

# Big Push Policies and Firm-Level Barriers to Employing Women: Evidence from Saudi Arabia

**PRELIMINARY AND INCOMPLETE**

Conrad Miller<sup>1</sup>, Jennifer Peck<sup>2</sup> and Mehmet Seflek<sup>\*1</sup>

<sup>2</sup>Department of Economics, Swarthmore College

<sup>1</sup>Haas School of Business, UC Berkeley

February 19, 2019

We study firm demand for female labor in Saudi Arabia, a country where laws and norms regulate interactions between men and women and female labor force participation is among the world's lowest. In 2011, Saudi Arabia imposed gender-neutral quotas for Saudi employment at private sector firms. We find this policy more than tripled the female share of Saudis working in the private sector within four years, with this increase concentrated at firms that were previously all-male. Motivated by these findings, we examine whether the potential fixed costs that Saudi firms must pay to employ both men and women affect firm personnel decisions. Fixed integration costs can potentially lead to multiple equilibria for female employment: firms only integrate if enough women are in the labor force, and women only join the labor force if enough firms can employ women. Consistent with fixed integration costs, we find that: firms are highly segregated with few firms employing small numbers of women; the distribution of female hiring across firms is similarly 'lumpy', and this pattern is concentrated at previously all-male firms; employment quotas lead to a disproportionate and persistent increase in female hiring; and this increase is concentrated at previously all-male firms. We also find that the quota policy induces integration at firms where quotas were not binding, consistent with demand externalities or changing norms. Our results suggest an important role for 'big push' demand-side policies in increasing female employment.

## 1 Introduction

While female labor force participation (FLFP) has increased substantially over the last half century, FLFP remains low in many parts of the world. The global labor force participation rate is 75.2% for working-age

---

<sup>\*</sup>We thank seminar participants at Berkeley Haas (BPP), Harvard (WiGi and Growth Lab), IMF, 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.

men, yet only 48.7% for women (World Bank, 2016). One common explanation for low FLFP is that social norms or culture prevent or discourage women from working outside the home. This research primarily focuses on how social norms and culture affect women's labor *supply* decisions (see e.g. Goldin (2014) and Bertrand (2011)), but the same factors may influence labor *demand*. In many labor markets, firms face laws or cultural norms that regulate interactions between men and women in the workplace or limit the types of work women can perform. For example, in many countries, women are prohibited from working at night and discouraged from commuting alone due to safety concerns (International Finance Corporation, 2013; World Bank, 2018). In order to hire female employees ('integrate'), previously all-male firms may need to build separate restrooms or other facilities for women and implement female-oriented human resource policies regarding maternity leave, childcare, and workplace harassment. Gender integration may also require costly changes in workplace organization and culture. How firms respond to these potential costs may have important implications for female labor market outcomes.

In this paper we study firm demand for female labor in Saudi Arabia, a country that strictly limits between-gender interactions in the workplace. Like many countries in the Middle East and North Africa, FLFP in Saudi Arabia is among the lowest in the world, 17.8 percent in 2011 (GaStat, 2011). At the same time, women in Saudi Arabia have experienced a transformation in their economic engagement over the last decade. An ambitious gender-neutral nationalization quota policy (called *Nitaqat*) launched in 2011 incentivized firms to hire more Saudi workers and many firms responded by hiring women, often for the first time. Between 2011 and 2015 female employment in the Saudi private sector increased over sixfold, and much of this growth was driven by the integration of previously all-male firms.

We argue that this striking employment response is largely the result of firms overcoming *firm-level* barriers to integration in response to the quota policy. These barriers are particularly salient in the Saudi context: labor regulations mandate separate offices and facilities for men and women, limit between-gender face-to-face interactions, and, for female employees, limit direct oversight of work by male managers. Women may also find working at a firm undesirable in the absence of female coworkers. These constraints introduce costs with a significant *fixed* component, which we label as *integration costs*—fixed costs that firms must pay to employ any women at all.<sup>1</sup> Paying these costs may be difficult to justify when employing only a small number of workers, but become less of a barrier when employing a large number of Saudis to meet employment quotas.

To examine this phenomenon we develop a simple search model based on Black (1995) to study the implications of fixed integration costs for female employment. On the demand side, firms only hire women if they expect to spread fixed integration costs over a sufficient number of female employees. Hence, a gender-neutral demand shock may induce some firms to integrate. On the supply side, women only enter the labor market if enough firms have integrated. This feedback loop can generate a demand externality—a firm's decision to integrate may induce other firms to integrate by increasing the supply of women

---

<sup>1</sup>Saudi firms often cite these integration costs when discussing obstacles to female employment. One business owner tells 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 their male and female employees (Fortune 2015).

searching in the labor market.

We test the model using administrative employer-employee data. The model generates four predictions that we test in our data: (1) the distribution of female workers across firms exhibits bunching at zero and a ‘missing middle’ of firms that employ small numbers of female employees; (2) the distribution of female *hiring* is similarly ‘lumpy’, and this pattern is concentrated at previously all-male firms; (3) quotas increase relative labor demand for women at all-male firms, but (4) do not affect relative demand at integrated firms.

We find support for all four predictions in the data. First, we find that Saudi female workers are substantially segregated across firms. We compare the actual distribution of female employment across firms to a simulated benchmark, where workers in the same labor market (as defined by location and occupation) are allocated randomly across firms. While our simulation predicts that 36% of firms would have zero female employees in 2009, we find that in fact 72% of firms had zero female employees. This bunching at zero female employees is offset by a ‘missing middle’ mass of firms with only a handful female employees.

Second, we find that the distribution of female hiring across firms following Nitaqat is also lumpy, particularly for firms that are segregated at baseline. Our simulation predicts that about 13% of baseline segregated firms will hire zero female employees after the implementation of Nitaqat, while 36% of firms hire no women in practice. Again, this bunching is offset by a missing mass of firms with small numbers of female hires. By contrast, the extent of a missing middle is substantially more muted for firms that are integrated at baseline.

Third, we show that the introduction of gender-neutral employment quotas led to a dramatic increase in the Saudi female employment in the private sector. Consistent with the model, we find that this increase was concentrated at firms that were further below the quota and segregated at the time the quotas were introduced. Among similar segregated firms, we find that firms that are further below the quota (and hence must hire more Saudis to meet the quota) integrate at faster rates. Fourth, among *integrated* firms, the rate at which firms hire women is unrelated to their distance from the quota.

Finally, we find a sharp increase in female hiring at segregated firms where quotas are not binding. These firms do not appear to be hiring more Saudis in general under Nitaqat. Instead, the behavior of these firms suggests they are hiring women due to a labor *supply* response. This is despite the fact that the gender wage gap is *reduced* following Nitaqat. This pattern is consistent with *demand externalities*—Nitaqat induces some firms to integrate via quotas, but there may be spillover effects to other firms as integration at some firms induces more women to search in the labor market. The relative availability of female labor may also increase because male labor supply is depleted, and not through a female supply response *per se*. The response from firms where quotas are not binding are also consistent with more direct firm-to-firm spillovers, where firms integrate in response to their peers integrating, perhaps due to changes in social norms. Unfortunately, we do not have a way to distinguish between these channels empirically.

The notion that gender integration involves substantial fixed costs has important implications for policy. In particular, our results suggest that ‘big push’ demand-side policies can substantially change firm hiring preferences and increase female labor force participation. Moreover, though we cannot test this directly here, our results 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 workspace space and facilities, organizational change—have a significant sunk component.

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 [Bursztyn et al. \(2018\)](#), who study social norms over women’s labor supply in Saudi Arabia. They show that men underestimate the share of their peers that support FLFP, and provide some evidence that correcting those misperceptions increases married men’s willingness to let their wives join 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. \(2018\)](#), who study how bride price traditions mediate educational investments in daughters’ education and household responsiveness to education policy.

This paper engages a 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 U.S. 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 ([Groschen, 1991](#); [Bayard et al., 2003](#); [Card et al., 2016](#)). There is little work explaining why some firms integrate and some firms do not. However, various explanations for female earnings differences are implicitly explanations for between-firm segregation as well, including skill differences and discriminatory employer or co-worker preferences. For example, [Pan \(2015\)](#) documents that the gender composition of occupations display ‘tipping points’, where occupations like bank tellers and typesetters quickly transitioned from being predominantly male to predominantly female. This pattern is consistent with men preferring to work in occupations with a higher male share of co-workers.

We also 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 non-convex 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 fixed, potentially one-time adjustment costs.

Finally, our paper relates to the literature on the impact of legislation like the Americans with Disabilities Act. This literature focuses on the employment impacts of the introduction of the ADA, and finds that costs associated with ADA compliance shifted the employment of disabled workers toward large firms and away from small and medium firms. [Acemoglu and Angrist \(2001\)](#) find that the ADA was associated with the largest negative employment impacts in the smallest firms that were subject to the regulations. [Hotchkiss and Rovba \(2003\)](#) find a similar shift in employment away from small and medium firms and toward large firms. They speculate that these patterns may be due to the ability of larger employers to spread fixed accommodation costs across a larger number of employees and the ability of larger and more prof-

itable employers to absorb these costs. This observation is consistent with our results on how employment and hiring patterns respond to fixed costs for subgroups of workers.

The remainder of the paper proceeds as follows. We describe the Saudi labor market in more detail in Section 2, including trends in FLFP, constraints women face in the labor market, and the Nitaqat employment quotas we exploit to study firm labor demand. In Section 3 we develop a simple search model and derive testable predictions. In Section 4 we describe our data. In Section 5 we test each model prediction using the data. Section 6 concludes.

## 2 Background

### 2.1 Women in the Saudi Workforce

Female employment in Saudi Arabia is low by international standards, with an employment rate of 16 percent among women age 15 or older before the 2011 labor reforms (World Bank, 2016). Although the labor force participation rate is also low, it has grown considerably over the last decades, up from 16 percent in 2000 to 20 percent in 2014.<sup>2</sup> Women’s employment has not kept pace with this growth in participation, leading to high unemployment among women: the unemployment rate for Saudi women was 33 percent at baseline in 2011, compared with 7 percent for men (GaStat, 2011). This is particularly true at higher education levels, and female jobseekers with bachelor’s degrees outnumber men by six to one (Figure 1). This is likely partially due to women’s traditional reliance on public sector work and low engagement in the private sector: even in 2014 women overwhelmingly worked in the public sector, with 74 percent of employed women working in girls’ schools in 2014 (Evidence for Policy Design, 2015).

[Figure 1 about here.]

Women appear to have trouble finding private sector work. Although women were 15 percent of the overall Saudi workforce, they formed only 10 percent of the Saudi workforce in the private sector. Correspondingly most private sector firms don’t employ women: in June 2011 86 percent of Nitaqat firms had no female employees. This was true even for larger firms, with 62 percent of firms with at least ten Saudi employees employing no Saudi women. This is likely due to a variety of factors on both sides of the market. Female employment in the public sector 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). Low female employment in the private sector likely 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.

### 2.2 Firm-level Costs of Employing Women

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

---

<sup>2</sup>Overall world female labor force participation is 50 percent; U.S. female labor force participation is 60 percent.

employment: Saudi Arabia is 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 costs of complying with other gender-based legal restrictions around women's work. Many of these costs are what we call "fixed" costs of employing any female workers, and do not depend on the number of female workers that firms employ. These include one-time switching costs (or integration 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.

### 2.2.1 Fixed Costs

The fixed costs of employing women are particularly striking in the Saudi context. In particular, it may be costly for firms to comply with government regulations regarding gender segregation in the workplace. The government requires that a firm that employs women must provide them separate workstations, a private space for women to pray and take breaks, convenient restroom access, and a separate entrance to the building or workplace. Meeting rooms have also been subject to changing guidance, and mixed-gender meetings must currently be visible to the rest of the office. Employing women exposes firms to inspections and potentially fines through the Ministry of Labor and Social Development and the Ministry of Municipal and Rural Affairs.<sup>3</sup> 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 beginning to employ female workers.

Although not always legally mandated, these types of investment costs are present in other contexts in which firms begin employing women for the first time. The World Bank's International Finance Corporation reports case studies from India, South Africa, and Thailand in which firms invested in separate facilities and overhauled uniform designs in order to begin attracting female workers (*International Finance Corporation, 2013*). As in these cases, Saudi firms may experience additional costs in providing a workplace environment that is acceptable to female workers. The IFC report also describes many other measures firms can take to attract female workers, including providing parental leave, offering flexible hours, and facilitating access to childcare. The report also suggests ways in which firms might provide more inclusive HR policies. Many of these may involve fixed costs, including learning about and implementing HR best-practices in these areas. Some of these are mandated for firms above a particular size: Saudi labor law requires firms that employ more than fifty women with at least ten children under age six must provide childcare. Firms with more than 100 women must provide a childcare center. Implementing these policies are associated with significant investment costs for the firm but can have permanent positive effects on recruiting and retaining female employees (see e.g. *Wooten (2001)*).

Firms may also face costs due to the lower mobility of female employees.<sup>4</sup> Some firms address this by providing group transportation for their employees, making this a fixed but ongoing cost. Although

---

<sup>3</sup>In 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*).

<sup>4</sup>Women were not permitted to drive in Saudi Arabia until June 2018.

particularly acute in Saudi Arabia, the issue of mobility is a common one for women globally. There are sixteen 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).

Historically low engagement 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. It may be difficult for firms to find female jobseekers with the desired skills, either because these skills aren't as common among women or because firms lack information on how to find these jobseekers. There are likely fewer professional women in their employee referral networks, for example, if their staff is all-male. It may also be more difficult for firms to assess the qualifications of female applicants, as women attend separate schools and tend to have shorter work histories. Furthermore, firms in Saudi Arabia 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.<sup>5</sup> Addressing these 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<sup>6</sup> or years of experience, that disqualify many female applicants. Overcoming these barriers may require firms to think flexibly about how they structure their tasks across occupations within the firm. The 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 PM. 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).

We can also think of fixed integration costs as including cultural shifts, both within firms and in the market more broadly, that make it easier for women to join and remain in the workforce. These long-run effects can include societal attitudes toward women's work, fertility, norms around work-family policies, and expectations about gender roles and childcare.<sup>7</sup>

---

<sup>5</sup>The guardianship requirement was lifted by the Saudi Ministry of Labor in 2008. There are still 18 countries where women must have permission to get a job (World Bank, 2018).

<sup>6</sup>Engineering, for example, was not offered to Saudi women as an undergraduate degree program until 2005.

<sup>7</sup>Evidence in other contexts suggests that these effects are not always straightforward. In Japan, for example, Brinton and

### 2.2.2 Variable Costs

In addition to the costs of employing women at all, employers also face differences in hiring costs when employing Saudi women. Wages tend to be significantly lower for Saudi women as compared with Saudi men: the average monthly full-time wage for women at baseline is SAR 3190 compared with SAR 6279 for men. Even when controls are added for education, location and occupation women earn 40 percent less than men.<sup>8</sup> Workers may also have different average productivity levels due to mobility, professional skills, retention, occupational restrictions, and potential for career development.

There may also be additional per-worker costs to employing women in the Saudi context. Saudi labor law requires firms to provide ten weeks of paid maternity leave to new mothers. Firms may also opt to provide other benefits to attract female workers, such as transportation allowances<sup>9</sup> or childcare benefits. The lack of historical female employment in the private sector may mean that female hires tend to require different types of training.

## 2.3 Integration Incentives from Nitaqat Employment Quotas

In this paper we examine the evidence that firms face fixed integration costs. We do this in part by examining how firms responded to Nitaqat Saudization quotas, a policy shock that provided an incentive to hire large numbers of Saudi employees. This provided an opportunity for firms to spread potential integration costs over a large number of hires, potentially provoking a large integration response.

### 2.3.1 Nitaqat Program Structure

The Nitaqat program is an ongoing nationalization quota policy first instituted in 2011.<sup>10</sup> The policy was designed to address growing national unemployment, which had reached 40 percent for Saudis in the 20-25 age group in 2011, in the context of the low participation of nationals in the private sector. At the time, foreign guest workers made up ninety percent 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 Saudi Ministry of Labor and Social Development (then the Ministry of Labor) 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 (500-2999 employees) and giant (3000+ employees). Industry classifications were made using economic activities registered with the Ministry of Commerce, and size group using the total number

---

Mun (2016) finds that interactions between norms around ideal employee behavior and gendered expectations about family care work produce extremely high barriers to career advancement for professional women.

<sup>8</sup>A Mincer regression of the log of private sector wages at baseline on employee characteristics indicates that Saudi women earn 40 percent less than men within occupations after controlling for educational attainment, years of potential experience, and location (all with dummy variables).

<sup>9</sup>A 2017 royal decree also pushes firms to provide transportation cost offsets for their female employees.

<sup>10</sup>See (Peck, 2017) for a more detailed description of the Nitaqat program and its effects.

of Saudi nationals registered as employees with the General Organization for Social Insurance (GOSI) and foreign nationals with visas sponsored by the company, as registered with the National Information Center. Firms with fewer than ten employees were not subject to Nitaqat regulations over this period, though they were added in to the regulation post-2012.

Within each cell of this industry by size classification, the Ministry then defined four different color bands corresponding with a company's Saudi employee percentage. A large-sized manufacturing firm, for example, faced the following cutoffs for the four color bands:

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 Ministry staff. The Ministry used its visa issuance and foreign recruitment services to enforce the program. These services were tied directly into the monitoring system and implemented automatically on firms that failed to meet their nationalization quotas.<sup>11</sup> 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 Ministry 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 had the effect of dramatically increasing the number of Saudis in the private sector (Peck, 2017), with firms complying with the program by increasing their Saudi employment.

### 2.3.2 Nitaqat Employment Effects

Nitaqat quotas were effective at increasing Saudi employment recorded in the private sector, though it did so at significant cost in terms of firm exit and expatriate employment (Peck, 2017). The program corresponded with a significant increase in employment of Saudis at intermediate education levels for both men and women (Figure 2)). Firm-level increases in Saudi employment were closely related to distance from the quota (Peck, 2017), and these employment increases were concentrated at firms below the quota, with a large discrete jump in male employment at these firms after the start of the quota enforcement (Figure 3a). Though the initial effect was smaller for female employment, the growth in female employment at firms below the quota continued to increase over the period compared to those firms above (Figure 3b). While the program appears to have increased the level of male employment at affected firms, it seems to

---

<sup>11</sup>Color band status in the system was based on a thirteen-week moving average of the number of Saudi workers registered with GOSI divided by the total number of workers.

have instead increased the growth rate of female employment at these same firms, leading to large and growing female employment effects.

[Figure 2 about here.]

[Figure 3 about here.]

Much of this increase in female employment was concentrated in clerical occupations (Table 1). The largest increase in female employment was for occupations in the “Clerical Support Workers” category, 37 percent of overall growth in female employment. Another significant contributor to women’s employment was the retail sector: 15 percent of the increase in female employment occurred in the “Services and Sales Worker” category, with 58 percent of these coded as “Sales Workers” (Table 1).

[Table 1 about here.]

## 2.4 Other Labor Policies

In addition to Nitaqat, the Saudi government has also pursued a slate of practical measures designed to increase women’s employment, 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.<sup>12</sup> Though not female-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 percent of Hafiz beneficiaries have been women (Evidence for Policy Design, 2017). The Ministry removed regulations requiring women to obtain permission from a male guardian to apply for private sector jobs.<sup>13</sup> Many firms still require a guardian’s approval, though the Ministry recently forbade this practice among government employers.<sup>14</sup>

We now turn to developing a basic model to capture the intuition regarding both the employment effects of the fixed cost to integrating as well as how quotas may affect integration in the presence of such fixed costs.

## 3 Model

In this section, we develop a simple search model to study the implications of fixed integration costs faced by firms for female employment. We adapt the Black (1995) search model with employer discrimination. Unemployed workers decide whether to search for a job in the labor market, and firms decide how many

---

<sup>12</sup>The decree was expanded to also cover stores selling women’s clothing and accessories beginning in January 2014. There are 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).

<sup>13</sup>Jafar AlShayeb, *Arab News* June 15, 2010 “Women’s rights gain focus in the Kingdom”

<sup>14</sup>Lulwa Shalhoub, *Arab News* May 5, 2017 “Saudi women no longer need guardians’ consent to receive services” <http://www.arabnews.com/node/1094681/saudi-arabia>

vacancies to post. Vacancies posted by segregated firms can only be filled by male searchers. At a (fixed) cost, a firm can integrate so that their vacancies can be filled by both male and female searchers. Integrated vacancies will produce more successful matches, but are only worth the up-front cost if enough women search. On the supply side, women only find it worthwhile to search if enough firms have integrated.

Suppose there are  $L$  workers that are one of two types:  $t \in \{M, F\}$ . Let  $\gamma > 0$  denote the fraction of workers that are type  $F$ . For simplicity, we assume both type  $M$  and type  $F$  have the same marginal product, denoted  $Y$ . We index workers by  $i$ .

Workers may choose not to enter the labor market and receive utility  $h_i$  from home production, or they may choose to sequentially search for a job. There are two types of firms, segregated and integrated. Vacancies posted by segregated firms can only be filled by type  $M$  workers. Vacancies posted by integrated firms can be filled by both type  $M$  and type  $F$  workers. Let  $\theta$  denote the share of vacancies posted by segregated firms.

To begin, we fix labor force participation and assume that all  $L$  workers search. The segregated share of vacancies  $\theta$  will be determined endogenously. Later, we endogenize labor supply so that unemployed workers decide whether to search in the labor market.

### 3.1 The Worker's Problem

Workers search over vacancies for a job, sequentially and randomly matching with new vacancies. For each vacancy the worker derives match-specific job satisfaction,  $\alpha$ . This  $\alpha$  has cumulative distribution function  $F_\alpha(\alpha)$  and probability density function  $f_\alpha(\alpha)$ . We assume that the inverse hazard function,

$$\mu(\alpha) \equiv \frac{1 - F_\alpha(\alpha)}{f_\alpha(\alpha)}$$

satisfies

$$\mu'(\alpha) < 0.<sup>15</sup>$$

Workers search sequentially until they receive an offer that provides a utility level at least as great as their reservation utility level,  $r^t$ . Type  $M$  workers receive wage offers  $w_I^m$  and  $w_S^m$  from integrated and segregated firms, respectively. Type  $F$  workers receive wage offer  $w_I^f$  from integrated firms. Thus, workers accept an offer when  $\alpha \geq r^t - w_j^t$ , where  $j = I, S$ .

For type  $M$  workers, the value of search is

$$U^m = \theta \mathbb{E} \max\{w_I^m + \alpha, U^m\} + (1 - \theta) \mathbb{E} \max\{w_S^m + \alpha, U^m\} - \kappa$$

where  $\kappa$  is the cost of search. With probability  $\theta$ , the worker is matched with an integrated firm and is offered wage  $w_I^m$ . With probability  $1 - \theta$ , the worker is matched with a segregated firm and offered wage

---

<sup>15</sup>Many common distributions satisfy this restriction, including the uniform, gamma, and normal distributions.

$w_S^m$ . This can be re-written as

$$U^m = \frac{\theta \int_{\hat{\alpha}_S^m}^{\infty} (w_S^m + \alpha) f_{\alpha}(\alpha) d\alpha + (1 - \theta) \int_{\hat{\alpha}_I^m}^{\infty} (w_I^m + \alpha) f_{\alpha}(\alpha) d\alpha - \kappa}{1 - \theta F_{\alpha}(\hat{\alpha}_S^m) - (1 - \theta) F_{\alpha}(\hat{\alpha}_I^m)},$$

where  $\hat{\alpha}_j^m \equiv r^m - w_j^m$  for  $j = I, S$ . The reservation utility that maximizes expected utility is set so that

$$r^m = U^m.$$

In words, the worker sets his reservation utility such that he is indifferent between accepting a job at his reservation utility level and continuing search. This implies that

$$\kappa = \theta \int_{\hat{\alpha}_S^m}^{\infty} (w_S^m + \alpha - r^m) f_{\alpha}(\alpha) d\alpha + (1 - \theta) \int_{\hat{\alpha}_I^m}^{\infty} (w_I^m + \alpha - r^m) f_{\alpha}(\alpha) d\alpha.$$

The left-hand side is cost of an additional search, while the right-hand side reflects the expected gains of additional search.

The expected number of searches that a type  $M$  worker will undertake,  $\epsilon^m$ , is

$$\epsilon^m = \frac{1}{\theta[1 - F_{\alpha}(r^m - w_S^m)] + (1 - \theta)[1 - F_{\alpha}(r^M - w_I^m)]}.$$

For type  $F$  workers, the value of search is

$$U^f = \theta \mathbb{E} \max\{w_I^f + \alpha, U^f\} + (1 - \theta)U^f - \kappa.$$

With probability  $\theta$ , the worker matches with an integrated firm and is offered wage  $w_I^f$ . With probability  $1 - \theta$ , the worker matches with a segregated firm, where she cannot fill the vacancy and hence must continue searching. As above, we get the following condition:

$$\kappa = (1 - \theta) \int_{r^f - w_I^f}^{\infty} (w_I^f + \alpha - r^f) f_{\alpha}(\alpha) d\alpha.$$

Note that

$$0 < \frac{dr^f}{d\theta} < 1$$

or, in words, an increase in the share of vacancies posted by integrated firms increases the reservation wage for type  $F$  workers. This is true because an increase in  $\theta$  reduces search costs for type  $F$  workers.

The expected number of searches for a type  $F$  worker is

$$\epsilon^f = \frac{1}{(1 - \theta)[1 - F_{\alpha}(r^f - w_I^f)]}.$$

Next, we move on to the firm's problem.

### 3.2 The Firm's Problem

Each firm  $j$  is exogenously assigned to post  $n_j > 0$  vacancies to hire these workers.<sup>16</sup> Let  $N = \sum_j n_j$  equal the total number of vacancies posted. Each firm will decide whether to integrate prior to posting the vacancies and pay fixed integration cost  $\phi$ , or remain segregated at no cost. Vacancies posted by integrated firms can be filled by both type  $M$  and type  $F$  workers; a segregated firm can only hire type  $M$  workers. Firms also set wages.

Integrated firms will set wages that maximize expected profits per applicant,

$$\pi_I^t = [1 - F_\alpha(r^t - w_I^t)](Y - w_I^t)$$

the probability an applicant accepts the offer times the profit derived from an accepted offer. The first order condition implies that profit maximization requires

$$Y - w_I^t - \mu(r^t - w_I^t) = 0.$$

Segregated firms will only hire type  $M$  workers, so expected profits per type  $M$  applicant are

$$\pi_S^m = [1 - F_\alpha(r^m - w_S^m)](Y - w_S^m)$$

which yields necessary condition

$$Y - w_S^m - \mu(r^m - w_S^m) = 0.$$

Together, this implies that  $w_I^m = w_S^m \equiv w^m$ . We can also suppress the subscript for  $w_I^f$ , so that  $w_I^f \equiv w^f$ .

#### 3.2.1 The Integration Decision

The expected profit for a segregated vacancy is

$$V_S = \frac{L(1 - \gamma)\epsilon^m \pi^m}{N}.$$

This is equal to the expected number of type  $M$  applications times the profit per type  $M$  applicant, divided by total vacancies.

The expected profit for an integrated vacancy is

$$V_I = \frac{L [\gamma \epsilon^f \pi^f + (1 - \gamma) \epsilon^m \pi^m]}{N}.$$

The difference in expected profit between a segregated and integrated vacancy, or the value of the

---

<sup>16</sup>Note that in our setting a 'vacancy' can be potentially filled by multiple workers. Hence, it may be more natural to think of a 'vacancy' as a 'job posting' here.

ability to hire type  $F$  workers, is

$$V_I - V_S = \frac{L\gamma\epsilon^f\pi^f}{N} > 0.$$

Hence, a firm will choose to integrate if the following condition is met:

$$n_i > \frac{\phi}{V_I - V_S}$$

The distribution of  $n_i$  across firms determines equilibrium  $\theta$ . Suppose that  $n_i$  is distributed across firms where the cumulative distribution function is denoted  $F_n(n)$ , where  $n_i \in [1, \bar{n}]$ . Then  $\theta$ , the share of vacancies that are segregated, is given by

$$\theta = \frac{\int_1^{n^*} n dF_n(n)}{N}$$

where

$$n^* = \frac{\phi}{V_I - V_S} = \frac{\phi N}{L\gamma\epsilon^f\pi^f}$$

This yields our first two predictions.

**Prediction 3.1** *The distribution of female workers across firms exhibits bunching at zero and a ‘missing mass’ of firms that employ a small number of female employees.*

This result is a direct consequence of our assumption that integration costs are fixed.

Relatedly, though we do not model firm hiring *dynamics*, the next prediction is another natural implication of fixed costs:

**Prediction 3.2** *A firm’s hiring of women exhibits state dependence where, for a given firm, the female share of new hires is higher if the firm already has female employees at that time.*

### 3.3 Modeling Nitaqat

We model Nitaqat as an increase in  $n_i$  for a subset of firms. The idea is that employment quotas force a firm to post more vacancies for Saudi workers.<sup>17</sup> Hence, under Nitaqat, share  $\eta$  of firms have binding quotas, and the new distribution for  $n_i$  is given by  $\tilde{F}_n(n)$ , where

$$\tilde{F}_n(n) = \eta F_n(n - q) + (1 - \eta) F_n(n)$$

Among firms where the quota is binding, the change from  $F_n(\cdot)$  to  $\tilde{F}_n(\cdot)$  will induce some firms to integrate. This implies that the female share of hires at below quota firms will increase relative to above quota firms which, with fixed labor supply, are not induced to integrate.

<sup>17</sup>As described above this seems to be consistent with how firms respond to quotas.

**Prediction 3.3** Among firms that were **segregated** at baseline, Nitaqat increases the type  $F$  share of employment at below quota firms relative to above quota firms.

By contrast, among firms that were integrated at baseline, they will remain integrated under Nitaqat. Hence, among these firms, we do not expect to see a relative increase in the female share of hires at below quota firms.

**Prediction 3.4** Among firms that were **integrated** at baseline, Nitaqat increases the type  $F$  share of employment **equally** at below quota firms relative and above quota firms.

Finally, though not a prediction we will test directly, note the following result: an increase in  $\gamma$  (the share of workers that are type  $F$ ) decreases  $\theta$ , the share of vacancies that are segregated.

### 3.4 Endogenous Labor Supply

Next, we allow workers to decide whether to enter the labor market at all. In particular, we will study how changes in  $F_n(n)$  that in turn affect  $\theta$  alter labor supply decisions.

Let  $\omega$  denote the type  $F$  share of *potential* workers, while  $\gamma$  is the type  $F$  share of workers in the labor force. The two will not necessarily be equal. As above, we will assume that all type  $M$  potential workers indeed search. For type  $F$  workers, let  $F_h(h)$  denote the CDF of home production utilities,  $h_i$ . Type  $F$  workers will search if

$$h_i \leq U^f = \frac{(1 - \theta) \int_{\hat{\alpha}^f}^{\infty} (w^f + \alpha) f_{\alpha}(\alpha) d\alpha - \kappa}{(1 - \theta) F_{\alpha}(\hat{\alpha}^f)}$$

Hence, the type  $F$  share of the labor force,  $\gamma$ , is given by

$$\gamma = \frac{F_h(U^f) \omega}{F_h(U^f) \omega + (1 - \omega)}$$

What happens if  $F_n(n)$  adjusts in a way that decreases  $\theta$ ? Then the value of search for type  $F$  workers,  $U^f$ , increases, and hence  $\gamma$  increases. That is, an increase in the demand for type  $F$  workers will in turn increase supply. Interestingly, as we've shown above, an increase in  $\gamma$  increases  $\theta$ . Hence, there is a feedback loop between  $\gamma$  and  $\theta$ .

**Prediction 3.5** Among firms that were **segregated** at baseline, Nitaqat increases the type  $F$  share of employment at **both** below quota firms and above quota firms.

In the following sections, we test each of these predictions empirically to determine whether our results are in line with our basic framework. We first discuss our data sources below.

## 4 Data

Our primary dataset is the administrative social security data obtained from the General Organization for Social Insurance (GOSI) of Saudi Arabia. These data (hereafter referred to as the GOSI data) contains

information on all Saudis employed in the private sector between January, 2009 and June, 2015<sup>18</sup>. The dataset is used to track Saudi employees for social security eligibility and withdrawal purposes. The dataset 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; as well as firm information such as their administrative identifiers and industries. In total, the GOSI data contains information on approximately 2.8 million unique individuals and 430 thousand firms.

The dataset 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 contain date ranges that specify when the individual was working in a particular firm in the noted occupation and paid the corresponding wage.

As such, the dataset provides information about the career profile of individuals, but is difficult to use for analyzing trends over time. Since the bulk of our analysis requires tracking hiring trends over time, we use the durations associated with each observation to transform the data into monthly observations. This provides an unbalanced monthly panel for each Saudi employee, which further allows us to determine whether an employee is a new entrant to the private sector labor force, switching firms, or is exiting the private sector in any particular month. To clean up potentially erroneous observations, we drop individuals with ages below 10 or above 100. We also drop entries for part-time work, which only affects 47 thousand 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 2-digit ISCO-08 group, reducing the number of occupations from 2,151 to 40. This significant drop in occupations is primarily due to variations in translation from the original Arabic to English, inconsistent naming, as well as changes to the GOSI classification scheme over time. Appendix Table A1 lists the top 10 most common ISCO-08 coded occupations in June, 2011, the baseline date of our analysis.

#### 4.1 Quota status

The GOSI data above does not contain information on the quotas that each firm was subject to as part of the Nitaqat program. In order to determine whether and the extent to which each firm faced quotas, we use administrative data from the Nitaqat program (hereafter referred to as the Nitaqat data). As described by Peck (2017), this dataset is used to track compliance with the national quota on hiring Saudis in the private sector. These data provides information on whether a given firm was subject to quotas during a given week, and, if so, whether they met the quotas for that particular week. These data provides weekly quota compliance information from June, 2011 (the start of the Nitaqat program) until December, 2013<sup>19</sup>.

---

<sup>18</sup>Our data only goes back to 2009 due to a change in how the data was stored and collected by GOSI. Unfortunately we are not able to obtain information prior to this year.

<sup>19</sup>Unfortunately we were only able to access Nitaqat data until this time period, and were not able to get more recent data to overlap with the data available from GOSI.

Firms are defined differently between the Nitaqat and GOSI datasets. 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 above in the Background section, 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)<sup>20</sup>. 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 harmonize the definition of the firm between the two datasets, 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 thousand 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 defines firms with a national or multi-city presence as separate commercial organizations depending on the geographic Ministry of Labor and Social Development (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 to be a legal commercial organization within a particular province.

We primarily use the Nitaqat data to obtain a list of firms and their quota status for the second week of June, 2011, when the program began assessing quotas and reporting color band status to firms. This gives us a sample of approximately 1.07 million firms at our baseline, over 990 thousand of which were exempt from the program for having fewer than 10 employees (i.e. status of white). Approximately 113 thousand of these firms appear in the GOSI data. The big drop in the number of baseline firms between the two datasets is primarily due to the fact that many white firms do not need to hire any Saudi employees, and therefore 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. Most of our analysis is restricted to firms existing in June, 2011 and their employees in order to keep the sample of firms consistent.

We merge these two datasets based on firm identifiers. Combining these two datasets therefore gives us an employee-firm matched dataset for Saudis with information on quota compliance status at our baseline in June, 2011. This allows us to track hiring (and therefore integration) and quota compliance trends for these Nitaqat-eligible firms.

To construct a measure of quota compliance, we create a distance metric that measures the number of

---

<sup>20</sup>An entity consisting of multiple branches (e.g. a national franchise) are counted as a single entity for each branch of the Ministry of Labour and Social Development labor office they are linked to.

Saudis that would have to be hired to meet the quota at baseline, or in June, 2011 when the quotas were first imposed<sup>21</sup>:

$$Distance_{ijs}^S = \max\{Saudi_{ijs}^* - Saudi_{ijs}, 0\}$$

where  $Saudi_{ijs}^*$  is the number of Saudis necessary to meet the relevant quota without changing the baseline number of expatriates for firm  $i$  in industry  $j$  with size group  $s$ . This metric allows us to create a standardized variable for the quota pressure that firms face.

In addition, we also further categorize firms by their integration status, or, in other words, whether or not they have hired their first Saudi female employee. Firms are divided into three categories depending on their hiring status:

- **Legacy integrator firms:** Firms that employed Saudi women in our first month of data, January 2009
- **Newly integrated firms:** Firms who have hired their first Saudi female employee after January 2009
- **Never integrated firms:** Firms which are always segregated (i.e. male Saudi employees only) for the duration of the GOSI data (Jan. 2009 to June, 2015)

This division of firms is used in our empirical analysis to determine which firms integrate and when they do so. In addition, it allows us to construct potential counterfactuals when considering the determinants of integration. Having described the data, we now turn to our empirical analysis to test our model's predictions.

## 5 Empirical Results

In this section we test the following predictions from the model: (1) the distribution of female workers across firms exhibits bunching at zero and a 'missing mass' of firms that employ a small number of female employees; (2) the distribution of female *hiring* follows a similar pattern, particular at non-legacy firms; (3) Nitaqat increases relative labor demand for women at firms segregated at baseline, but (4) not for firms integrated at baseline. In addition, we present suggestive evidence that binding Nitaqat quotas at some firms induce integration at firms where quotas are not binding.

### 5.1 Gender Segregation Across Firms at Baseline

We first test whether, at baseline, Saudi female employment is distributed across firms in a manner consistent with prediction 3.1 of our model. In the model, the presence of fixed integration costs generates a

---

<sup>21</sup>Firms could alternatively meet the quota by downsizing the number of expatriates that they hire. Since we are primarily concerned with firms that integrate by employing Saudi females, we do not pursue this hypothesis extensively. On this note, Peck (2017) shows that Red and Yellow firms' employment of expatriates is less responsive to quota pressure than hiring of Saudi employees.

‘missing middle’ in this distribution, where few firms employ a small number of female employees relative to a counterfactual with no fixed integration costs. In particular, there exist a set of firms that would hire a small number of female employees in this counterfactual, but do not integrate when fixed integrated costs exceed the variable labor cost savings of integration.

We test this prediction by comparing the actual distribution of female employment across firms to a benchmark with no fixed integration costs. To construct the benchmark, we simulate the distribution of female employment across firms, taking each firm’s size as given but randomly re-assigning workers across firms. This exercise is similar to the approach taken in [Hellerstein and Neumark \(2008\)](#), where the authors measure workplace segregation by race and ethnicity in the United States. The authors simulate the degree of workplace segregation by race that would occur by chance alone and compare that to observed segregation. Across simulation specifications, we vary the units within which workers are randomly assigned. We randomly reassign workers within groups defined by worker and job characteristics, including firm location and worker occupation. Each firm’s simulated number of female employees depends solely on its size, these job and worker characteristics, and random chance.

The first benchmark bundles all workers into a single group, reassigning all workers at random to different firm positions. The second benchmark differentiates workers and jobs by location, reassigning workers to positions at firms within local labor markets. The third benchmark differentiates workers and jobs by occupation. We classify occupations by ISCO-08 2-digit occupation groups, of which there are 40 in total. The fourth benchmark differentiates workers and jobs by both location and occupation. This means, for example, that this simulation would only reassign an accountant in Dammam to other accountant positions in the same city held in the data by other workers. Differentiating workers by occupation potentially obscures the role of fixed integration costs if integrating firms reorganize to focus on occupations that are better suited to women, but accounts for the fact that some occupations are extremely unlikely to be occupied by women in Saudi Arabia and so may provide a more reasonable benchmark.

This benchmark should not be interpreted as a counterfactual *per se*. In particular, we set total female employment in our benchmark to equal actual female employment and, in the absence of fixed integration costs, we would likely expect aggregate female employment to increase. Instead, the benchmark represents how we would expect a fixed set of male and female workers to be distributed across firms in the absence of fixed integration costs. In all simulations, we restrict the analysis to firms with at least five Saudi employees at baseline.

Figure 4 compares the distribution of female employment across firms to our initial benchmark. Panel A reports the share of firms with exactly zero female employees. Panel B reports the share of firms with at least one female employee across a range of bins. We report these shares in separate panels to allow for different scales in the figures.

[Figure 4 about here.]

We find strong evidence of a ‘missing middle’ of firms with a small number of female employees. While in our richest simulation about 38% of firms have zero female employees, about 72% of firms have zero female employees in practice. This mass is shifted from the set of firms with one to four female

employees. While 47% of firms have one to four female employees in our richest simulation, only 14% of firms have this number of female employees in practice. This pattern is consistent with our model. A share of segregated firms would hire a small number of female employees if not for the fixed costs associated with integration.

In Table 2 we describe the simulated distribution of female employment across firms under various conditions. Column (1) contains the actual distribution of female employment across firms. Column (2) shows the baseline simulation. In column (3) we fix the probability that an employee in a given occupation is female to the share of workers in that industry that are female. In column (4) we set probabilities based on occupation and location. The simulated distributions are similar across all specifications. Across simulations, 29-38% of firms have zero female employees, while 47-55% have between 1-4 female employees. There are far more firms with zero female employees and far fewer firms with 1-4 female employees than we would expect if labor demand and supply by gender depended solely on occupation and firm location.

[Table 2 about here.]

By comparing the realized distribution of female employment to the distribution we would ‘expect’ in the absence of firm-level frictions, we find a far higher share of firms are segregated than expected. Our fixed costs model implies that the gap between realized and expected segregation will be largest for firms that we expect to hire few women in the first place, so that fixed integrated costs are more likely to exceed the integration benefits. We test this prediction in Panel C of Figure 4. We first divide firms into bins based on their expected number of female employees given their location and occupational composition. For each bin we then calculate the percentage of firms that are integrated, both in the actual and simulated data.

Consistent with our fixed costs model, we find substantially larger discrepancies between actual and simulated rates of integration for firms with a low expected number of female employees. For firms with between 1-5 expected female employees, 87% are integrated in the simulation while 34% are integrated in practice. By contrast, for firms with 25 or more expected female employees, 100% are integrated in the simulation while 88% are integrated in practice.

This comparison provides some sense for the distribution of fixed costs across firms. Within each bin, the gap in the percentage of firms that are integrated between the realized and simulated distributions suggests the share of firms facing fixed costs that exceed the variable cost savings associated with employing the corresponding expected number of female employees.

## 5.2 Gender Segregation in Hiring

We next test whether female *hiring* is distributed across firms in a manner consistent with prediction 3.2 of our model. As in section 5.1, we test whether there is a ‘missing middle’ in the distribution of female employment across firms at endline, restricting to employees hired after June 2011. We restrict to firms that exist at both baseline and endline, and conduct separate analyses for legacy and non-legacy firms. We construct a benchmark via simulation by fixing the number of employees hired by each firm and randomly re-assigning the gender of each worker hired.

Figure 5 compares the distribution of female hires across non-legacy firms to our richest simulation benchmark.

[Figure 5 about here.]

Again, we find strong evidence of a “missing middle” of firms with a small number of female recent hires. While in our simulation about 13% of firms have zero recent female hires, about 36% of firms have zero female recent hires in practice. This mass is shifted from the set of firms with one to four female recent hires. While 52% of firms have one to four female recent hires in our simulation, only 32% of firms have this number of female recent hires in practice.

Figure 6 compares the distribution of female hires across legacy firms to our simulation benchmark.

[Figure 6 about here.]

Among legacy firms, evidence of a ‘missing middle’ is substantively weaker. In our simulation, 26% of firms have 1-4 female recent hires in our simulation, while 21% of firms have this number of female recent hires in practice.

We also test whether firm hiring of women exhibits state dependence as predicted by the model. In the model, firm preferences over potential hires depend critically on whether the firm has paid fixed integration costs. Following integration, the marginal cost of employing an additional woman decreases substantially. The logic of the model predicts that, once a firm has female employees, it will hire woman at a higher rate.

To summarize firm hiring patterns, we plot the female share of hires at integrating firms in the months following a firm’s first observed female hire.<sup>22</sup> We limit to firms with at least 5 Saudi employees in the month prior to integration. Prior to integrating, we observe integrating firms in the GOSI data for an average of 28 months. We split firms by whether they integrate prior to or following Nitaqat. We do this to confirm that hiring patterns we observe are not specific to Nitaqat and the potential influx in female labor supply that coincides with that policy.

The event studies are shown in Figure 7. We plot the female share of new hires in six-month increments before and after a firm’s first female hire, averaged across all firms meeting the sample restrictions described above. Hiring patterns are similar at firms that integrate prior to and following the implementation of Nitaqat, though the female share of hires are generally larger at firms that integrate following Nitaqat. By construction, among hires made 7 to 12 months and 1 to 6 months 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. Among hires made in the 6 months following integration, including the first female hire, about 40% are female at early integrators and about 50% are female at late integrators. This drops to about 30% in the following six-month period, and converges to about 25% for both early and late integrators thereafter.

[Figure 7 about here.]

---

<sup>22</sup>In this exercise, we exclude firms that have female employees when they are first observed in the GOSI data.

Overall, the pattern we observe is consistent with the logic of the model. Following integration, firm hiring patterns change dramatically. Integrating firms transition immediately from employing no women at all to hiring women on a regular basis.

### 5.3 Nitaqat Increases Relative Demand for Women at Segregated Firms

Turning to predictions 3.3 and 3.4, we next test whether Nitaqat increases relative labor demand for Saudi women among segregated firms (but not among integrated firms). Our model predicts that Nitaqat will induce some previously segregated firms to integrate by increasing their target number of Saudi employees. All else equal, firms that experience larger increases in their target number of Saudi employees are more likely to integrate. To support these predictions, we show that: (1) the female share of the Saudi private sector workforce experiences a trend break that coincides precisely with the implementation of Nitaqat; (2) this increase is concentrated at newly-integrating firms; (3) this increase is concentrated at firms further under quota pressure at baseline.

Figure 8 Panel A plots the female share of the Saudi private sector workforce over time. The vertical line marks June 2011, when Nitaqat is first implemented. From January 2009 to June 2011, the female share of the Saudi private sector workforce hovers around 10%. At the onset of Nitaqat, there is a clear trend break and the female share begins to increase dramatically. By December 2013, the female share increases to about 27%, nearly a three-fold increase, and stagnates thereafter. Recall that the total Saudi private sector workforce is also increasing over time. The total number of Saudi women in the private sector increases from approximately 71,000 in June 2011 to 463,000 in June 2015.

Figure 8 Panel B plots the share of firms with any Saudi female employees, among firms with any Saudi employees. The pattern is similar to the pattern observed in Panel A. From January 2009 to June 2011, the share of firms with any Saudi female employees is relatively stagnant at about 12%. Between the onset of Nitaqat and December 2013, this share increases to about 50%, and levels from there.

[Figure 8 about here.]

We take this clear trend break as strong *prima facie* evidence that Nitaqat caused a substantial increase in the female share of the private sector workforce, consistent with the model. A sharper prediction of the model is that this increase is at least in part driven by an increase in relative demand for female workers at firms that had not previously integrated. We next show that this increase in female share is concentrated at newly-integrated firms. We also provide evidence that this increase is at least in part driven by demand rather than labor supply by exploiting variation across firms in quota pressure.

We first document the female share of Saudi workers over time at four distinct sets of firms. We first restrict to firms that exist at baseline and have at least one Saudi employee at that time. We also restrict to firms in the Red, Yellow, Green, and Platinum Nitaqat color bands. We then split the remaining firms on two dimensions: integration status and Nitaqat quota status. Firms are divided by integration status based on whether they are integrated as of January 2009. We label firms that have had a female employee over this employee as *legacy* (integrated at baseline) firms. Our model predicts that the increase in relative demand for female workers is concentrated at *non-legacy* firms. Legacy firms have already paid fixed

integration costs, and so should already be willing to hire both men and women at the margin. Firms are also divided based on their color band at baseline. We group firms into those below their Nitaqat quota ('below quota'), Yellow and Red firms, and firms that satisfy their quota ('above quota') at the onset of Nitaqat, Green and Platinum firms. The model predicts a larger demand response at below quota firms, where Nitaqat has a larger and more binding effect on their target number of Saudi employees.

[Figure 9 about here.]

[Figure 10 about here.]

We find evidence that, consistent with Peck (2017), Nitaqat induces below quota firms to increase their hiring of Saudis. In Figure 9, we plot total Saudi hires by quarter, separately by legacy and quota status. Panel A plots quarterly hires for below quota and above quota non-legacy firms. The vertical line marks Q1 of 2011; Nitaqat is implemented in June 2011, the end of Q2 2011. Prior to Nitaqat, total hires at below and above quota firms move roughly in step. Following the implementation of Nitaqat, hiring at above quota firms remains on a similar trend. We interpret this as *a priori* evidence that Nitaqat did not have a significant effect on demand for Saudi labor at these firms. By contrast, hiring at below quota firms increases sharply. From Q2 through Q4 of 2011, Saudi hiring at below quota firms doubles relative to hiring at above quota firms. Saudi hiring at below quota firms stagnates thereafter but remains elevated relative to above quota firms. Panel B plots hiring at legacy firms, which follow a similar pattern. Below and above quota firms follow similar trends prior to Nitaqat, and above quota firms remain on a similar trend following the implementation of Nitaqat. Hiring increases dramatically at below quota firms.

We next document the female share of Saudi hires at these four sets of firms in Figure 10. Panel A plots the female share of Saudi hires at non-legacy firms, while Panel B does the same for legacy firms. There are two clear patterns to note. First, there is a dramatic increase in the female share of employees at non-legacy firms, and this increase is larger at below quota firms. For below quota non-legacy firms, the female share increases from about 5% in Q1 2011 to about 22% at endline. By endline, the female share of hires at these firms approaches that of legacy firms. At above quota non-legacy firms, the female share increases from 5% to about 16%. This is despite the fact that, as indicated by Figure 9, Nitaqat does not appear to lead above quota firms to hire more Saudis in general. In the context of our model, this type of response is consistent with a demand externality. The relative availability of female labor may also increase because male labor supply is depleted, and not through a female supply response *per se*. The response from above quota firms is also consistent with more direct firm-to-firm spillovers outside of the scope of our model, where firms integrate in response to their peers integrating, perhaps due to changes in social norms. Second, changes in the female share of hires are more modest at legacy firms, and without a clear difference between above and below quota firms.

One concern with interpreting the patterns documented in Figure 10 is they may be driven by compositional changes in the set of firms that are hiring. We complement this aggregate analysis with a simple firm-level difference-in-difference model. We estimate firm-level models of the form:

$$Y_{it} = \alpha_i + \tau_t + \sum_j \beta_t \text{Below}_i \times \mathbf{1}_{t=j} + \epsilon_{it} \quad (1)$$

where  $i$  indexes firms,  $t$  indexes quarters,  $Y_{it}$  is the number of male or female hires made by firm  $i$  in quarter  $t$ ,  $\alpha_i$  are firm fixed effects and  $\tau_t$  are quarter fixed effects.  $\text{Below}_i$  is an indicator for whether firm  $i$  is below the Saudi quota at the onset of Nitaqat, so that the  $\beta_t$  coefficients reflect differential hiring at below quota firms over time. As above, we restrict the analysis to Red, Yellow, Green, and Platinum firms. Descriptive statistics for the firms included in this analysis are presented in Appendix Table A2.

We plot the  $\beta_t$  coefficients from equation (1) in Figure 11. In Panel (a), we plot the coefficients from equation (1) where the outcome is number of male Saudi hires; in Panel (b), the outcome is female Saudi hires. For both outcomes, there is little evidence of pre-trends; male and female hiring do not differentially change at below quota firms prior to Q2 of 2011. For both male and female hiring, there is a clear trend break in Q2 of 2011. Nitaqat is implemented in June 2011, the end of that quarter. For male hiring, below quota firms hire about 1.25 more men per quarter by Q3 of 2011. For reference, the average number of male hires per quarter at above quota firms is 1.8. Male hiring remains elevated for several quarters, but steadily declines at below quota firms (relative to above quota firms) until 2013 or so, when the difference in male hiring returns to pre-Nitaqat hiring.

For female hiring, the pattern is quite different. Female hiring at below quota firms increases (relative to above quota firms) beginning in Q2 of 2011, and by Q4 of 2011 those firms are hiring nearly 0.2 more women per quarter. For reference, the average number of female hires per quarter at above quota firms is 0.26. By contrast to the pattern for male hiring, female hiring *remains* elevated at below quota firms throughout the period we have data. In Q2 of 2015, the last quarter where we have data, the  $\beta$  coefficient remains above 0.1.

This pattern of coefficients is consistent with fixed costs of integration. Below quota firms initially hire more men and women to comply with the quota. They continue to hire more women because they have already paid the fixed costs of integration.

[Figure 11 about here.]

In Table 3 we compare characteristics of legacy firms, firms that integrate after January 2009, and firms that never integrated. Here we restrict to firms that exist at both baseline and endline.

[Table 3 about here.]

Overall, 68.5% of the total increase in female employment occurs in newly-integrated firms. If we restrict to firms that employed at least one Saudi at baseline, newly-integrated firms account for 50% of the increase.

The differential patterns across firms displayed in Figure 10 suggests that quota pressure increased relative labor demand for female Saudi employees. However, it is possible that below quota firms differ from above quota firms in ways that would lead to differential changes in female share even in the absence of differential pressure under Nitaqat. For example, below quota firms may come from industries or hire in occupations that are more likely to hire women over this period. Moreover, below quota firms may show larger changes in female share simply because they have larger gross flows of Saudi workers.

To address these concerns, we test whether below quota firms integrate at a faster rate following Nitaqat than otherwise similar above quota firms. To do this, we estimate discrete-time hazard models for

the probability of a firm hiring its first female employee, following [Jenkins \(1995\)](#). In these models, a firm’s exposure is their number of hires between the baseline and endline. This accounts for the fact that firms that do more hiring are more likely to integrate because they have more opportunities to do so.

Let  $w$  index hires after Nitaqat,  $i$  index firms, and  $W_i \geq 0$  denote the number of hires a firm makes without integrating. The sample includes all hires made after June 2011 at firms that are “at risk” of integrating, i.e. all hires at firms that have not previously hired a woman. We initially restrict to firms that have not hired a Saudi woman prior to June 2011. We also restrict to firms that have at least 1 Saudi employee in May 2011.<sup>23</sup> We estimate a discrete hazard  $h_{iw} = \mathbb{P}(W_i = w | W_i \geq w)$  where it assumed that  $h_{iw}$  follows a logistic distribution. In other words, restricting to the sample described above, we estimate a hire-level logistic regression:

$$\begin{aligned} \mathbb{P}(\text{Female}_{iw} = 1 | \text{Below Quota}, X, Z; W_i \geq w) = & \Lambda(\beta \text{Below Quota}_i & (2) \\ & + \gamma \text{Log (Distance Below Quota + 1)}_i \\ & + X_{iw}\delta + Z_i\lambda | W_i \geq w) \end{aligned}$$

where the outcome  $\text{Female}_{iw}$  is an indicator for whether hire  $w$  at firm  $i$  is female,  $\text{Below Quota}$  is an indicator for below quota,  $\text{Distance Below Quota}$  is the number of Saudi employees a firm would need to hire to meet their Nitaqat quota at baseline,  $X$  are job-level characteristics corresponding to the hired position, and  $Z$  are firm-level characteristics. Job-level controls include the female share of Saudis employed in that industry by occupation combination at baseline and the female-male log earnings gap in that cell at baseline, where a negative number indicates that men have higher average earnings in that position. Firm-level controls include log firm size at baseline, where size includes both Saudi and foreign workers.  $Z$  also includes firm Saudi workforce characteristics at baseline, including average age, share married, share with secondary education, share with tertiary education. Descriptive statistics for each hire are presented in [Appendix Table A3](#).

[Table 4 about here.]

The estimated marginal effects evaluated at means are presented in [Table 4](#). Columns (1) and (2) include only firm-level controls; columns (3) and (4) include both job- and firm-level controls. Column (2) includes industry fixed effects and column (4) includes both occupation and industry fixed effects.

Across specifications, the coefficients on  $\text{Below Quota}$  and  $\text{Log (Distance Below Quota + 1)}$  are positive and statistically significant. Conditional on firm and job characteristics, the rate of integration is increasing in the number of Saudis a firm must hire to meet the quota. For  $\text{Below Quota}$ , the marginal effect ranges from 0.0035 to 0.0042. For  $\text{Log (Distance Below Quota + 1)}$ , the marginal effect ranges from 0.0017 to 0.0033. A 10 log point increase in the  $\text{Distance Below Quota}$  is associated with a 0.0002 to 0.0003 increase in the hazard rate.  $\text{Below Quota}$  firms have an average value for  $\text{Log (Distance Below Quota + 1)}$  of 2.23 (and an average  $\text{Distance Below Quota}$  of 46). Using estimates from column (4), the average  $\text{Below Quota}$

<sup>23</sup>Results are similar if we include firms that exist in May 2011 but do not have any Saudi employees.

firm is predicted to have a hazard rate that is about  $0.0035 + 2.23 * 0.0032 = 0.011$  larger than that of a comparable Above Quota firm, or 20% higher than the baseline hazard rate at Above Quota firms (0.056).

We interpret this as evidence that Nitaqat's causal effect of the female share of the Saudi workforce is at least in part driven by changes in firm demand. We also find that firms are more likely to integrate when hiring in jobs that historically have higher female shares and larger wage gaps. The coefficient of -0.0181 on Log Wage Gap in column (4) implies that a 10% increase in the wage gap is associated with about a 0.0018 increase in the hazard rate, a 3.2% increase relative to the baseline hazard rate.

One concern with the above results is that there may be unobserved differences between above and below quota firms that generate differences in integration rates between the two sets of firms, rather than their positions relative to the quota *per se*. As an additional check, we conduct a placebo test. In the model, Nitaqat increases relative demand for female workers; however, this effect is only present among firms that have not previously integrated. Integrated firms have already paid their fixed integration and so will continue along the integrated production frontier regardless of the quota pressure they face. However, if the pattern we observe in Table 4 is driven by unobserved differences between firms rather than quota pressure, we may expect to see the same pattern among firms that have already integrated by May 2011.

We re-estimate (2), this time restricting to firms to. As above, the event of interest is the first female hire after the implementation of Nitaqat. The results are presented in Table 5.

[Table 5 about here.]

By contrast to Table 4, we find no relationship between Below Quota and integration in this sample. This finding is both consistent with the model and supports a causal interpretation of the baseline hazard results.

## 6 Ghost Employment

One potential concern is that firms might falsify their employee records with GOSI to meet their quotas after Nitaqat, so these 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.<sup>24</sup> This ghost employment would cause our analysis to overstate the degree to which firms hire Saudi women in response to employment quotas.

---

<sup>24</sup>There 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.

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.<sup>25</sup>

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.<sup>26</sup>

Figure 12 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 percent for men and 47 percent 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 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.

[Figure 12 about here.]

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

---

<sup>25</sup>One 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.

<sup>26</sup>Firms 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.

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 13 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 13 about here.]

## 7 Conclusion

Although women’s employment in Saudi Arabia has historically been very low, the country is currently experiencing a transformation in women’s economic engagement. Female labor force participation rates grew from 10.1 percent to 17.8 percent from 2000 to 2017. As in many countries, however, employment opportunities for women have lagged behind their growth in participation. We documented a particular type of demand-side constraint to employment opportunities, namely that many private-sector firms hire only men and may face fixed costs to employ female workers. This fixed cost may be due to legal and cultural norms about workplace interaction, a lack of physical infrastructure to comply with legal regulations such as restroom accessibility and office space, or to a lack of hiring experience by the firms.

This paper investigates the integration of previously all-male firms. We develop a simple search model to study the effects of fixed integration costs on female employment. The model predicts that firms will only hire women if they expect to spread these fixed integration costs over a sufficient number of female employees. We derive a set of four testable predictions from this framework: (1) the distribution of female workers across firms exhibits bunching at zero and a ‘missing mass’ of firms that employ small numbers of female employees; (2) the distribution of female *hiring* is similarly ‘lumpy’, and this pattern is concentrated at previously all-male firms; (3) quotas increase relative labor demand for women at all-male firms, but (4) do not affect relative demand at integrated firms. We also find suggestive evidence that binding gender-neutral quotas at some firms induce integration at firms where quotas are not binding, consistent with demand externalities or changing norms.

Our empirical analysis tests these predictions using administrative employer-employee data and the introduction of the Nitaqat policy. We document that both baseline integration patterns and the response to the Nitaqat quotas all support the idea that firms face fixed integration costs. In particular, we find that firms are substantially segregated (particularly at baseline), and that there are relatively few firms with 1-4 female employees compared with our simulated benchmark. We also find that firm hiring behavior changes dramatically following integration: firms that integrate continue to hire women at a substantial rate in the months following their first female hire. Nitaqat quotas also have the expected effect on female

employment: segregated firms facing larger Saudi hiring requirements were more likely to integrate, and already-integrated firms did not adjust their employee gender mix in response to the quotas.

The idea that firms face these fixed integration costs has important implications, not only for female employment dynamics but also for policy. In particular, this evidence suggests that policies like Nitaqat that push firms to overcome these integration costs can have permanent effects. Addressing these costs directly may also have large effects on firms' ability to make these investments, potentially leveraging the employment effects of these policies. More detailed work identifying these costs would help inform labor policy to ease this transition and further increase female employment.

## References

- Acemoglu, Daron and Joshua D. Angrist**, "Consequences of Employment Protection? The Case of the Americans with Disabilities Act," *Journal of Political Economy*, October 2001, 109 (5), 915–957.
- 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," February 2018.
- 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.
- Betrand, Marianne**, "New Perspectives on Gender," in David Card and Orley Ashenfelter, eds., *Handbook of Labor Economics*, Vol. 4, Elsevier, 2011, chapter 7, pp. 1543–1590.
- Black, Dan A.**, "Discrimination in an Equilibrium Search Model," *Journal of Labor Economics*, April 1995, 13 (2), 309–334.
- Brinton, Mary C. and Eunmi Mun**, "Between State and Family: Managers' Implementation and Evaluation of Parental Leave Policies in Japan," *Socio-Economic Review*, 2016, 14 (2), 257–281.
- 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.
- 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.
- 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.
- GaStat**, *Labor Force Survey 2011*, General Authority for Statistics, Kingdom of Saudi Arabia, 2011.
- Goldin, Claudia**, “A Grand Gender Convergence: Its Last Chapter,” *American Economic Review*, April 2014, 104 (4), 1091–1119.
- 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.
- Hellerstein, Judith K. and David Neumark**, “Workplace Segregation in the United States: Race, Ethnicity, and Skill,” *The Review of Economics and Statistics*, 2008, 90 (3), 459–477.
- Hotchkiss, Julie L. and Ludmilla Rovba**, “Employment,” in Julie L. Hotchkiss, ed., *Labor Market Experience of Workers with Disabilities: The ADA and Beyond*, W. E. Upjohn Institute for Employment Research, 2003, chapter 2, pp. 21–47.
- International Finance Corporation**, *Investing in Women’s Employment : Good for Business, Good for Development*, World Bank, 2013. <https://openknowledge.worldbank.org/handle/10986/16257>.
- Jenkins, Stephen P.**, “Easy Estimation Methods for Discrete-Time Duration Models,” *Oxford Bulletin of Economics and Statistics*, February 1995, 57 (1), 129–136.
- 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>.
- 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.

**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.

**Wooten, Lynn Perry**, “What Makes Women-Friendly Public Accounting Firms Tick? The Diffusion of Human Resource Management Knowledge Through Institutional and Resource Pressures,” *Sex Roles*, Sep 2001, 45 (5), 277–297.

**World Bank**, *World Development Indicators 2016*, The World Bank, 2016.


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

Figure 1: Saudi Unemployment by Gender and Education, 2011




*Note:* This figure reports the official Saudi unemployment statistics from the Saudi Government’s Labor Force Survey for 2011 by gender and education (GaStat 2011). Panel (a) shows the share of Saudi nationals in the labor force who are unemployed, and panel (b) the number of unemployed Saudi nationals in each category. The category Primary- includes respondents whose highest education level is listed as illiterate, read and write, or primary. Masters+ includes those with a Master’s degree (or equivalent “higher diploma”) or a doctorate. All other categories correspond directly with those in the Labor Force Survey.

Figure 2: Change in Private Sector Employment, June 2011 to June 2015




*Note:* This figure reports the changes in private sector Saudi employment counts by gender and education between June 2011 and June 2015 as reported in the GOSI data. Education is reported for 78 percent of the observations across the two years. Masters+ includes those with a Master's degree or a PhD. Primary- includes those with elementary schooling and those listed as illiterate. All other categories correspond with those collected by GOSI.

Figure 3: Percent Change in Total Saudi Employment Relative to June 2011




*Note:* This set of figures compares the percent change in Saudi employees by firms' quota status relative to our baseline of June 2011, separately for male and female employees. Firms above the quota are Green and Platinum firms, and firms below are Yellow and Red firms. Source: GOSI

Figure 4: Distribution of Female Employment Across Firms, January 2009



*Note:* This set of figures compares the distribution of female employment across firms in January 2009 to a simulated distribution where workers are randomly re-assigned across jobs within location by occupation cells. Panel A plots the share of firms with zero female employees in both the actual and simulated distributions. Panel B plots the distribution of female employees across firms. Overall, a Kolmogorov-Smirnov test rejects equality of distributions at the 0.01 significance level. Panel C groups firms by their expected number of female employees, and compares the share of firms that have any female employees by bin. More details on the simulation and construction of these figures is provided in Section 5.1.


Figure 5: Distribution of Female Hiring Across Firms, Non-Legacy Firms




*Note:* This set of figures compares the distribution of female employment across non-legacy firms in June 2015 among employees hired since June 2011 to a simulated distribution where workers are randomly re-assigned across jobs within location by occupation cells. We restrict to firms that employed any Saudis in January 2009 and have at least 5 employees in June 2015 that were hired since 2011. Non-legacy firms are firms that did not employ a Saudi woman in January 2009. Panel A plots the share of firms with zero female employees in both the actual and simulated distributions. Panel B plots the distribution of female employees across firms. Overall, a Kolmogorov-Smirnov test rejects equality of distributions at the 0.01 significance level. Panel C groups firms by their expected number of female employees, and compares the share of firms that have any female employees by bin. More details on the simulation and construction of these figures is provided in Section 5.2.

Figure 6: Distribution of Female Hires Across Firms, Legacy Firms

(a) Percentage of Firms with Zero Female Employees      (b) Distribution of Female Employment Across Firms with Any Female Hires




(c) Any Female Hires by Expected Number of Female Hires



*Note:* This set of figures compares the distribution of female employment across legacy firms in June 2015 among employees hired since June 2011 to a simulated distribution where workers are randomly re-assigned across jobs within location by occupation cells. We restrict to firms that employed any Saudi in January 2009 and have at least 5 employees in June 2015 that were hired since 2011. Legacy firms are firms that employed at least one Saudi woman in January 2009. Panel A plots the share of firms with zero female employees in both the actual and simulated distributions. Panel B plots the distribution of female employees across firms. Overall, a Kolmogorov-Smirnov test rejects equality of distributions at the 0.01 significance level. Panel C groups firms by their expected number of female employees, and compares the share of firms that have any female employees by bin. More details on the simulation and construction of these figures is provided in Section 5.2.


Figure 7: Female Share of Hires at Integrating Firms




*Note:* This figure plots the percent female of hires made at integrating firms in six-month periods 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. We plot the composition of hires separately for firms that integration before and after the implementation of Nitaqat in June 2011. 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.

Figure 8: Saudi Women in the Private Sector Over Time

(a) Percent Female of Saudi Private Sector Workforce




(b) Percent of Firms with Any Saudi Female Employees



Note: Panel A plots the female share of Saudi employment in the private sector over time. Panel B plots the share of firms with any Saudi female employees, among firms with any Saudi employees. Plotted data are monthly.


Figure 9: Gross Saudi Hires Over Time by Legacy Status




*Note:* This figure plots the number of Saudis hired in the private sector over time at Yellow and Red ('below quota') firms and Green and Platinum ('above quota') firms. The data plotted are quarterly. Panel A restricts to non-legacy firms with any Saudi employee in January 2009. Legacy firms are firms with any Saudi female employee in January 2009. Panel B restricts to legacy firms.

Figure 10: Female Share of Hires by Legacy Status

(a) Non-Legacy Firms




(b) Legacy Firms




*Note:* This figure plots the female share of Saudi hires in the private sector over time at Yellow and Red ('below quota') firms and Green and Platinum ('above quota') firms. The data plotted are quarterly. Panel A restricts to non-legacy firms with any Saudi employee in January 2009. Legacy firms are firms with any Saudi female employee in January 2009. Panel B restricts to legacy firms.

Figure 11: Female Share of Hires by Legacy Status

(a) Male Hiring



(b) Female Hiring




Note: This figure plots the  $\beta_t$  regression coefficients from the model (1) and their 95% confidence intervals. The model (1) is described in more detail section 5.3. In Panel A, the outcome is male hires. In Panel B, the outcome is female hires. The average number of male and female hires per quarter at above quota firms is 1.8 and 0.26.

Figure 12: 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 dummy variables.

Figure 13: 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. Dummy variables are used to flexibly control for age, education, and marital status of new hires. Dashed lines show the 95% confidence interval for month dummy variables.

Table 1: Occupations with Largest Female Employment Growth

Occupation	ISCO Code	June 2011		June 2015		$\Delta$ Female Employees	
		Female	Male	Female	Male	Female	% $\Delta$
<i>ISCO-08 Two-digit Occupations</i>							
Business and administration associate professionals	33	4,692	31,855	81,061	99,076	76,369	1,628
General and keyboard clerks	41	8,895	75,511	81,687	122,781	72,792	818
Customer services clerks	42	9,238	55,027	47,423	92,759	38,185	413
Sales workers	52	3,242	29,701	33,167	55,204	29,925	923
Numerical and material recording clerks	43	804	12,257	28,134	28,042	27,330	3,399
Chief executives, senior officials and legislators	11	1,144	18,066	17,974	47,273	16,830	1,471
Administrative and commercial managers	12	1,868	19,411	18,186	40,820	16,318	874
Teaching professionals	23	9,653	6,410	24,810	9,542	15,157	157
Unclassified	99	2,389	45,993	12,852	134,328	10,463	438
Business and administration professionals	24	937	7,514	11,351	15,998	10,414	1,111
<i>ISCO-08 Occupation Groups</i>							
Clerical support workers	4	19,573	150,113	167,432	259,541	147,859	755
Technicians and associate professionals	3	6,543	60,926	91,595	156,807	85,052	1,300
Service and sales workers	5	15,301	110,045	59,714	182,749	44,413	290
Professionals	2	15,360	48,216	59,360	88,831	44,000	286
Managers	1	3,640	42,682	39,101	97,605	35,461	974
Elementary occupations	9	8,357	155,279	37,946	321,338	29,589	354
Craft and related trades workers	7	1,419	49,590	4,522	62,611	3,103	219
Plant and machine operators, and assemblers	8	795	37,361	3,162	53,964	2,367	298
Skilled agricultural, forestry and fishery workers	6	141	222	132	321	-9	-6

*Note:* The first half of this table represents the occupations (as defined by ISCO two-digit codes) with the largest increase in female employment from June 2011 to June 2015. The second half of this table provides this same information at a more aggregated level at the ISCO one-digit (or major group) level. Source: GOSI

Table 2: Distribution of Female Employment Across Firms, January 2009

	Actual	Baseline	Location	Occupation	Combined
<i>Number of Female Employees</i>					
0	72.3	32.4	32.0	37.8	38.2
1	5.6	27.2	25.1	25.6	24.9
2	3.5	14.1	14.2	12.0	12.0
3	2.9	7.3	7.8	6.0	6.1
4	2.1	4.3	4.7	3.5	3.6
5	1.8	2.8	3.1	2.4	2.4
6-10	4.5	6.1	6.7	5.5	5.5
11-24	3.5	3.9	4.2	4.1	4.2
25+	4.0	2.0	2.2	3.1	3.2

*Note:* This table summarizes the share of firms with different numbers of female employees, both in the actual data and simulated data described in further detail in section 5.1. We simulate the distribution of female employment across firms, taking each firm’s size as given but randomly re-assigning workers across firms. Across simulation specifications, we vary the units within which workers are randomly assigned. The first simulation (‘Baseline’) bundles all workers into a single group, reassigning all workers at random to different firm positions. The second simulation (‘Location’) differentiates workers and jobs by location, reassigning workers to positions at firms within local labor markets. The third simulation (‘Occupation’) differentiates workers and jobs by their location. The fourth benchmark (‘Combined’) differentiates workers and jobs by both location and occupation.

Table 3: Firm Descriptive Statistics by Integration Status

	Never integrated		Legacy integrators		Newly integrated	
	June, 2011	June, 2015	June, 2011	June, 2015	June, 2011	June, 2015
<i>All baseline firms</i>						
Mean firm size	1.5	2.6	62.0	89.6	1.6	5.4
Mean female share	0	0	41.7	46.7	2.8	58.9
Median female share	0	0	33.3	46.5	0	60
Female share of Employment	0	0	13.2	20.8	1.3	33.0
Pct. of Fem. Increase		0		31.6		68.4
% of firms in Retail		6.7		18.6		10.7
% of firms in Construction		14.9		17.2		16.9
% of firms in Manufacturing		2.7		6.7		2.9
% of firms in Oil, Gas, & Petrochemicals		0.08		0.6		0.07
% White		62.4		21.4		57.3
% Red		25.4		29.1		30.0
% Yellow		3.6		15.1		4.6
% Green		7.8		26.0		7.6
% Platinum		0.7		8.4		0.5
# of firms		46,666		4,748		61,112
<i>Firms with at least one Saudi employee at baseline</i>						
Mean firm size	5.7	7.3	66.6	95.6	8.1	17.8
Mean female share	0	0	41.7	45.3	2.8	37.7
Median female share	0	0	33.3	44.4	0	33.3
Female share of Employment	0	0	13.2	20.5	1.3	22.6
Pct. of Fem. Increase		0		50.0		50.0
% of firms in Retail		12.4		19.6		19.5
% of firms in Construction		24.3		17.7		33.1
% of firms in Manufacturing		6.8		7.0		8.8
% of firms in Oil, Gas, & Petrochemicals		0.2		0.6		0.3
% White		38.3		18.0		18.2
% Red		30.6		29.7		36.9
% Yellow		7.4		15.8		14.8
% Green		21.0		27.7		27.7
% Platinum		2.8		8.8		2.4
# of firms		11,994		4,423		12,135

*Note:* This table present descriptive statistics of firms subject to Nitaqat quotas in June, 2011, the first month the quotas were instituted ("baseline firms"). The table is divided into three firms integration categories: firms which have not hired women from January 2009 to June 2015 ("Never integrated"), firms which hired at least one Saudi female employee before June 2011 when the quotas were introduced ("Legacy Integrators"), and lastly, firms which hired at least one Saudi female employee on or after June 2011, when the quotas were imposed ("Newly integrated"). The information is further broken up by firms who did not have a Saudi in June 2011, our baseline, and those who had at least one. The "Pct. of Fem. Increase" line indicates the percentage of the increase in female hiring between June 2011 and June 2015 each category of firms hired. The other percentages listed indicate the fraction of firms in that integration category that correspond to the relevant indicator (ie. "% Red" under the "Newly Integrated" column indicates the percentage of firms that are Newly Integrated that are designated red for Nitaqat purposes). Source: GOSI

Table 4: Segregated Firms: Do Below Quota Firms Integrate More Quickly Than Above Quota Firms?

	(1)	(2)	(3)	(4)
Below Quota	0.0042*** (0.0011)	0.0039*** (0.0010)	0.0041*** (0.0010)	0.0035*** (0.0009)
Log (Saudi Distance + 1)	0.0021*** (0.0004)	0.0033*** (0.0004)	0.0017*** (0.0004)	0.0032*** (0.0004)
Baseline firm characteristics:				
Avg. Age	0.0000 (0.0001)	0.0000 (0.0001)	0.0000 (0.0001)	0.0000 (0.0001)
Share Married	0.0029* (0.0016)	0.0028* (0.0015)	0.0032** (0.0016)	0.0028* (0.0014)
Log firm Size (NQ)	-0.0164*** (0.0004)	-0.0162*** (0.0004)	-0.0158*** (0.0004)	-0.0154*** (0.0004)
Share Sec. Edu	0.0052*** (0.0012)	0.0040*** (0.0012)	0.0042*** (0.0012)	0.0034*** (0.0011)
Share Ter. Edu	-0.0028 (0.0024)	0.0028 (0.0024)	-0.0103*** (0.0026)	0.0012 (0.0022)
Fem. Share of Ind./Occ.			0.0256*** (0.0035)	0.0568*** (0.0046)
Log Wage Gap			-0.0123*** (0.0012)	-0.0181*** (0.0020)
local labor market fixed effects?	✓	✓	✓	✓
Industry fixed effects?		✓		✓
Observations	199418	199418	199418	199418

*Note:* This table presents marginal effect estimates for (2), a discrete 'time' hazard model for the first female hire following the implementation of Nitaqat. Each observation is a hire that firm makes, so that 'duration' in this hazard is measured in number of hires. The model is discussed in further detail section 5.3. The models summarized in this table are estimated using firms that (1) had at least once Saudi employee as of May 2011 and (2) had never employed a Saudi woman between January 2009 and May 2011. 'Saudi Distance' measures the number of Saudis that a firm would have to be hired to meet the quota in June 2011 when the quotas were first imposed. 'Saudi Distance' is zero for firms that satisfy the quota in June 2011. We take the natural log of this distance metric plus one. Covariates 'Average Age', 'Log Firm Size', 'Share Secondary Education', and 'Share Tertiary Education' are measured at the firm as of May 2011. 'Log Firm Size' is measured in the Nitaqat data and includes both Saudi and non-Saudi employees. Covariates 'Female Share of Industry/Occupation' and 'Log Wage Gap' refer to the specific job for which a hire is made. 'Female Share of Industry/Occupation' is the female share of employees and 'Log Wage Gap' is the log wage gap in a given industry by occupation cell as measured in the data from January 2009 to May 2011. The hazard rate for Above Quota firms included in the estimation sample is 0.056.

Table 5: Integrated Firms: Do Below Quota Firms Hire Women More Quickly Than Above Quota Firms?

	(1)	(2)	(3)	(4)
Below Quota	0.0019 (0.0022)	0.0003 (0.0019)	0.0012 (0.0020)	0.0010 (0.0017)
Log (Saudi Distance + 1)	0.0007 (0.0009)	0.0009 (0.0007)	-0.0003 (0.0009)	0.0007 (0.0007)
Baseline firm characteristics:				
Avg. Age	0.0006*** (0.0002)	0.0003* (0.0002)	0.0004** (0.0002)	0.0003* (0.0002)
Share Married	0.0062 (0.0045)	0.0101** (0.0042)	0.0115*** (0.0040)	0.0111*** (0.0037)
Log firm Size (NQ)	-0.0147*** (0.0009)	-0.0125*** (0.0008)	-0.0125*** (0.0008)	-0.0112*** (0.0007)
Share Sec. Edu	0.0207*** (0.0054)	0.0123*** (0.0047)	0.0121** (0.0052)	0.0097** (0.0042)
Share Ter. Edu	0.0579*** (0.0073)	0.0432*** (0.0062)	0.0221*** (0.0065)	0.0331*** (0.0053)
Fem. Share of Ind./Occ.			0.0796*** (0.0072)	0.0998*** (0.0096)
Log Wage Gap			-0.0044 (0.0027)	-0.0211*** (0.0036)
local labor market fixed effects?	✓	✓	✓	✓
Industry fixed effects?		✓		✓
Observations	58549	58549	58549	58549

*Note:* This table presents marginal effect estimates for (2), a discrete 'time' hazard model for the first female hire following the implementation of Nitaqat. Each observation is a hire that firm makes, so that 'duration' in this hazard is measured in number of hires. The model is discussed in further detail section 5.3. The models summarized in this table are estimated using firms that (1) had at least once Saudi employee as of May 2011 and (2) had employed a Saudi woman between January 2009 and May 2011. 'Saudi Distance' measures the number of Saudis that a firm would have to be hired to meet the quota in June 2011 when the quotas were first imposed. 'Saudi Distance' is zero for firms that satisfy the quota in June 2011. We take the natural log of this distance metric plus one. Covariates 'Average Age', 'Log Firm Size', 'Share Secondary Education', and 'Share Tertiary Education' are measured at the firm as of May 2011. 'Log Firm Size' is measured in the Nitaqat data and includes both Saudi and non-Saudi employees. Covariates 'Female Share of Industry/Occupation' and 'Log Wage Gap' refer to the specific job for which a hire is made. 'Female Share of Industry/Occupation' is the female share of employees and 'Log Wage Gap' is the log wage gap in a given industry by occupation cell as measured in the data from January 2009 to May 2011.

Table A1: 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:* The table presents the number of Saudi employees by the 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. We do not have a way of ascertaining what these occupations may have been.

Table A2: Descriptive Statistics for Difference-in-Difference Sample Firms, June 2011

	Non-Legacy		Legacy	
	Below	Above	Below	Above
Mean firm size	13.5	20.1	140.6	116.9
Mean female share	2.4	2.4	37.4	35.3
Median female share	0	0	30.8	25
% of firms in Retail	1.5	1.7	0.3	0.4
% of firms in Construction	32.1	37.0	17.3	16.9
% of firms in Manufacturing	1.8	2.1	1.8	2.2
% Red	73.5	0	66.7	0
% Yellow	26.5	0	33.3	0
% Green	0	86.8	0	73.1
% Platinum	0	13.2	0	26.9
# of firms	6,679	4,036	1,001	777

*Note:* This table describes the firms included in the firm-level difference-in-difference model described in section 5.3. The estimation sample is restrict to Red, Yellow, Green, and Platinum firms that have at least one Saudi employee in January 2009. Legacy firms are firms with any Saudi female employee in January 2009.

Table A3: Descriptive Statistics for Hazard Model Hires

	Segregated in May 2011	Integrated in May 2011
Below Quota	0.7 (0.5)	0.6 (0.5)
Saudi Distance + 1	24.9 (121.5)	144.7 (530.5)
Log(Saudi Distance + 1)	1.2 (1.6)	1.5 (2.4)
Fem. Share of Ind./Occ.	0.06 (0.10)	0.06 (0.1)
Log Wage Gap	-0.002 (0.3)	0.02 (0.3)
Avg. Age	22.0 (14.2)	29.1 (5.7)
Share Married	0.1 (0.2)	0.2 (0.1)
Log Firm Size (NQ)	5.0 (1.8)	6.6 (1.8)
Share Sec. Edu	0.4 (0.3)	0.6 (0.2)
Share Ter. Edu	0.04 (0.1)	0.07 (0.1)

*Note:* This table describes the firm- and job-level characteristics for the firms included in the hire-level hazard model described in section 5.3. The estimation sample is restrict to Red, Yellow, Green, and Platinum firms that have at least one Saudi employee in May 2011. We further split the sample by whether the firm had any Saudi female employee by May 2011. Firm characteristics are measured as of May 2011. ‘Log Firm Size’ is measured in the Nitaqat data and includes both Saudi and non-Saudi employees. ‘Saudi Distance’ is zero for firms that satisfy the quota in June 2011. We take the natural log of this distance metric plus one. Covariates ‘Female Share of Industry/Occupation’ and ‘Log Wage Gap’ refer to the specific job for which a hire is made. ‘Female Share of Industry/Occupation’ is the female share of employees and ‘Log Wage Gap’ is the log wage gap in a given industry by occupation cell as measured in the data from January 2009 to May 2011.