

Missing Women, Integration Costs, and Big Push Policies in the Saudi Labor Market

By CONRAD MILLER AND JENNIFER PECK AND MEHMET SEFLEK*

In settings where social norms promote gender segregation, firms may find it costly to employ both men and women. These integration costs may hinder women's employment. We develop a methodology to test for the presence of fixed integration costs and estimate counterfactual women's employment at all-male firms where these costs bind. We apply our approach in Saudi Arabia and find that integration costs bind for the majority of firms. We show that Nitaqat, a gender-neutral quota program that incentivized the hiring of Saudi nationals at private sector firms, induced firms to integrate and dramatically increased Saudi women's employment.
JEL: J16, J23, J71, O53

Women's employment rates are particularly low in the Middle East and North Africa (MENA) and South Asia.¹ These low employment rates are attributed in part to social norms regarding gender roles and their effects on labor supply decisions (e.g., Fernandez and Fogli, 2009; Alesina, Giuliano and Nunn, 2013). In these regions, strong preferences for gender segregation are common (Jayachandran, 2015). Saudi Arabia is a particularly extreme example: women made up only 15.9 percent of the Saudi labor force in 2018 (GaStat, 2011). Saudi cultural norms around gender segregation are particularly strict and, until recently, labor regulations explicitly mandated gender-segregated workplace facilities. Increasing women's employment, however, has recently been a key economic and social goal for the Kingdom: since 2011, a series of ambitious labor reforms has corresponded with rapid growth of women's employment in the Saudi private sector. Some of these policies have been consequential for women's employment even when not gender-specific. The most significant of these reforms was Nitaqat, a gender-neutral quota program designed to increase the number of Saudi nationals working in the private sector. The program, launched in 2011, has corresponded

* Miller: Haas School of Business, University of California, Berkeley and NBER, ccmiller@berkeley.edu. Peck: Swarthmore College, jpeck1@swarthmore.edu. Seflek: Haas School of Business, University of California, Berkeley, mseflek@berkeley.edu. We thank Hadar Avivi and Aaron Metheny for excellent research assistance. We thank Peter Blair, Patricia Cortes, Alessandra Gonzalez, Supreet Kaur, Pat Kline, Jonathan Kolstad, Peter Kuhn, David Levine, Jonathan Leonard, Jeremy Magruder, Steve O'Connell, 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, Stanford (SIEPR), and Texas A&M 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.

¹The female share of the labor force ranges from a high of 46.4% in sub-Saharan Africa to 20.3% and 24% in the Middle East and North Africa (MENA) and South Asia (World Bank, 2018b).

with a dramatic increase in the number of Saudi women working in the private sector: from just 56,000 in 2010 to 606,000 by the end of 2017, increasing the female share of the Saudi private sector workforce from 8% to 32%.

We argue that this dramatic transformation occurred in part because gender-neutral Nitaqat quotas led firms to overcome firm-level barriers to gender integration, where integration is defined here as employing both men and women. These barriers are the largely fixed costs of accommodating social norms and regulations that require a physically and socially segregated workplace. Workers, customers, and regulators may expect firms to establish gender-segregated facilities, including restrooms, entrances, and workspaces. Firms may also segregate tasks or teams to limit interactions between male and female employees. For male-dominated firms, hiring women may necessitate changes in their workplace culture. These social expectations constrain the production process for integrated firms. We study the consequences of these integration costs for Saudi women's employment and argue that Nitaqat pushed firms to integrate by incentivizing them to hire more Saudis, reducing the per-worker burden of the fixed costs. Paying these fixed integration costs may be difficult to justify when hiring only a small number of workers, but become less onerous when hiring a large number of Saudis to meet hiring quotas.

The distribution of women's employment across firms provides *prima facie* evidence that integration is costly. In January 2009, our first month of data, 73% of private sector firms with at least 5 Saudi employees employ no Saudi women.² This is substantially larger than what one would expect by chance, even with the low female share of employment.³ This strongly suggests that firms face an extensive margin decision of whether to integrate their workforce.

Motivated by a simple model of firm hiring, we develop a methodology that uses the distribution of women's employment across firms to assess whether and how integration costs constrain women's employment at firms. We apply and validate the methodology using administrative data on Saudi citizens' employment in the private sector. We find that the majority of Saudi firms employ only men because they face binding integration costs. We then show that Nitaqat led to a dramatic increase in female share of the Saudi workforce at least in part by inducing firms to integrate.

We first build a 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 both male and female candidates for all vacancies. To hire from both pools, the firm must pay a fixed integration cost to accommodate social preferences for gender segregation. In the Saudi context,

²We describe these data and sample restrictions we apply in more detail in Section II.D.

³In January 2009, 8% of Saudis in the private sector are women. If the gender of each employee were independent draws from a binomial distribution where the probability an employee is female is 0.08, the share of all-male firms that would occur by random chance is 34%.

these costs include establishing legally required gender-segregated facilities, for example. 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.

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 . We apply our test to social security data covering Saudi nationals in the private sector from 2009 to 2015.⁴ In our data, X_i includes occupation, industry, and the location of the job. To test the null hypothesis that all firms are ex-ante integrated, we estimate $\theta(X_i)$ using data on employees at all firms, simulate the distribution of women’s employment across firms using this estimate, and compare the simulated and observed distributions. Under the null, we would expect some firms to have zero female employees by chance alone. However, if integration costs bind for *some* firms, we show that we should see an excess mass or “bunching” of firms with zero female employees. We 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 women. We simulate that 43% of firms in our sample should employ only men under the null hypothesis. In practice, 73% of firms in our sample employ only men.

We next estimate $\theta(X_i)$ using data from integrated firms only. 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.

A key concern with our approach is that an estimate of $\theta(X_i)$ based on integrated firms may not produce accurate estimates for counterfactual employment of women at segregated firms. With sufficient unobserved heterogeneity in candidate pools or preferences across firms, the observed distribution of female employment—including a mass of all-male firms—can be rationalized in the absence of any integration costs at all. Moreover, it is not clear that the relationship between a firm’s ex-ante integration status and its expected number of female employees under integration is *causal*, as our model implies. We address these concerns using 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

⁴A limitation of these data is that they do not cover non-Saudi workers, who make up a substantial share of the private sector workforce. In particular, some firms that we identify as all male might in fact employ non-Saudi women. However, non-Saudi women are rare in the private sector, and firms that employ both Saudi nationals (a requirement for inclusion in our data) and non-Saudi women tend to also employ Saudi women. We discuss this issue in more detail in Section II.D.

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 women. 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. 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 the exogenous variation in Saudi employment across firms generated by the Nitaqat nationalization quotas. As shown in Peck (2017), firms generally responded to Nitaqat quotas by employing more Saudis. Firms vary in their distance from their quota at baseline, generating exogenous variation in Saudi hiring across firms. We use this variation in hiring incentives 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 hiring. Consistent with the model, we also find that (previously all-male) below quota firms integrate at higher rates and have a larger female share of hires than above quota firms. 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 women’s 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 women’s wages and employment. We interpret the aggregate effects of Nitaqat as evidence that integration costs reduce aggregate women’s employment in Saudi Arabia. In particular, 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 suggest that Nitaqat increased relative aggregate demand for women by inducing more firms to integrate.

The workplace segregation we document is consistent with employees preferring to work with coworkers of the same gender as in Becker (1957). In our framework, firms can pay a fixed cost to accommodate these preferences and employ both men and women, e.g. by establishing gender-segregated facilities or teams. A distinguishing prediction of our model is that ex-ante integration rates are increasing in a firm’s total hires. This prediction is borne out in the data. Moreover, following a firm’s first female hire, we do not see the wages or separation rates of incumbent male employees increase as the logic of the Becker (1957) model of coworker discrimination would suggest if workplace segregation were driven by male preferences for male coworkers. Women’s preferences for female coworkers could generate our findings if their utility is highly nonlinear in their number of female coworkers and, for example, women particularly dislike working in firms where they are the sole female employee.⁵

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 like Nitaqat that incentivize firms to integrate can substantially change firm hiring preferences at the margin.⁶ 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 women’s 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.⁷

Related Literature.—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, Lelarge and Reenan (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 based on observed bunching at zero in the distribution of female employment across firms. While in traditional approaches the presence of bunching is often visually clear from a density plot alone, a key challenge in our setting is that identifying bunching requires a model of counterfactual female employment at firms. Our setting requires this structure because fixed integration costs imply excess mass in the number of firms with zero female employees, a corner solution where we may expect a

⁵In this case, one can think of the cost of a “cluster hire” as an approximately fixed integration cost.

⁶In the literature on racial discrimination in the US, there is evidence that affirmative action policies and other shocks to minority hiring can have long-term effects on minority employment even after the policies end (Miller and Segal, 2012; Miller, 2017; Whatley, 1990).

⁷Policies that increase exposure to integrated workplaces may also change gender attitudes, which could in turn affect integration costs (Dahl, Kotsadam and Rooth, 2018).

mass of firms even in the absence of fixed integration costs. Our simulation-based approach to constructing a counterfactual is similar to Ellison and Glaeser (1997) and Augereau, Greenstein and Rysman (2006). Our methodology can be applied to measure integration costs on the basis of other worker characteristics including disability (Acemoglu and Angrist, 2001), language (Lang, 1986; Hellerstein and Neumark, 2008), ethnicity (Hjort, 2014; Glover, Pallais and Pariente, 2017), and religion.

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 and how policy interacts with social norms (e.g., Fernandez, 2013; Ashraf et al., 2019). Most closely related is Bursztyn, Gonzalez and Yanagizawa-Drott (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.

We also contribute to the literature on workplace segregation and its implications for labor market inequality. As some firms pay more than others, this segregation can have important implications for gender earnings inequality (Groshen, 1991; Bayard et al., 2003; Card, Cardoso and Kline, 2016). While prior research has shown that skill differences, occupational preferences (Goldin, 1986), and gender-based perceptions of prestige (Pan, 2015) 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.

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. A series of papers document that firms tend to change employment in a manner consistent with nonconvex 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.

I. A Model of Firm Hiring

In this section we describe a simple, partial equilibrium model of an individual firm’s hiring strategy to study the implications of integration costs for women’s employment. The 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 a fixed number of applications from two types of candidates: type

F and 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 type I extreme value distribution with scale parameter β .⁸ β indexes how much candidate quality varies within group. 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 , e.g., the cost of establishing gender-segregated facilities. We assume the firm must choose their hiring strategy prior to observing their candidates.

Define U_*^M , U_*^F , and U_*^I as the expected value of the highest U_j value among group M candidates, group F candidates, and the combined pool, respectively. 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

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

As we show in Appendix B, the left hand side of this expression can be expressed as

$$(2) \quad \begin{aligned} U_*^I - U_*^M &= -\beta \log[1 - \theta] \\ &\approx \beta\theta. \end{aligned}$$

where θ denotes the probability that the firm's preferred candidate from the combined pool is type F . Combining (2) with (1), an approximate condition for the firm to pay the fixed integration cost and hire from the combined pool is

$$(3) \quad n\theta > \frac{c}{\beta}.$$

The left-hand side of (3) is the firm's expected number of type F hires if it were

⁸The CDF is $F(\epsilon_j) = \exp(\exp(-\epsilon_j/\beta))$. It follows that $Var(\epsilon_j) = \frac{\beta^2 \pi^2}{6}$ and $E(\epsilon_j) = \beta\gamma$, where γ is Euler's constant.

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 in our data.

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 women’s employment. If turnover rates are similar for men and women—as we show they are in the Saudi private sector—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.

II. 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 the counterfactual employment of women at segregated firms. We apply our methodology using administrative data from Saudi Arabia. In this section, we describe the Saudi context and data.

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 was 56.8% (World Bank, 2018*b*); 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.⁹ 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.¹⁰

⁹Even 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).

¹⁰Unemployed women with college degrees outnumbered men by almost four to one in 2008.

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 (Table 2). Even when controls are added for education, location, and occupation, women earn about 40% less than men in January 2009.¹¹

The Saudi private sector is composed primarily of male, non-Saudi expatriate workers (see Tables 1 and 2). Among Saudis, women form a small share of private sector employment, with Saudi women making up 8.5% of Saudi employees in January 2009. Low employment of women 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, women’s 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.

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. 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.¹²

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 (whether Saudi or non-Saudi) 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

¹¹A 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).

¹²These integration costs are sometimes explicitly cited when discussing obstacles to women’s 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).

Table 1—: 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) via Saudi Arabian Monetary Agency (SAMA) (Saudi Arabian Monetary Agency, 2019).

Development (MLSD) and the Ministry of Municipal and Rural Affairs (Khoja and Thomas, 2018). 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.

Even outside of legal requirements, Saudi firms may incur similar costs to accommodate social preferences for gender segregation. Moreover, integration may be costly if employees prefer to work with coworkers of the same gender regardless of accommodations as in Becker (1957). In that case, integrated firms may have to increase compensation to offset the disutility of working with members of a different gender. A key prediction of the model outlined in Section I that is not present in the Becker (1957) model of coworker discrimination is that ex-ante integration rates are increasing in a firm's total hires. Moreover, if segregation is driven by male preferences for male coworkers, we should see the wages or sepa-

Table 2—: 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 (General Organization for Social Insurance, Kingdom of Saudi Arabia, 2015). Data include only Saudi nationals in the private sector. These statistics describe 1,396,962 Saudis employed in the private sector in January 2009. Section II.D provides more details about the dataset used for this table.

ration rates of incumbent male employees increase following a firm's first female hire. We test this prediction in Section III.C. Women's preferences for female coworkers could generate predictions similar to our model if their utility is highly nonlinear in their number of female coworkers. For example, if women particularly dislike working in firms where they are the sole female employee, firms may integrate by hiring a 'cluster' of 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 per-

mission when recruiting female workers.¹³ 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.¹⁴¹⁵ Addressing these HR issues involves learning 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)¹⁶ 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.¹⁷ Some firms address this by providing group transportation for their employees, a lumpy, ongoing cost.

C. *Nitaqat Nationalization Quotas*

We analyze the Nitaqat quota policy, which generated exogenous variation in Saudi employment across firms, through the lens of the model. The Nitaqat program is an ongoing gender-neutral nationalization quota policy first instituted in 2011.^{18,19} The policy was designed to address growing national unemployment, which in 2011 had reached 40% for Saudis in the 2025 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. These quotas varied by firm industry and size and assigned firms to

¹³The 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, 2018a).

¹⁴Some of these adjustments are mandated for firms above a particular size: Saudi 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.

¹⁵Addressing workplace harassment is an ongoing issue in labor markets with high levels of gender integration as well as an important barrier to further integration in markets with low integration. The current literature suggests that there is significant firm-level heterogeneity in the prevalence of harassment, and firms can take steps to create an organizational culture that prevents harassment (Willness, Steel and Lee, 2007; Hart, Crossley and Correll, 2018). Developing the procedures for reporting and responding to complaints and creating credible messaging from leadership about organizational values can be thought of as an investment in both prevention and firm reputation. This likely affects recruiting as well as retention.

¹⁶Engineering, for example, was not offered to Saudi women as an undergraduate degree program until 2005.

¹⁷Women were not permitted to drive in Saudi Arabia until June 2018.

¹⁸See Peck (2017) for a more detailed description of the Nitaqat program and its effects.

¹⁹Other contemporary labor policies are described in Appendix C.

four color bands according to their level of compliance: firms in the Green and Platinum bands were in-compliance with nationalization quotas, while firms in the Red and Yellow bands were required to increase their share of Saudi workers.²⁰ 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. Compliance was monitored using a system integration social security data for Saudi workers and visas for expatriate workers from the National Information center. Firms in the Red and Yellow bands faced restrictions on their ability to renew existing visas, obtain new visas, and access the MLSA’s foreign recruitment services; Green and Platinum firms were given access to a streamlined visa renewal service. The MLSA 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 non-compliance were phased in, starting just three months later in September 2011.

Overall, Nitaqat quotas were effective at increasing Saudi employment in the private sector (Peck, 2017), with firms complying with the program by increasing their Saudi employment. 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 2013.²¹

D. Data

We test for integration costs in the Saudi context using administrative social security data from the General Organization for Social Insurance (GOSI) (General Organization for Social Insurance, Kingdom of Saudi Arabia, 2015). These data contain information on all Saudis employed in the private sector between January 2009 and June 2015. They record 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 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 D. In total, the GOSI dataset contains information on approximately 2.8 million unique individuals and 430,000 firms. We restructure this dataset into an unbalanced monthly panel for each full-time Saudi employee and standardize the occupations using two-digit codes from the International Labour Organization’s ISCO-08 codes.

²⁰Cutoffs for each band were set based on pre-Nitaqat Saudization rates so that slightly less than half of firms in each industry by size group would be classified as Green or Platinum.

²¹*Saudi Gazette*, September 9, 2012, “Nitaqat percentage linked to salary.”

We also test our model using firm responses to Nitaqat Saudi employment quotas. We use the Nitaqat data (Ministry of Labor and Social Development, Kingdom of Saudi Arabia, 2015) 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.²² The details of merging the two data sets are described in more detail in Appendix D.

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 E 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, our GOSI data do not include information on non-Saudi workers, so our references to the composition of workers throughout the paper refer only to Saudi employees. This means that firms that we identify as "all-male" may in fact employ non-Saudi women. Other government data suggest that this group is small: of the 11.5 million people employed in 2015 only 773,000, or 7 percent, were non-Saudi women. Of these, 94 percent worked in private households as domestic workers or in public sector jobs in education and health. Less than one percent of workers outside households, health, and education were non-Saudi women (GaStat, 2015). In our analysis we consider a firm segregated if it reports employing Saudi men but not Saudi women in the GOSI dataset. Across 2012 and 2013 there are 2.1 million such firms in the Nitaqat data, only 1.5 percent of which reported employing non-Saudi women. Firms that employ non-Saudi women are also very likely to also employ Saudi women: among all firms that employed both non-Saudi women and Saudis in either 2012 or 2013, 75 percent employed Saudi women. Among firms that employed both non-Saudis (women or otherwise) and Saudis, only 39 percent employed Saudi women. This suggests that the integration costs associated with employing either Saudi or non-Saudi women likely have substantial overlap.

²²The 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.

III. 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.

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 measure ex-ante integration rates as a function of expected female employment under integration. We also test our central assumption in several ways described below.

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.²³ Under the null, 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.²⁴

While the model is framed in terms of firm hiring, the test we first develop uses cross-sectional, firm-level data on women's *employment*. If turnover rates are similar for men and women—as they are in Saudi Arabia—then a firm's female share of hires will equal its female share of employees.²⁵ Otherwise, the female share of employees will equal the duration-weighted female share of hires. We conduct a similar test in Section III.C that examines hiring rather than employment.

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

²³A special case would be that integration costs are zero (i.e., do not exist).

²⁴By contrast, if some firms *do* face binding integration costs, we show in Appendix F 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.

²⁵In Saudi Arabia, turnover rates are similar for Saudi men and women. The 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.

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.²⁶ 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. Table 3 compares the characteristics of all firms and those with at least five Saudi employees. The two sets of firms have comparable industry compositions.

Table 3—: 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 III.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.

²⁶In other words, this restriction reduces the number of firms that are potentially ex-ante integrated but ex-post segregated.

We estimate a logistic regression model of the form

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

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 4 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 gender across positions.

Table 4—: 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 III.A, 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 G.1. Each estimated function is applied to all jobs in firms meeting the sample criteria described in Section III.A. The *Pairwise R²* values are the *R²* 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.

SIMULATION RESULTS. — Next, we simulate the distribution of women’s employment across firms using our estimate, $\hat{\theta}^0(X_i)$, and compare the result to the distribution we observe.²⁷

²⁷In 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 women’s employment in Figure 1. 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.



(a) Percentage of Firms with Zero Female Employees (b) Percentage of Firms with > 0 Female Employees

Figure 1. : Distribution of Female Employment across Firms

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 III.A and III.A. 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 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 1% significance level.

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 the 1% significance level.²⁸ Overall, the pattern is consistent with binding integration costs at many firms.

However, our simulated distribution may also fail to match the observed distribution because we have misspecified $\theta(X_i)$. Relative to traditional bunching approaches, misspecification is a particularly important concern in our context because, even in the absence of fixed integration costs, we expect a mass of all-male firms. Hence, the presence of a mass of all-male firms alone does not imply

²⁸We compare the distributions of the number of female employees, not the binned data presented in Figure 1.

an ‘excess’ mass. $\theta(X_i)$ is misspecified if there may be 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 gender 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. 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 . In Panel A of Appendix Table A.1 we compare the observed distribution to the simulated distribution for several specifications of $\theta(X_i)$, where we vary the set of job characteristics we include in X_i . 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. While adding controls for job location makes little difference, adding controls for 2-digit occupation increases this value to 39%. Adding 1-digit industry increases this value further to 43%. This value is unchanged if we replace occupation and industry fixed effects with 1-digit occupation by 1-digit industry interactions. Hence, while the simulated percentage of firms that are all male depends somewhat on the job characteristics we include in constructing $\theta(X_i)$, including all of our observable job characteristics brings the value nowhere close to the observed value, 73%. Unobservable job characteristics would need to be highly influential relative to observable job characteristics to match the distribution of female employment in the data.²⁹

We conduct more direct tests of our specification of $\theta(X_i)$ in Sections III.C and III.D.

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

²⁹The reasoning behind this argument is similar to that of Altonji, Elder and Taber (2005) and Oster (2019).

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 $\theta(X_i)$. The bias is small in practice 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. Second, we estimate a more parametric model that accommodates ex-ante integration as a potentially unobserved firm state.

We first limit the data to ex-post integrated firms when estimating $\theta(X_i)$. Table 3 summarizes the characteristics of ex-post integrated firms, using the same sample restrictions described in Section III.A. 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 4 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%.

The second approach we take is to directly model the distinction between ex-ante and ex-post integrated firms and to structurally estimate $\theta(X_i)$. We jointly estimate $\theta(X_i)$ and the probability that a firm is ex-ante integrated as a function of firm characteristics. Estimation details are provided in Appendix G. Column (3) of Table 4 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 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 π_j is 0.65, indicating 65% of firms are ex-ante segregated.

APPLYING ESTIMATES OF $\theta(X_i)$. — We next examine how integration rates relate to a firm’s expected number of female employees if ex-ante integrated. 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.³⁰ 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.

Panel A of Figure 2 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

³⁰For this construction, we assume that a firm’s job mix does not depend on its integration status.

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 2 estimates this chance factor.

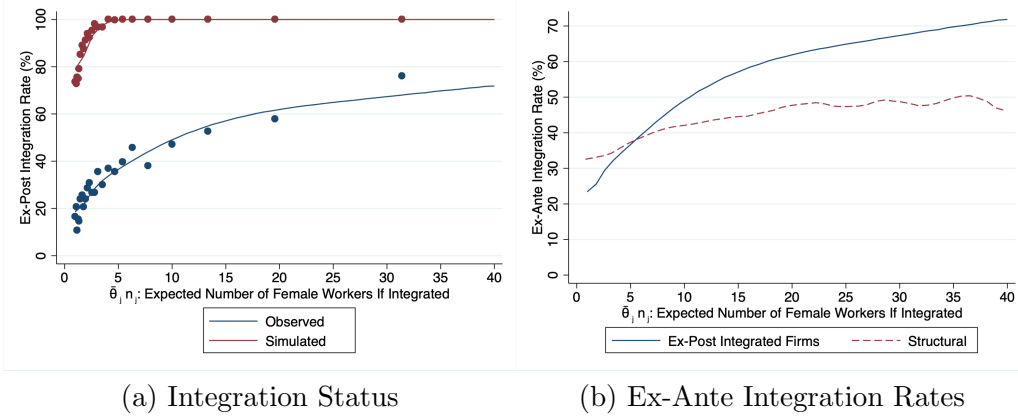


Figure 2. : Integration Rates by $\bar{\theta}_j n_j$

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 III.B. “Structural” plots the average estimated values of π_j (described in Appendix G) 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 . The sample includes a cross-section of firms from 2009 with at least five Saudi employees. Sample selection and simulation details are described in Sections III.A and III.A.

We use the relationships illustrated in Panel A of Figure 2 to estimate *ex-ante* integration rates as a function of $\bar{\theta}_j n_j$. The ratio of the actual and simulated ex-post integration rates provides an estimate of the ex-ante integration rate, as a function of $\bar{\theta}_j n_j$.³¹ Panel B of Figure 2 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

³¹To 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 2 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 $\hat{\theta}^{EP}(X_i)$.

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 $\bar{\theta}_j n_j$.

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. These simulations are described in more detail in Appendix G. Panel B of Appendix Table A.1 compares the simulated distribution of female employment to the observed distribution for various specifications of $\theta(X_i)$. Our baseline (and most saturated) specification, where X_i includes job location, two-digit occupation, and one-digit industry fixed effects, matches the observed distribution. Across all simulations, we fail to reject equality of the simulated and observed distributions in a Kolmogorov-Smirnov test at the 1% significance level. The average p-value is 0.75. 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.

C. Using Panel Structure to Validate the Model

The panel structure of the Saudi data permits additional tests of the model. In particular, panel data allow us to further probe our key assumption that $\theta(X_i)$ dictates counterfactual employment of women for ex-ante segregated firms. In Appendix G 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 1 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.

We find that firm transitions from segregated to integrated are consistent with the model. Prior to integrating, we observe transitioning firms in the GOSI data for an average of 39 months. Six months after a newly integrated firm's first female hire, 26% of their hires are women. This share 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.

Second, we test for *state dependence*: the hiring behavior of firms that have already paid their integration costs should 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 can 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.

For baseline all-male firms, the contrast between the simulated and observed distributions of female employment is similar to that shown in Figure 1. By contrast, the simulated distribution for baseline integrated firms matches the observed distribution relatively well. 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.

We also examine how incumbent male employees respond to their employer’s first female hire. If between-firm gender segregation were purely driven by employee preferences for same-gender coworkers as in Becker (1957), then we should observe an increase in separation rates or in wages for these employees to compensate for the disutility they incur when working in an integrated firm. In Appendix Figure A.1 we plot average monthly separation rates and wages for male employees that joined transitioning firms more than 18 months prior to their first female hire. Neither outcome appears to meaningfully respond to a firm’s first female hire. This suggests that gender segregation is not driven by the preferences of male employees, or that firms can accommodate these preferences while employing both men and women, e.g. by establishing gender-segregated facilities or teams.

D. Using Policy Variation to Verify the Threshold Rule

We next test whether the Nitaqat employment quotas induce firm integration and increase hiring of women 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.

Panel A of Figure 3 plots the average number of Saudi hires each half-year, separately for below and above quota firms. Prior to Nitaqat, hiring 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. 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. We use this policy-induced variation in hiring (n) to test the predictions of the model.

In particular, we test whether Nitaqat increases (1) integration rates and (2) the female share of hires at below quota firms relative to above quota firms. 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.

Appendix Table A.2 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.

First, 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 B of Figure 3. 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.

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, pooling both firms that do and do not integrate. There is no mechanical reason that an increase in total hires would increase the female share of hires. We plot the female share of hires in Panel B of Figure 3. 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—24 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 III.C. Averaging across post-Nitaqat hires within firms, the average estimated value of $\hat{\theta}_j$ is 0.25.

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

$$(4) \quad Y_{jt} = \alpha_j + \tau_t + \beta \text{Post}_t \times \text{Below}_j + \epsilon_{jt},$$

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_i is an indicator for a 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 (4) for the same three outcomes: total hires, integration status, and female share of hires.

The DD estimates are overlaid on Figure 3. 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.

IV. Aggregate Effects of Integration Costs: Evidence from Nitaqat

The results above do not consider what would happen in the aggregate if integration costs were eliminated in the labor market. 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 women's 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 discuss these aggregate effects in more detail in Appendix H.

To assess the aggregate consequences of integration costs, we would ideally use exogenous variation in integration costs across labor markets. Lacking such variation, we examine the labor market response to Nitaqat, which induced many firms to integrate and hire women (see Appendix Figure H.1). Following the introduction of Nitaqat in 2011, the female share of the Saudi private sector workforce increased from 10% to 27% in 2015. This increase is concentrated in firms that were previously all-male. We also find that this increase in the female share of Saudis in the private sector is not offset by a decrease in the public sector: the female share of Saudis working in the public sector instead increases over this period, from 33% in 2011 to 40% in 2015.

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 the increase in demand for labor bids up wages for males, then we may expect a demand increase for relatively cheaper female labor. We find, however, that the gender wage gap *decreases* following Nitaqat. Moreover, the establishment of the effective minimum wage in 2013 reduces the wage gap even further, and the female share of the workforce remains elevated (see Appendix Figure H.2). The fact that both women's 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.

Lastly, it is also possible 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). One reason for this low rate may be that

households perceive that few firms are willing to hire women in the first place. Since we do not have data on labor supply decisions, we look at the response to Nitaqat for firms that had integrated prior to the policy's implementation. These firms are already employing a mix of men and women and face an increase in the relative price of female labor due to Nitaqat. In the absence of a supply response, we would expect to see the female share of employment at these firms weakly decreasing. Instead, for *both* sets of firms, there is a marked increase in the female share of employment beginning with Nitaqat's integration, implying that there may have been an outward shift in women's labor supply.³²

V. 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 also find evidence that integration costs depress aggregate demand for female labor. In particular, we document that Nitaqat—a gender-neutral policy in Saudi Arabia that had the unintended consequence of inducing many firms to integrate—increased women's employment and wages.

Integration costs may seem particularly likely to exist and bind in Saudi Arabia, a country with uniquely explicit restrictions on between-gender interactions. However, there is suggestive evidence that integration costs also bind in other countries with strong social norms for gender segregation. In particular, in recent World Bank survey data (The World Bank, 2019) on manufacturing firms, 50% of medium firms and 25% of large firms in MENA and South Asia are all male, a far higher share of firms than one would expect by chance alone (see Appendix Table A.3). 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.³³ Integration costs may therefore constrain women's employment more broadly across the world, providing an important barrier to growth as well as an important driver of global employment dynamics.

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

³²Appendix Figure H.3 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.

³³These data are from World Bank Enterprise Surveys (The World Bank, 2019)

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: big-push policies like Nitaqat could have large equilibrium-switching effects, and complementary policies that address labor supply may help magnify these impacts.

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? Investments in overcoming integration costs may have longer-term implications for productivity in addition to employment.

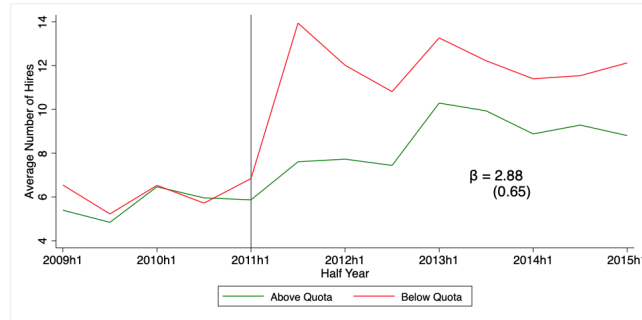
REFERENCES

- Acemoglu, Daron, and Joshua D. Angrist.** 2001. "Consequences of Employment Protection? The Case of the Americans with Disabilities Act." *Journal of Political Economy*, 109(5): 915–957.
- Alesina, Alberto, Paola Giuliano, and Nathan Nunn.** 2013. "On the Origins of Gender Roles: Women and the Plough." *Quarterly Journal of Economics*, 128(2): 469–530.
- Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber.** 2005. "An Evaluation of Instrumental Variable Strategies for Estimating the Effects of Catholic Schooling." *Journal of Human Resources*, 40(4): 791–821.
- Ashraf, Nava, Natalie Bau, Nathan Nunn, and Alessandra Voena.** 2019. "Bride Price and Female Education." *Journal of Political Economy*.
- Augereau, Angélique, Shane Greenstein, and Marc Rysman.** 2006. "Coordination versus Differentiation in a Standards War: 56K Modems." *RAND Journal of Economics*, 37(4): 887–909.
- Bayard, Kimberly, Judith Hellerstein, David Neumark, and Kenneth Troske.** 2003. "New Evidence on Sex Segregation and Sex Differences in Wages from Matched Employee-Employer Data." *Journal of Labor Economics*, 21(4): 887–922.
- Becker, Gary.** 1957. *The Economics of Discrimination*. University of Chicago Press.
- Bursztyn, Leonardo, Alessandra L. Gonzalez, and David Yanagizawa-Drott.** 2018. "Misperceived Social Norms: Female Labor Force Participation in Saudi Arabia."


- Card, David, Ana Rute Cardoso, and Patrick Kline.** 2016. "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women." *Quarterly Journal of Economics*, 131(2): 633–686.
- Dahl, Gordon B., Andreas Kotsadam, and Dan-Olof Rooth.** 2018. "Does Integration Change Gender Attitudes? The Effect of Randomly Assigning Women to Traditionally Male Teams." Unpublished manuscript.
- Ellison, Glenn, and Edward L. Glaeser.** 1997. "Geographic Concentration in U.S. Manufacturing Industries: A Dartboard Approach." *Journal of Political Economy*, 105(5): 889–927.
- Evidence for Policy Design.** 2015. "Back to Work in a New Economy: Background Paper on the Saudi Labor Market." Harvard Kennedy School, Harvard University.
- Fernandez, Raquel.** 2013. "Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century." *American Economic Review*, 103(1): 472–500.
- Fernandez, Raquel, and Alessandra Fogli.** 2009. "Culture: An Empirical Investigation of Beliefs, Work, and Fertility." *American Economic Journal: Macroeconomics*, 1(1): 146–177.
- Fry, Erika.** 2015. "This CEO is helping Saudi women break a gender barrier." *Fortune*. September 14, 2015.
- Garicano, Luis, Claire Lelarge, and John Van Reenan.** 2016. "Firm Size Distortions and the Productivity Distribution: Evidence from France." *American Economic Review*, 106(11): 3439–79.
- GaStat.** 2011. *Labor Force Survey 2011*. General Authority for Statistics, Kingdom of Saudi Arabia.
- GaStat.** 2015. *Labor Force Survey 2015*. General Authority for Statistics, Kingdom of Saudi Arabia.
- General Organization for Social Insurance, Kingdom of Saudi Arabia.** 2015. "GOSI Social Security Payment Database 2009-2015." (Last accessed 2015-05-15).
- Glover, Dylan, Amanda Pallais, and William Pariente.** 2017. "Discrimination as a Self-Fulfilling Prophecy: Evidence from French Grocery Stores." *Quarterly Journal of Economics*, 132(3): 1219–1260.
- Goldin, Claudia.** 1986. "Monitoring Costs and Occupational Segregation by Sex: A Historical Analysis." *Journal of Labor Economics*, 4(1): 1–27.

- Gourio, François, and Nicolas Roys.** 2014. "Size-dependent regulations, firm size distribution, and reallocation." *Quantitative Economics*, 5(2): 377–416.
- Groshen, Erica L.** 1991. "The Structure of the Female/Male Wage Differential: Is It Who You Are, What You Do, or Where You Work?" *Journal of Human Resources*, 26(3): 457–472.
- Hamermesh, Daniel S., and Gerard A. Pfann.** 1996. "Adjustment Costs in Factor Demand." *Journal of Economic Literature*, 34(3): 1264–1292.
- Hart, Chloe Grace, Alison Dahl Crossley, and Shelley J. Correll.** 2018. "Leader Messaging and Attitudes toward Sexual Violence." *Socius*, 4.
- Hellerstein, Judith K., and David Neumark.** 2008. "Workplace Segregation in the United States: Race, Ethnicity, and Skill." *The Review of Economics and Statistics*, 90(3): 459–477.
- Hjort, Jonas.** 2014. "Ethnic Divisions and Production in Firms." *Quarterly Journal of Economics*, 129(4): 1899–1946.
- Jayachandran, Seema.** 2015. "The Roots of Gender Inequality in Developing Countries." *Annual Review of Economics*, 7: 63–88.
- Khoja, Sara, and Sarit Thomas.** 2018. "Labour and Employment Compliance in Saudi Arabia." In *International Labour and Employment Compliance Handbook*, ed. del Rey Salvador and Robert J. Mignin. Kluwer Law International BV, Netherlands.
- Kleven, Henrik.** 2016. "Bunching." *Annual Review of Economics*, 8: 435–464.
- Kuhn, Peter, and Kailing Shen.** 2013. "Gender Discrimination in Job Ads: Evidence from China." *Quarterly Journal of Economics*, 128(1): 287–336.
- Lang, Kevin.** 1986. "A Language Theory of Discrimination." *Quarterly Journal of Economics*, 101(2): 363–382.
- Miller, Amalia R., and Carmit Segal.** 2012. "Does Temporary Affirmative Action Produce Persistent Effects? A Study of Black and Female Employment in Law Enforcement." *Review of Economics and Statistics*, 94(4): 1107–1125.
- Miller, Conrad.** 2017. "The Persistent Effect of Temporary Affirmative Action." *American Economic Journal: Applied Economics*, 9(3): 152–90.
- Ministry of Labor and Social Development, Kingdom of Saudi Arabia.** 2015. "Nitaqat Administrative Dataset 2011-2015." (Last accessed 2015-05-15).
- Oster, Emily.** 2019. "Unobservable Selection and Coefficient Stability: Theory and Evidence." *Journal of Business and Economic Statistics*, 37(2): 187–204.


- Pan, Jessica.** 2015. "Gender Segregation in Occupations: The Role of Tipping and Social Interactions." *Journal of Labor Economics*, 33(2): 365–408.
- Peck, Jennifer R.** 2017. "Can Hiring Quotas Work? The Effect of the Nitaqat Program on the Saudi Private Sector." *American Economic Journal: Economic Policy*, 9(2): 316–347.
- Sara Hamdan.** 2012. "Saudi Arabia Signals Openness to Women Seeking Work." *New York Times*. September 5, 2012.
- Saudi Arabian Monetary Agency.** 2019. "Annual Statistics 2019." (Last accessed 2020-07-15).
- Saudi Gazette.** 2012. "Nitaqat percentage linked to salary." *Saudi Gazette*. September 9, 2012.
- The World Bank.** 2019. "Enterprise Surveys, 2006-2018." (Last accessed 2019-03-19).
- Varejão, José, and Pedro Portugal.** 2007. "Employment Dynamics and the Structure of Labor Adjustment Costs." *Journal of Labor Economics*, 25(1): 137–165.
- Whatley, Warren C.** 1990. "Getting a Foot in the Door: "Learning," State Dependence, and the Racial Integration of Firms." *Journal of Economic History*, 50(1): 43–66.
- Willness, Chelsea R., Piers Steel, and Kibeom Lee.** 2007. "A Meta-Analysis of the Antecedents and Consequences of Workplace Sexual Harassment." *Personnel Psychology*, 60(1): 127–162.
- World Bank.** 2018a. *Women, Business and the Law 2018*. The World Bank.
- World Bank.** 2018b. *World Development Indicators 2018*. The World Bank.



(a) Number of Hires



(b) Ex-Post Integration Rates



(c) Female Share of Hires

Figure 3. : Integration Rates and Female Share of Hires Over Time

Note: This set of figures compares the number of hires, ex-post integration rates, and 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. There are 2,224 below quota firms and 1,559 above quota firms satisfying our sample criteria. The vertical line marks the first half of 2011. Nitaqat is implemented in June 2011. Panel A plots the average number of hires by half-year. Panel B plots the share of firms that are ex-post integrated by half-year. Panel C 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 of the female share of hires for that period. Each panel includes OLS coefficient estimates for equation (4), 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. Robust standard errors are clustered at the firm level.