
2015	11.7%	4.7%	79.2%	2.2%
<i>Occupational distribution among group in 2015</i>				
Managers	8.0	8.1	0.7	0.5
Professionals	5.6	6.3	7.7	13.6
Technicians	8.0	12.5	7.7	29.8
Clerical	23.8	39.8	0.6	1.3
Sales	9.8	20.5	5.5	0.8
Service	24.1	7.5	29.3	44.8
Agriculture	0.2	0.1	6.6	0.1
Industrial, Chemical, and Food Industries	1.8	1.6	2.5	8.4
Engineering Support	16.0	3.3	39.3	0.4
Armed Forces and Security	2.6	0.2	0.1	0.1
Other	0.0	0.0	0.0	0.3

Note: The top half of this table tabulates the distribution of private sector workers by year. The second half tabulate the occupational distribution of each subgroup of private sector workers in 2015. Numbers exclude domestic workers. Source: Saudi Ministry of Labor and Social Development (MLSD).

TABLE A3
SAUDI WORKERS SUMMARY STATISTICS,
JANUARY 2009

	Men	Women
Share of employment (%)	91.6	8.4
Age	32.1 (10.1)	30.3 (7.6)
Married	24.6	32.9
<i>Education level (%)</i>		
Less than Secondary	5.5	4.1
Secondary	40.0	42.9
University	5.6	32.6
Missing	48.9	20.4
Monthly Wage (Riyals)	7206 (8016)	3308 (4178)

Source: General Organization for Social Insurance (GOSI) administrative data. Data include only Saudi nationals. More data details are discussed in Section III.D.

TABLE A4
EMPLOYEES BY ISCO-08 OCCUPATION, JUNE, 2011

ISCO Code	ISCO Category	Frequency	Percent
96	Refuse workers and other elementary workers	104,744	14.4
41	General and keyboard clerks	84,406	11.6
54	Protective services workers	65,032	9.0
42	Customer services clerks	64,265	8.9
99	Unclassified	48,382	6.7
33	Business and administration associate professionals	36,547	5.0
52	Sales workers	32,943	4.5
74	Electrical and electronic trades workers	26,754	3.7
83	Drivers and mobile plant operators	25,296	3.5
21	Science and engineering professionals	23,465	3.2
	Total	511,834	70.5

Note: This table presents the number of Saudi employees in the ten most common ISCO-08 2-digit occupation group. The large number of unclassified occupations is due to the significantly large number of cases where the GOSI occupation verification process was still processing or was incomplete.

TABLE A5
FEMALE EMPLOYMENT AND INTEGRATION RATES ACROSS COUNTRIES

	EMP_F		$EMP_F - EMP_M$	
	(1)	(2)	(3)	(4)
<i>Ex-ante integration rate:</i>				
Overall ($\hat{\theta}^{EP}$)	0.388** (0.083)	0.243** (0.097)	0.420** (0.065)	0.264** (0.073)
Overall ($\hat{\theta}^S$)	0.572** (0.137)	0.262~ (0.152)	0.582** (0.113)	0.237~ (0.119)
Representative firm ($\bar{\theta}_j^{EP} \times n_j = 10$)	0.461** (0.085)	0.321** (0.101)	0.459** (0.069)	0.324** (0.075)
Region FEs		✓		✓
Observations	65	65	65	65

Notes: This table reports the OLS estimates of a regression of labor force participation measures on ex-ante integration rates. EMP_f is the percentage of women over the age of 15 that are employed in the years the manufacturing surveys were conducted, averaged across years. $LFP_F - LFP_M$ is the difference between female and male employment rates in the same years. The top panel reports the results from a regression of employment rates on average ex-ante integration rates in surveyed firms, where ex-ante integration rates are derived using ex-post integrated firms in Section IV.B.3. The middle panel uses the structural estimates of overall ex-ante integration rates as in Section IV.B.2. The bottom panel reports the results from a regression on predicted ex-ante integration rates in each country for a “representative” firm with $\bar{\theta}_j \times n_j = 10$ (i.e., a firm with ten expected female employees if integrated), where $\bar{\theta}_j$ is derived using $\hat{\theta}^{EP}(X_i)$. Columns (2) and (4) report results for regressions that include region fixed effects. Robust standard errors are in parentheses.

~ significant at 10% level; * significant at 5% level; ** significant at 1% level.