

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Economics

School of Economics

8-2022

Intra-firm hierarchies and gender gaps

Nicolo DALVIT
World Bank

Aseem PATEL
University of Essex

Joanne TAN
Singapore Management University, joannetany@smu.edu.sg

Follow this and additional works at: https://ink.library.smu.edu.sg/soe_research



Part of the [Behavioral Economics Commons](#), [Industrial Organization Commons](#), and the [Organization Development Commons](#)

Citation

DALVIT, Nicolo; PATEL, Aseem; and TAN, Joanne. Intra-firm hierarchies and gender gaps. (2022). *Labour Economics*. 77, 1-16.

Available at: https://ink.library.smu.edu.sg/soe_research/2594

This Journal Article is brought to you for free and open access by the School of Economics at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Economics by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.



ELSEVIER

Contents lists available at ScienceDirect

Labour Economics

journal homepage: www.elsevier.com/locate/labecoIntra-firm hierarchies and gender gaps[☆]Nicolò Dalvit^a, Aseem Patel^b, Joanne Tan^{c,*}^a World Bank, 1818 H Street, NW Washington, DC, USA^b University of Essex and Sciences Po, 28 Rue des Saints-Pères, Paris, France^c Singapore Management University, 90 Stamford Road, Singapore

ARTICLE INFO

JEL classification:

G38
J16
J3
L2

Keywords:

Gender gap
Firm organisation
Gender quota
Trickle-down effect

ABSTRACT

We study how changes in female representation at the top of a firm's organisation affect gender-specific outcomes across hierarchies within firms. We start by developing a theoretical model of a hierarchical firm, where gender representation in top organisational layers can affect gender-specific hiring and promotion probabilities at lower layers. We then exploit a recent French reform that imposed gender representation quotas in the boards of directors and test the model's predictions in the data. Our empirical results show that the reform was successful in reducing gender wage and representation gaps at the upper layers of the firm, but not at lower firm layers. A Panel VAR analysis confirms that the trickle-down effects of this policy were limited and suggests that quotas targeting middle management, rather than corporate boards, may have a more widespread effect across the firm.

1. Introduction

Gender differences in labour force participation, education and political participation have narrowed in recent decades. Despite these significant advances, gender disparities in the labour market remain large.¹ This is particularly evident in top leadership positions, where women continue to be severely under-represented and where gender wage gaps remain substantial. To address these issues, many European countries

[☆] We thank Jean-Marc Robin, Ghazala Azmat, Pierre Cahuc, Cecilia Garcia-Peñalosa, Vincent Sterk, seminar participants at Sciences Po, the University of Michigan, the 2018 Econometric Society European Winter meeting, the 2018 Latin American Meeting of the Econometric Society and the 2019 and 2020 meetings of the Society of Labor Economists. We also thank the editors Maia Güell and Claudia Olivetti, as well as two anonymous referees for their comments. We thank Sciences Po for funding this project. This work was supported by a public grant overseen by the French National Research Agency (ANR) as part of the Investissements d'avenir program (reference: ANR-10-EQPX-17 Centre d'accès sécurisé aux données CASD). The views expressed in the paper are those of the authors and should not be attributed to the World Bank. Any fault is our own.

* Corresponding author.

E-mail addresses: nicolo.dalvit@gmail.com (N. Dalvit), aseem.patel@sciencespo.fr (A. Patel), joannetanym@smu.edu.sg (J. Tan).

¹ The literature on the evolution and causes of the gender wage gap is well-documented. A good summary of this literature can be found in [Blau and Kahn \(2016\)](#). Reasons cited to explain why such a wage gap exists in the first place includes differences in human capital, experience ([Olivetti, 2006](#)), choice of occupation and industry ([Mulligan and Rubinstein, 2008](#)), labour force participation decisions ([Goldin et al., 2017](#)) and discrimination ([Becker, 1971](#)).

<https://doi.org/10.1016/j.labeco.2021.102029>

Received 29 September 2020; Received in revised form 24 June 2021; Accepted 29 June 2021

0927-5371/© 2021 Elsevier B.V. All rights reserved.

such as Norway, Italy, Belgium and France have passed laws introducing gender quotas in the corporate boards of directors of publicly-listed firms.²

In this paper, we investigate if greater female representation in top management positions raises relative female employment and wages in the rest of the firm. In other words, does having more women with decision-making power at the top of the firm hierarchy lead to improved outcomes for women further down the firm hierarchy?

To answer this question, we first construct a stylised model of statistical discrimination. We model firms as organisations with distinct hierarchical layers. Within each layer individuals perform tasks that vary in complexity. Hiring and promotion decisions for the individuals in each layer are taken by managers higher up in the firm hierarchy. Managers only have imperfect information about potential candidates since their skills are only partly observed at the point of making these decisions.³

We introduce gender differences by allowing agents in the model, i.e. workers and their managers, to be one of two types: male or female. Following the recent work of ([Flabbi et al., 2019](#)), we assume that female managers assess the skills of other females with greater precision compared to male managers. Under the plausible assumption that profits from a worker-job match are concave in workers' skills, our model predicts that female managers are relatively more likely to hire and pro-

² The state of California in the US has also implemented similar gender quota laws on the corporate boards of publicly-listed firms in the state.

³ There is a separate theoretical literature that considers a matching model of the labour market in the presence of uncertainty and learning about ability that provides a unified framework for analysing the dynamics of jobs and wages within firms and in the labour market. See, for instance, [Pastorino \(2015\)](#).

mote female candidates. Consequently, a greater share of women in a given layer increases the share of women hired and promoted in the adjacent subordinate layer, which in turn affects the next layer below. As such, the impact of higher female representation at the top of the firm on gender gaps may trickle down the firm hierarchy. The extent of these trickle-down effects is determined by exogenous job turnover and promotion rates within firms.

We then test our model's predictions empirically by estimating the effect of the share of women on corporate boards on gender gaps at each layer of the firm hierarchy using French administrative data. To do so, we exploit a 2010 reform imposing gender quotas on corporate boards in France. This law, called the *Loi Copé-Zimmerman*, mandated publicly-traded firms to have at least 20% women on the board of directors by the 1st of January 2014 and 40% by the 1st of January 2017.⁴ Firms that failed to comply with the reform would have been fined, dissolved or banned from paying existing directors. Crucially, the implementation of the *Loi Copé-Zimmerman* generates instruments for the share of women on corporate boards, which is likely endogenous. Following (Ahern and Dittmar, 2012) and (Bertrand et al., 2018), we instrument the shares of female board members using the share of female board members in 2010, the year prior to the implementation of the reform. The intuition behind our instrumental variable strategy is simple: companies that started with a lower share of women on their board have had to increase their share of women relatively more to comply with the quota.

We find that, consistent with the predictions of our model, a rise in female board membership (layer 1) narrows gender wage and employment gaps at the top layers of the firm, namely among senior executives and professionals (layer 2) as well as among middle management (layer 3). For instance, an increase in the share of women on corporate boards by 10 percentage points (pp) has a statistically significant effect of raising the share of women by 2.4 and 2.5 pp in layers 2 and 3 respectively. In addition, it lowers the gender wage gap in layer 2 significantly by approximately 2.4 percent.

In our analysis of the reform period, we find that no statistically significant impact is observed on gender and representation gaps at the lowest layer (layer 4), which comprises of administrative, sales, security and blue-collar workers. To explore the extent of trickle-down effects to lower layers in the long term, we run a panel vector auto-regression analysis (P-VAR) on our data from 1999 to 2016. We find that even over a ten-year period, an increase in female board share has limited trickle-down effects on the lowest layers of the firm. A counterfactual exercise using the estimates from our P-VAR suggests that an increase in the share of women in middle management (layer 3) would have a greater impact on gender gaps at the lowest layer of the firm. As such, while corporate board quotas do mitigate gender gaps at upper layers of firms where these gaps are widest, other policies should be considered if the aim is to improve the labour market outcomes of a broader section of women.

We make two main contributions in this paper. First, we construct a theoretical model, which formalizes how an increase in female representation at the top of the firm hierarchy can affect gender-specific outcomes such as gender wage gaps and representation gaps at lower layers. Second, we contribute to the literature by studying the impact of gender quota reforms on gender-specific outcomes at each hierarchical layer of the firm. Previous literature has highlighted that these changes have no statistically significant effects at the firm level. In contrast, we show that changes in gender quotas improving female representation at the top of the firm hierarchy have differing effects on gender-specific outcomes across the lower layers of the firm hierarchy.

Related Literature A growing literature has focused on the role of female leadership in determining labour market outcomes for females,

⁴ The text of the law and further details can be found here: <https://www.entreprises.gouv.fr/politique-et-enjeux/mixite-et-egalite-professionnelle-dans-entreprises>

albeit with little consensus. One strand of this literature has identified a positive impact of female leadership on labour market outcomes of women. For instance, Cornell and Welch (1996) find that women in general are likely to be more sensitive to the issue of gender representation, discriminate less and be able to better assess their female co-workers relative to men. In the corporate setting, authors including (Cardoso and Winter-Ebmer, 2010), Bhide (2019), Cardoso and Winter-Ebmer (2010) and Kunze and Miller (2017) have found a general positive impact of female leadership on the wages and representation of female subordinates. A similar literature in political economy has discussed the positive impact of female political leadership on gender bias (e.g. Beaman et al., 2009 and Gagliarducci and Paserman, 2012).

On the other hand, Adams and Funk (2012) have presented evidence showing that women that do break the glass ceiling and make it into top management positions might behave as their male counterparts and may not necessarily have a positive impact on female labour market outcomes. Likewise, recent evidence from academic committees shows that having a higher share of women can have a negative effect on the probability that a woman is hired (Bagues et al., 2017, Deschamps, 2018).

We contribute to this on-going literature, both theoretically and empirically, by examining if a greater share of women on corporate boards significantly mitigates gender wage and representation gaps in firms. At a theoretical level, our work contributes to two strands of the literature. First, our model borrows from the literature on internal labour markets (e.g. Pastorino, 2015) and from papers studying the role of female leadership in determining the gender-specific outcomes of subordinates (e.g. Flabbi et al., 2019). Specifically, Flabbi et al. (2019) propose a signal extraction model where employers have incomplete information about workers' productivity, which is in turn influenced by their gender. They assume that executives are better-equipped to assess the skills of employees of the same gender. This is in line with recent works in the sociolinguistic literature (e.g. Canary et al., 2009 and Scollon et al., 2011) and survey evidence (e.g. Angier and Axelrod, 2014 and Ellison and Mullin, 2014) suggesting the presence of communication frictions between men and women at work. We adopt a similar framework in this paper. However, in our model, we allow for different managerial layers within firms and for the dynamic evolution of the gender composition in each layer. This helps us to endogenise potential trickle-down effects induced by changes at the top of the organisation on gender gaps across the rest of the firm hierarchy.

Second, our work relates to a recent strand of the literature aimed at explaining the internal organisation of firms, the formation of hierarchies within firms as well as the assignment of workers across firm hierarchies (see Caliendo et al., 2015, Garicano and Van Zandt, 2012 and references therein). Yet, this set of papers has so far abstracted from the role that organisational structure has to play in determining gender-specific outcomes within firms. We complement this literature by exploring the role that organisational hierarchies play in propagating gender gaps.

At an empirical level, our work relates to the literature examining the role of female representation on the board of directors on firm-level gender outcomes, exploiting two similar board quota reforms (see Bertrand et al., 2018 for Norway and Maida and Weber, 2019 for Italy).⁵ Both these papers use a similar identification strategy, originally employed in (Ahern and Dittmar, 2012), and find no evidence of spillover effects on the representation of women at the top on women in the rest of the firm.⁶ We contribute to this strand of literature by distinguishing the

⁵ Recent work by Drechsel-Grau et al. (2020) uses personnel data from one of the largest European manufacturing firms and shows that a lower number of female managers increases gender gaps and thus constitutes a structural disadvantage for women.

⁶ A related literature has considered the effect of female leadership on measures of overall firm performance, with mixed evidence. For example, Adams and Ferreira (2009) and Ahern and Dittmar (2012) study the effect of the gender composition of boards on firms' valuation and operating performance,

effect of this policy at different hierarchical layers within the firm. In particular, we conjecture that since corporate board members are more likely to interact with workers at the upper layers of the firm hierarchy, the impact of board quota reform should be more pronounced at these upper layers. Furthermore, not only does the framework of internal firm hierarchies allow us to distinguish the differing effect of the policy on the top and the bottom layers of the firm but it also facilitates predictions on the trickle-down effects of such a reform.

Outline The rest of the article is organised as follows. In Section 2, we present the administrative data used in the paper. In Section 3, we describe the reform. In Section 4, we present our theoretical model. Section 5 introduces our empirical strategy and displays our empirical results. In Section 6, we conclude our analysis and discuss further avenues of research.

2. Data

In this section, we first provide additional details about the data used in this paper. Next, we describe how we classify workers within firms into different hierarchical layers. Finally, we provide some descriptive evidence on the level and the evolution of the gender representation and wage gaps over our sample period.

2.1. Data description

Our data is formed by merging two data sets, the French administrative firm-employee data called DADS Postes and BoardEx.⁷ The DADS Postes is based on mandatory annual reports filed by all firms with at least one employee. Our data therefore includes all private sector French workers except the self-employed between 1999 and 2016. For each worker, the DADS reports the gross and net wages, hours worked, occupation, gender, age and the identity of the firm in which the was employed in. Although the data does not include worker identifiers, it tells us the worker's employment status, wage and occupation title in the previous year if the worker was employed in the same firm. As a result, we are able to observe the entire workforce of a given firm. Each firm in DADS Postes is assigned a unique identifier which allows us to track it over time and to merge it with the BoardEx data.

The BoardEx data provides us valuable information about the composition of the board of directors for firms that have ever been publicly-listed. The BoardEx data runs from 1999 to 2017 and contains information on the gender, age, experience, education and position of the members of the board of directors, as well as their year of entry and exit. Merging the BoardEx data and DADS Postes enables us to study the impact of the share of women in the board of directors on subordinates in the firm.

Sample selection. For our analysis we eliminate firms with less than one full time equivalent employee. We keep all employees in private sector firms with non-zero salaries and hours worked, aged between 26 and 64. We provide additional details about our sample selection in Appendix D.

2.2. Classifying occupations into layers

To study the effect of female board representation across firm hierarchies, we categorise occupations into different layers using our French administrative data. We borrow the concept of a layer from the theory of firm hierarchy proposed by [Garicano \(2000\)](#). According to this theory, a layer is a group of employees who perform a similar set of tasks within the organisation. Following [\(Caliendo et al., 2015\)](#), we assign employees

while [\(Matsa and Miller, 2011\)](#) consider the probability of downsizing during the Great Recession.

⁷ DADS stands for Déclaration Annuelle des Données Sociales in French. DADS Postes is a restricted data set and is administered by the French National Statistical Institute (INSEE).

Table 1
Descriptive statistics.

	Gender Wage Gap		Share of Women	
	1999	2016	1999	2016
Layer 1	1.32	1.22	0.18	0.21
Layer 2	1.19	1.15	0.22	0.36
Layer 3	1.12	1.08	0.33	0.44
Layer 4	1.08	1.07	0.34	0.53

Note: The gender wage gap is calculated as the ratio of the average male to female hourly wage for each of the four layers. The share of women employed is given by the total number of women divided by the total number of people employed for each of the four layers. Both these quantities are calculated yearly for the universe of workers in private firms.

to different managerial layers based on their occupations and organise these layers into a hierarchy. The purpose of this exercise is not to separate employees in a firm according to the functional characteristics of the tasks they perform (e.g. engineers, lawyers, accountants) but rather on the basis of their hierarchical level in the organisation, that is, the number of layers of subordinates that they have below them.

Recent research has shown that classifying the data by layers is economically meaningful.⁸ In our specific context, organising the data in this manner is insightful because we are not only able to investigate the direct effect of the reform at the top of the firm hierarchy, where gender gaps are widest, but also the indirect effects on employees in subordinate layers. Here, we follow [\(Caliendo et al., 2015\)](#) and use the occupation codes observed in our data to categorise workers into four layers.⁹

1. Chief Executive Officers, top management official, firm owners, members of the Board of Directors (occupation codes 21–23)
2. Senior executives and professionals, comprising of senior management and senior technical professionals (occupation codes 31–38)
3. Intermediate professions, middle management and technicians (occupation codes 42–48)
4. Employees such as administrative staff, security workers, sales workers and blue-collar workers (occupation codes 52–56 and 62–69)

2.3. Gender gaps across firm hierarchical layers

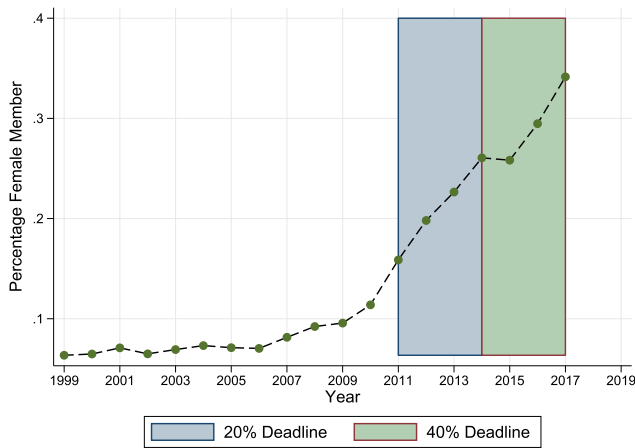
Table 1 documents the gender wage gap across different layers in 1999 and 2016. We calculate the gender wage gap for each hierarchical layer as the ratio of the average male to female hourly wages. The share of women employed in a given layer is calculated as the total number of women divided by the total number of male and female employees.

Gender wage gaps. Table 1 shows that the gender wage gap is higher in top layers of the firm hierarchy, with men earning on average 32 and 19 percent more than women in layers 1 and 2 respectively in 1999, compared to 8 percent more in layer 4 in the same year. While the gender wage gaps have narrowed in all layers in 2016, we still observe that they remain wider in the upper layers.

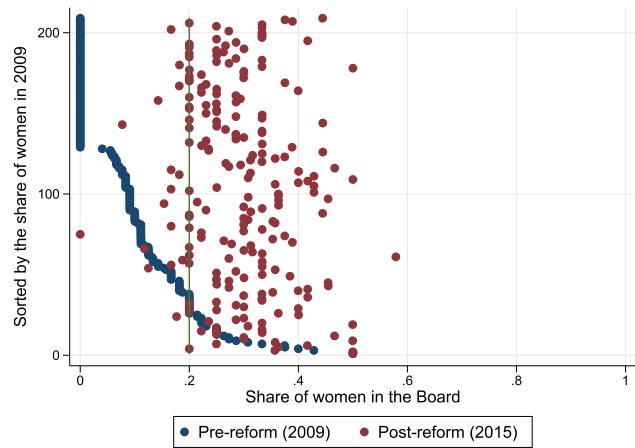
Gender representation gaps. We draw two conclusions from Table 1. First, the share of women is lower in top layers (1 and 2) compared to bottom layers (3 and 4). For instance, in 1999, the share of women in layers 1 and 2 was 18% and 22% respectively, compared to 33% and 34% in layers 3 and 4. Second, there has been a noticeable improvement in female representation in all four layers between 1999

⁸ For more details refer to [\(Caliendo et al., 2015\)](#) and the references therein.

⁹ Occupation codes in the French data are based on the Socio-Professional Categories (PCS). Additional information on these occupational codes can be found at <https://www.insee.fr/fr/information/2497958>.



(a) Mean share of women in the BoD



(b) Effect of reform on the share of women

Fig. 1. Compliance with gender quota reform.

and 2016, with more pronounced changes being observed in layers 2, 3 and 4.

Having discussed the classification of workers into firm hierarchies and examined gender wage and employment gaps across each hierarchical layer, the next section delves into the details of the gender quota reform in France.

3. The gender quota reform in France

On the 20th of January 2010, the French Parliament voted in favour of a law imposing a quota on female board membership for publicly-listed firms. This law was promulgated by the French president a year later on the 27th of January 2011. According to the law, women had to constitute 40 percent of corporate boards before 2017, with an interim deadline of 20 percent before 2014.

The law had an immediate impact on the share of women on corporate boards in publicly-listed firms, as can be seen in Fig. 1. Fig. 1a shows the share of women on the board of directors of firms of the treated sample. The treated sample refers to firms that were publicly-listed in 2010 and remained so for the rest of the reform period. This contrasts with the intent-to-treat sample, which comprises all firms that were publicly-listed in 2010, regardless of whether they de-listed after the reform. For the treated sample of firms, the average share of women on the board of directors hovered just above 10 percent in 2010 and increased sharply from 2011, reaching over 25 percent by 2014 and around 34 percent by 2017.¹⁰

Next, we investigate how many firms in the treated sample met the 20 percent quota prior to and after the first deadline of the reform in 2014. Fig. 1b plots the share of women on the board of directors of each of the publicly-listed firms in our treated sample of firms. Approximately, 50 percent of the publicly-listed firms in our data had no women in their board prior to the reform. The rest of the firms had a share of women between 10 and 50 percent. By 2015, approximately 95 percent of the firms complied with the reform in our sample. Evidence from these figures therefore suggests that the quota policy has succeeded in opening the doors of boardrooms to women.

A key concern is whether firms selectively de-listed after 2010 to avoid the gender quota.¹¹ By regressing the probability of de-listing on

¹⁰ The trend for the intent-to-treat sample is very similar and is available upon request.

¹¹ In the case of Norway, Ahern and Dittmar (2012) found that firms with a lower share of female board members prior to the reform in 2003 were more likely to de-list, perhaps to avoid the gender quota.

the pre-reform share of female board members in Table A.1 in the Appendix, we find that French firms with a lower share of female directors in 2010 were not more likely to de-list. Another concern may be that firms met the quota by assigning female directors to peripheral roles. In addition, newly-hired female board members may be less experienced or may be juggling several board appointments. This may then limit the impact of the reform on gender gaps within firms. As shown in Fig. A.1 in the Appendix, we find that the shares of female CEOs and executive directors increase modestly after the reform. Yet, the share of female directors in HR committees rises substantially, suggesting that these new female members may have substantive roles on the board. In addition, from Fig. A.2, the experience gap between men and women widens only modestly and the number of board appointments of female directors remains fairly stable. As such, just as for the case of Norway, the gender quota did not bring about a large decline in the age and experience of female directors.¹² The above trends are similar for both the intent-to-treat and treated firms. In summary, the evidence suggests that the Copé-Zimmerman law faced strong compliance from publicly-listed firms in France.

In the next section, we present a theoretical model that aims to explain how greater female representation on corporate boards may impact gender wage and representation gaps across the firm hierarchy.

4. Theoretical framework

Using a simple theoretical model, we examine the predicted effects of a rise in female representation at the top of the firm on gender wage and employment gaps along the firm hierarchy. The model considers firms as a hierarchical organisation, where internal and external labour markets co-exist and frictions slow down the ability of firms to adjust to any change in their environment. The framework builds on previous work on gender inequality, internal labour markets and organisational hierarchies. The theoretical framework will serve as the basis for the empirical analysis we carry out in Section 5.

4.1. The firm

We describe a big firm as a hierarchical organisation, consisting of a set of K layers of increasing hierarchical responsibility. Within each layer, there is a continuum of tasks, characterized by level of complexity $\omega > 0$. Each worker within a layer performs a task. We use the terms ‘job’ and ‘task’ interchangeably. A newly-hired worker is assigned a task in

¹² See Bertrand et al. (2018) and Ahern and Dittmar (2012).

a layer, while incumbent workers are promoted when they are assigned a more complex task within a layer. Since movements across layers are rare in our data, we rule out promotions across layers in the model.¹³ We also assume that workers in a given layer k are directly managed by workers in the layer just above, layer $k - 1$. Specifically, workers in layer $k - 1$ are responsible for hiring workers in layer k and decide on promotions across tasks within layer k . Wages are set at the layer-level according to a piece-rate contract θ_k , with tasks of complexity ω paying a wage $w_k = \theta_k \omega$, $\theta_k \in (0, 1)$. The wage of an incumbent worker thus increases as she is promoted to jobs of increasing complexity within a layer.

For simplicity, we assume that the number of jobs in each layer is fixed. Jobs are exogenously destroyed at a rate δ . Whenever a job is exogenously destroyed a new job with type ω_k^n is created and a vacancy is posted on the external labour market. We further assume that contracts are binding and separations occur only if the worker is hit by an exogenous separation shock. Workers performing a job ω receive an internal promotion opportunity for a job $\omega^p = \gamma \omega$ with exogenous arrival rate δ^p where $\gamma > 1$ is the productivity of the new job relative to the current one.¹⁴ Note that as $\gamma > 1$, new jobs of type ω_k^n will be the jobs in layer k with the lowest level of complexity, which we refer to as entry-level jobs.

When a worker, either internal or external, is considered for a job in layer k her application is assigned to a randomly-picked worker in layer $k - 1$, who is then responsible for the hiring decision. We define a worker in charge of a recruitment of an external candidate (or deciding on worker promotions) as the hiring manager. Applicants applying for a job of type ω draw their job-specific skills level from a common known distribution. However, the realisations of workers' job-specific skills are not perfectly observed. A candidate's true job-specific skills, q_i , is given as follows

$$q_i = s_i + \epsilon_i \quad (1a)$$

$$\epsilon \sim \mathcal{N}(0, \sigma^2) \quad (1b)$$

$$s \sim \mathcal{N}(0, \sigma_s^2) \quad (1c)$$

where ϵ_i is a random term independent from s_i . Specifically, s_i and ϵ_i can be interpreted as the observed and unobserved component of a worker i 's job-specific skills, respectively. The variance of s is denoted by σ_s^2 , while that of ϵ is denoted σ^2 . Eq. (1) implies that for a hiring manager observing s_i , the value of a worker's true job-specific skills is distributed as $q|s_i \sim \mathcal{N}(s_i, \sigma^2)$.

When a new worker is hired for or promoted to a job of type ω , the firm pays a one-time productivity penalty, $\omega \chi_k(q)$. We think of this penalty as a retraining cost, where $\chi_k(q)$ that is assumed to be a strictly decreasing, strictly convex function of q . Once the productivity penalty is paid, production starts and per-period profits from the firm-worker match are given by

$$\pi_k(\omega) = (1 - \theta_k)\omega \quad (2)$$

We assume that whenever a worker is promoted the incumbent manager gets a compensation equal to the continuation value of the destroyed match. Accordingly, a manager is indifferent between retaining a worker and allowing her to move to a new unit, implying that the value of a match for the hiring manager does not depend on the internal promotion

probability.¹⁵ Given a subjective discount factor β , the expected value of hiring a worker i for a job ω in layer k is then given by

$$\Pi_k(\omega|s_i) = \frac{1}{1-\beta(1-\delta)}\pi_k(\omega) - \omega E[\chi_k(q)|s_i] \quad (3)$$

where the first term on the right-hand side of Eq. (3) is the continuation value of the match and the second term is the expected retraining cost paid by the firm.¹⁶ The characteristics of a candidate i affect the value of a match via their expected retraining costs. Given our assumption that $\chi_k(q)$ is strictly decreasing in q , $\Pi_k(\omega|s_i)$ is strictly increasing in s_i .¹⁷ Finally, we show in Appendix B that as long as the cost function $\chi_k(q)$ is strictly convex in q , $\Pi_k(\omega|s_i)$ is a decreasing function of σ , the standard deviation of the unobserved component of a worker's productivity.¹⁸

External Hires

External hires occur whenever a job in layer k is exogenously destroyed. A new job is then created in the same layer. A hiring manager is randomly picked among workers in layer $k - 1$ and searches for candidates by posting vacancies on the external labor market. Search is sequential and we assume that candidates are screened at no cost. A hiring manager posting a vacancy on the external market can screen a randomly-drawn candidate from the pool of unemployed workers. After screening a worker i , the hiring manager observes s_i and takes the hiring decision. If the worker is not hired, the hiring manager randomly draws a new candidate from the pool of unemployed workers and the process goes on until the vacancy is filled. The hiring manager therefore hires the worker i if the expected value of hiring, $\Pi_k(\omega_k^n|s_i)$, is bigger than the value of waiting, which we define as $V_k(\omega_k^n)$. Since $\Pi_k(\omega_k^n|s_i)$ is strictly increasing in s_i , the hiring manager hires worker i if and only if $s_i > \bar{s}_k$, where \bar{s}_k is the value of s_i such that $\Pi_k(\omega_k^n|s_i) = V_k(\omega_k^n)$.

Internal Promotions Incumbent workers receive internal promotion opportunities with per-period probability δ^p . Promotions happen within layers and, if approved, result in an internal transfer. If promoted a worker working in a job of type ω is then transferred to a job of type $\omega^p = \gamma \omega$. Whenever a worker in layer k is considered for a promotion, a randomly-chosen worker in layer $k - 1$ is selected to be the hiring manager for the promotion decision. If the promotion is not approved, the worker remains in her current job. When assessing a promotion, the new hiring manager observes the candidate i 's observed component s_i and takes the promotion decision. For simplicity, we do not consider the possibility of the hiring manager learning about the skills of their employees over time.¹⁹ A promotion is approved if the expected value of worker i in the job ω^p is higher than the present discounted value of the future per-period profits in her current job ω

$$\Pi_k(\omega^p|s_i) > \frac{1}{1-\beta(1-\delta)}\pi_k(\omega) \quad (4)$$

¹⁵ For matches that are not destroyed at the beginning of a period, the manager expects to receive the continuation value of the match if the worker is not promoted. If the worker is promoted, we assume the manager receives a transfer equal to the continuation value. This implies that the expected value of a match will not depend on the promotion probability as the manager will receive the same expected value whether the worker remains in the current match or is promoted.

¹⁶ The derivation of Eq. (3) is done in Appendix C.

¹⁷ This is a direct consequence of the fact that the distribution of q conditional on a higher s_i stochastically dominates the distribution of q conditional on lower values of s_i . This implies that the expected retraining cost is strictly decreasing in s_i . With $\omega > 0$ the result on the value of a match follows.

¹⁸ The fact that $\chi_k(q)$ is a cost implies that when it is strictly convex function of q , the value of a match conditional on s_i is strictly concave in q .

¹⁹ Cullen and Perez-Truglia (2018) have presented evidence highlighting the importance of employees' social interactions with their managers for the employees' career progression. Introducing dynamic learning in the model would affect the promotion probability for agents over time. The implications of dynamic learning will then depend on the underlying assumptions concerning the learning rate of male and female managers. If female managers interact more with other female subordinates and learn quickly about their true level of skills then this would lead to higher promotion rates for women over time.

¹³ As shown in Table A.2, most incumbent employees stay in their layer across years. This is unsurprising, since each hierarchical layer is broad and moving across hierarchical layers would involve a large increase in responsibilities and job complexity.

¹⁴ For simplicity we assume that γ is fixed.

and it is not approved otherwise. Similar to the case of an external hire, as $\Pi_k(\omega^p | s_i)$ is strictly increasing in s_i , the hiring manager hires a worker i if and only if $s_i > \bar{s}_k^p$, where \bar{s}_k^p the value of s_i such that $\Pi_k(\omega^p | s_i) = \frac{1}{1-\beta(1-\delta)} \pi_k(\omega)$.²⁰

4.2. Gender and discrimination

We next introduce gender in our model and assume that workers can either be males, denoted as m , or females, denoted as f . To capture how differences in gender composition can generate gender wage and representation gaps at different levels of the organisation, we use a simple model of statistical discrimination. While statistical discrimination can be modeled in several ways, we borrow from (Flabbi et al., 2019) and assume that male managers are less accurate, relative to female managers, in their assessment of female workers' abilities.²¹ Specifically, while male and female managers see the same observable skills s_i of an applicant, their prior on the distribution of the unobserved component ϵ differs. More specifically, let us define $\sigma_{m,f} > \sigma$ as a male manager's prior on the standard deviation of a female's candidate unobserved abilities. We assume that female managers correctly assess the distribution of the unobserved skills of female candidates and their prior over the dispersion of ϵ is thus equal to σ . For simplicity, we do not model differences in the case of male applicants and assume that both female and male managers correctly assess the distribution of ϵ for male candidates.²²

How does having a male rather than a female hiring manager affect the probability that a female candidate is hired? Differences in the perceived dispersion of ϵ among hiring managers imply that for a given value of observed skill-level, s_i , a male hiring manager is more uncertain, relative to a female manager, about the true skill-level, q_i , of a female candidate. Consequently, conditional on a female candidate's observed skills s_i , a male hiring manager's prior over the distribution of the candidate true skill level is such that $q \sim \mathcal{N}(s_i, \sigma_{m,f}^2)$. On the other hand, a female hiring manager will correctly perceive $q \sim \mathcal{N}(s_i, \sigma^2)$.

As discussed in Section 4.1, as long as the profit function is strictly concave in q , an increase in σ reduces the expected profits from a match conditional on a given s_i . Consequently, given that the profit function is strictly increasing in s , a female candidate will need a higher draw of s to be considered profitable by a male hiring manager. Formally, let us use $g \in (m, f)$ to denote the gender of an applicant and $h \in (m, f)$ to denote the gender of a hiring manager. Finally, let us denote $\bar{s}_k(g, h)$ ($\bar{s}_k^p(g, h)$) as the reservation value for a hiring manager of gender h when assessing an external (internal) candidate of gender g . The following conditions are then implied by our model:

1. $\bar{s}_k(f, m) > \bar{s}_k(m, m)$ and $\bar{s}_k^p(f, m) > \bar{s}_k^p(m, m)$
2. $\bar{s}_k(f, m) > \bar{s}_k(f, f)$ and $\bar{s}_k^p(f, m) > \bar{s}_k^p(f, f)$

Condition 1 states that a male manager's reservation value for a female candidate will be *higher* than the reservation value for a male candidate. Condition 2, on the other hand, states that a male manager's reservation value for a female candidate will be *higher* than a female manager's reservation value for a female candidate.

Intuitively, our model implies that if a female candidate with observed skills s_i is assessed by a male manager, her probability of being hired or promoted is lower than that of a male candidate with the same s_i and lower than her chance of being hired or promoted when assessed by a female manager. In addition, the higher the share of men in layer

$k - 1$, the more likely a female candidate for layer k will be assessed by a man instead of a woman. Hence, the model predicts that firms with a lower share of women in $k - 1$ will have larger gender wage and employment gaps in layer k . In the remaining part of this section, we discuss these points more formally.²³

Probability of Being Hired on the External Market. Let us define the share of women and men in layer k as $l_k(f)$ and $l_k(m)$, respectively. By construction, $l_k(f) = 1 - l_k(m)$. The probability that an external candidate i of gender g considered for a newly-created job in layer k is hired is given by

$$P_k(g \text{ candidate hired}) = \sum_{h \in m, f} l_{k-1}(h) P(s_i > \bar{s}_k(g, h)), \quad (5)$$

which is the sum of the probability of being hired conditional on the gender h of the hiring manager, weighted by the share of hiring managers of each gender. Given that $\bar{s}_k(f, m) > \bar{s}_k(f, f)$, a female candidate's probability of being hired will positively depend on the relative share of female workers in layer $k - 1$, denoted $l_{k-1}(f)$. As such, the greater the share of women in layer $k - 1$, the greater the share of women hired in layer k .

Next, let us define the share of external candidates of gender g as $l_0(g)$. The probability that a new job in layer k is filled by an external candidate of gender g is given by

$$P_k(g \text{ new hire}) = \frac{l_0(g) P_k(g \text{ candidate hired})}{l_0(f) P_k(f \text{ candidate hired}) + l_0(m) P_k(m \text{ candidate hired})}, \quad (6)$$

which is equal to the probability that the first external candidate is of gender g , conditional on being hired. Note that $P_k(f \text{ new hire})$ increases in both $P_k(f \text{ candidate hired})$ and $l_0(f)$. As a result, the probability that a newly-created job in layer k is filled by a female external candidate is, other things equal, increasing in both the share of females in layer $k - 1$, $l_{k-1}(f)$, and in the share of females among external candidates, $l_0(f)$.

Probability of Being Promoted. Similarly, the probability that an internal candidate i of gender g being considered for a promotion in layer k is promoted is given by

$$P_k(g \text{ candidate promoted}) = \sum_{h \in m, f} l_{k-1}(h) P(s_i > \bar{s}_k^p(g, h)). \quad (7)$$

Once again, given that $\bar{s}_k^p(f, m) > \bar{s}_k^p(f, f)$, a female candidate's probability of being promoted will positively depend on the relative share of female workers in layer $k - 1$.

Average Wage Change. Wages in our model are defined as a piece rate θ_k of the degree of productivity of a job, $w = \theta_k \omega$. Therefore, in a given layer k , an incumbent worker experiences wage growth only if promoted to more complex tasks. The wage change conditional on a promotion for a worker of gender g earning w in layer k is thus given by

$$\Delta w = \theta_k (\gamma - 1) \omega. \quad (8)$$

Hence, the expected wage change for an incumbent worker of gender g working in a job ω is given by

$$\mathbb{E}_k[\Delta w | \omega, g] = \theta_k (\gamma - 1) \omega \left[\underbrace{\delta^p \sum_{h \in m, f} l_{k-1}(h) P(s > \bar{s}_k^p(g, h))}_{= P_k(g \text{ candidate promoted})} \right], \quad (9)$$

where the term in squared brackets is the probability of promotion for a worker of gender g in layer k , which is given by the probability of

²⁰ Note that as both sides of Eq. (4) are multiplicative in ω , \bar{s}_k^p does not depend on ω .

²¹ Other models of discrimination would produce qualitatively similar results as long as male hiring managers are less likely than female managers to hire a female candidate.

²² Assuming that male managers assess male abilities more accurately than female managers would simply reinforce our model's predictions on the effect of an increase in the share of female managers on gender gaps.

²³ While in our model we assume that the hiring decision is taken by a single hiring manager, we would obtain qualitatively similar conclusions if we instead assumed that the decision was taken by a hiring committee. If the gender composition of the hiring committee reflected that of layer $k - 1$, then raising the share of women in layer $k - 1$ would also have a positive impact on the outcomes of women in layer k . However, the quantitative effect would depend on how individual preferences are aggregated in the hiring committee.

promotion conditional on receiving a promotion opportunity times the probability of receiving a promotion opportunity, δ^p .

4.3. Propagation of gender gaps

One of the main goals of our analysis is to assess how a change in gender representation at the top of an organisation propagates to lower hierarchical layers. Our model allows us to formally inspect these trickle-down effects. In this section we present our model's predictions regarding how an increase in the share of women at the top trickles down within the organisation. We then use our model's testable predictions to guide our empirical analysis.

Gender Representation Gap. Let us use the hat notation to define a firm-level variable in the next period. For big firms the law of motion for the share of workers of gender g in layer k is then given as follows

$$\hat{l}_k(g) = (1 - \delta)l_k(g) + \delta P_k(g \text{ new hire}), \quad (10)$$

where the first term refers to the persistent effect of the gender composition among incumbent workers and the second term captures the effect of the gender of composition of new hires.²⁴ Eq. (10) shows that changes in the composition of hiring managers will only affect $\hat{l}_k(g)$ gradually via their effect on the composition of new hires. Accordingly, the speed at which changes in the hiring process transmit to $\hat{l}_k(g)$ depends on the turnover rate δ . A higher δ and thus a higher share of newly-hired workers in the total workforce imply a quicker adjustment of $\hat{l}_k(g)$ to the gender composition in layer $k - 1$.

Gender Wage Gap. Finally, in a big firm, the law of motion for the average wage by gender g in layer k is given by

$$\begin{aligned} \hat{W}_k(g) = & \underbrace{\frac{l_k(g)}{\hat{l}_k(g)}(1 - \delta)(1 - \delta^p)W_k(g)}_{\text{Incumbents not considered for promotions}} + \underbrace{\frac{1}{\hat{l}_k(g)}\delta P_k(g \text{ new hire})\theta_k\omega_k^g}_{\text{New hires}} \\ & + \underbrace{\frac{l_k(g)}{\hat{l}_k(g)}(1 - \delta)\delta^p[W_k(g) + (\gamma - 1)W_k(g)P_k(g \text{ candidate promoted})]}_{\text{Incumbents considered for promotions}}. \quad (11) \end{aligned}$$

The first term is the average wage among current incumbents of gender g not considered for promotion weighted by their share, $(1 - \delta)(1 - \delta^p)l_k(g)/\hat{l}_k(g)$, among workers of gender g in the next period. The second term is the wage among new hires, $\theta_k\omega_k^g$, weighted by the share of new hires among workers of gender g in the next period. Finally, the third term is the average wage among current incumbent workers of gender g considered for a promotion to more complex tasks weighted by their share, $(1 - \delta)\delta^p$, among workers of gender g in the next period.

Defining the gender wage gap in layer k as

$$\Gamma_k = \frac{W_k(m)}{W_k(f)}, \quad (12)$$

the evolution of the gender wage gap in layer k follows from Eqs. (11) and 12.

Eq. (11) shows how a change in the gender composition of hiring managers in layer $k - 1$ affects $\hat{W}_k(g)$. As discussed earlier, keeping δ and δ^p constant, a higher share of female managers in layer $k - 1$ has two effects. First, it increases the share of female workers among new hires in layer k and second it increases the share of incumbent women promoted to more complex tasks in layer k . Eq. (11) shows that the second effect increases the average wage for women in layer k (last term). Under our assumption that female and male managers are equally capable of assessing male candidates, the change has no effect on the

average wage of men in layer k . As a result, the effect on promotions reduces the gender wage gap in a given layer k .

An increase in the share of females among new hires, on the other hand, has a different effect on \hat{l}_k . By increasing the share of females in entry-level (and thus low paid) jobs, ω^n , an increase in the share of female workers among hiring managers decreases $\hat{W}_k(f)$. Conversely, as female new hires substitute male new hires the change has the opposite effect on the number of males in entry-level jobs, with a positive effect on $\hat{W}_k(m)$. At least initially, the effect on new hires thus increases the gender wage gap. Eq. (11) shows that the effects from both channels are gradual, as $\hat{W}_k(f)$ and $\hat{W}_k(m)$ slowly adjust to their new equilibrium at a speed that is increasing in δ and δ^p . Note also that the negative effect on $\hat{W}_k(f)$ coming from new hires channel slowly loses importance as the share of females in layer k , $l_k(f)$, increases to its new equilibrium (where $l_k(f) = \hat{l}_k(f)$) and newly-hired workers are gradually promoted to higher paying jobs.

Empirical Implications Given the expressions for the gender representation and employment gaps presented in Eqs. (10) and 12, the model has three main empirical implications relevant for our analysis. Namely, a rise in the share of women in layer $k - 1$, $l_{k-1}(f)$, should lead to a

- i) *Change in the share of women in layer k :* An increase in the share of women in layer $k - 1$, $l_{k-1}(f)$, should increase the share of female workers in layer k , $l_k(f)$. As long as the exogenous job turnover rate δ is smaller than 1, this change happens gradually, driven by an increase in the share of women among new hires, while $l_k(f)$ progressively adjusts to its new equilibrium.
- ii) *Change in the gender wage gap in layer k :* From Eq. (11), a greater $l_{k-1}(f)$ affects $\hat{W}_k(f)$ via two channels. First, it gradually raises the average wage among incumbent women in layer k by increasing their gender-specific promotion probability. This first effect gradually increases $\hat{W}_k(f)$ and hence lowers the gender wage gap in layer k . Second, a greater $l_{k-1}(f)$ increases the share of newly-hired women in layer k , who are paid an entry-level wage. This second channel mitigates the effect on promotions and reduces $\hat{W}_k(f)$. The second effect is initially stronger and slowly loses importance as $l_k(f)$ approaches its new equilibrium level, gradually increasing the number of women in layer k and thus reducing the share of newly-hired among the women working in the layer. Depending on the relative size of the two effects, an increase in $l_{k-1}(f)$ can initially either increase or decrease the gender wage gap in layer k , with the latter likely to dominate at least over longer horizons.
- iii) *Trickle-down effect:* An increase in the share of women in layer $k - 1$, $l_{k-1}(f)$, should gradually affect the gender representation and gender wage gap in layer $k + 1$, through its effect on $l_k(f)$. Similarly, the increase in $l_{k-1}(f)$ should gradually affect the gender representation and gender wage gap in layer $k + 2$, through its indirect effect on $l_{k+1}(f)$ via $l_k(f)$. Repeating this argument, the effect of a change in $l_{k-1}(f)$ should gradually trickle-down to lower layers, with a lag that increases with their hierarchical distance from layer $k - 1$. The speed of the trickle-down effects depends on the exogenous job turnover and promotion rates, δ and δ^p .

5. Empirical strategy

Our model predicts that changes at the top of the firm will slowly trickle down the firm's hierarchical ladder. Therefore, we expect an increase in the share of women on corporate boards to directly narrow the gender gaps in upper layers of the firm hierarchy. However, no direct or immediate impact is expected on gender gaps in the lowest layers. Indeed, we expect trickle-down effects on lower layers to be observed after some time lag. In this section, we investigate our model's predictions empirically. We first present our identification strategy, followed by a discussion of the estimation results.

²⁴ The derivation of the equations presented in Section 4.3 is done under the assumption that firms are large enough for us to approximate means with their asymptotic values. In Eqs. (10) and 11, this approximation is used both for means defined at layer level, as well for those defined on the subset of workers being promoted and for those hit by the exogenous separation and promotion shocks.

5.1. Identification strategy

To assess the impact of female corporate board representation on gender gaps on each layer of the firm hierarchy, we estimate the following:

$$y_{kjt} = cons + \alpha FSB_{jt} + \beta' X_{kjt} + \gamma_{kj} + \delta_t + \epsilon_{kjt} \tag{13}$$

where subscript kjt refers to layer k , firm j at time t . y_{kjt} refers to the outcome variable, FSB_{jt} refers to the share of female board members, X_{kjt} refers to time-varying control variables, γ_{kj} refers to unobservable fixed effects and δ_t is a set of year dummies. Lastly, ϵ_{kjt} refers to the error term. The coefficient α measures the impact of FSB_{jt} on y_{kjt} and is therefore our parameter of interest. In the empirical analysis below, we estimate Eq. (13) for each year.

The main challenge in estimating the effect of the share of female board members on any outcome of interest is endogeneity. Specifically, ϵ_{kjt} could potentially be contemporaneously correlated with FSB_{jt} due to omitted variables such as time-varying management style that we fail to capture. This would in turn bias our estimates of α . More precisely, it is likely that

$$E(FSB_{jt} \times \epsilon_{kjt}) \neq 0$$

To address the potential endogeneity bias, we follow (Ahern and Dittmar, 2012) and (Bertrand et al., 2018) in using the pre-quota share of female board members, interacted with year dummies, as instruments for the current share of female board members. The key intuition for the instrument is as follows: the lower the share of female board members prior to the reform, the further treated firms are from the quota and the more they must increase the number of women on their board. Otherwise put, the gender quota law gives us the first-stage in our identification strategy by inducing firms with low shares of female board members to increase their share of women so as to meet the quota. We also assume that the pre-reform share of female board members does not directly affect current gender gaps, apart from their impact on current female board share.

First-stage. The first-stage regression is therefore

$$FSB_{jt} = \rho \times FSB_{j2010} \times \varphi_t + u_{jt} \tag{14}$$

Where FSB_{j2010} refers to the share of women on the board in 2010, prior to the reform, φ_t refers to year dummies and u_{jt} to the error term. For the period between 2011 and 2016, we expect ρ to be negative, since firms with lower pre-quota shares of female corporate board members increase their share of female board members more than firms with higher pre-quota shares. This is confirmed by Table 2, which shows that firms with higher pre-quota shares of female board members experience a significantly lower increase in female board shares relative to 2011. $FSB_{j2010} \times \varphi_t$ therefore appears to be a valid instrument for FSB_{jt} .

5.2. Empirical results

Having discussed the empirical strategy, we now present the regression results. While we adopt the same instrumental variables as (Ahern and Dittmar, 2012) and (Bertrand et al., 2018), our firm-layer level specification in Eq. (13) allows us to differentiate the impacts of an increase in the share of female board members on each layer of firm hierarchy. Following our model's predictions, we expect an increase in the share of women at the corporate board level to primarily narrow gender gaps in the upper layers. No significant impact is immediately expected on gender gaps among the lowest layers. Instead, some time lag is expected before any trickle-down effect is observed for these layers.

In line with our model, the outcomes we consider are i) gender representation gaps and ii) gender wage gaps. Gender wage gaps are measured by the log ratio of average male to female real hourly wages while gender representation gaps are measured by the share of women employed. Consistent with Eq. (13), these outcomes are measured for layer

Table 2
Effect of pre-reform female board share on current female board share (FSB_{jt}).

	(1)	(2)
FSB_{j2010}	0.652*** (0.054)	-
$2012 \times FSB_{j2010}$	-0.108** (0.044)	-0.110** (0.044)
$2013 \times FSB_{j2010}$	-0.223*** (0.053)	-0.232*** (0.053)
$2014 \times FSB_{j2010}$	-0.367*** (0.065)	-0.374*** (0.065)
$2015 \times FSB_{j2010}$	-0.382*** (0.058)	-0.392*** (0.057)
$2016 \times FSB_{j2010}$	-0.443*** (0.070)	-0.449*** (0.069)
Industry-Time effects	Yes	Yes
Firm fixed effects	No	Yes
Observations	1133	1133

Note: *** $p < 0.01$, ** $p < 0.5$, * $p < 0.1$. The base year is 2011. All regressions include industry-year effects. Standard errors are reported in parentheses and are clustered at the firm and year levels. Sample consists of treated firms that were publicly-listed in 2010 and remained so till 2016. Observations are weighted by total employment in the firm.

Table 3
Descriptive statistics - regression sample.

	Gender wage gap		Share of women	
	Mean	St. Deviation	Mean	St. Deviation
Layer 2	0.287	0.424	0.367	0.207
Layer 3	0.039	0.470	0.502	0.291
Layer 4	0.036	0.571	0.560	0.249

Note: The gender wage gap is given by the log ratio of average male to female real hourly wages. The share of women employed is given by the total number of women divided by the total number of people employed.

k , where $k \in \{2, 3, 4\}$, firm j at time t . Descriptive statistics for the outcome variables in our regression sample are given in Table 3.

Table 4 presents the OLS and IV regression results of the impact of female board share on gender wage and representation gaps at the firm-layer level.²⁵ From column 2, an increase in FSB_{jt} significantly narrows the gender wage gap in layer 2. For instance, a rise in female board share by 10 percentage points would lower the ratio of male to female wages by 2.36 percent.²⁶ As for layers 3 and 4, while the coefficients on FSB_{jt} in column 2 are negative, they are statistically insignificant. Concerning the impact of FSB_{jt} on the share of women employed, the IV results presented in column 4 indicate that the effects are only statistically significant for layers 2 and 3, but not for layer 4. Specifically, a rise in female board share by 10 percentage points raises the share of women employed in layers 2 and 3 by 2.39 and 2.51 percentage points respectively. While positive, the coefficient on FSB_{jt} for layer 4 is insignificant.²⁷

Overall, the results from Table 4 support our model's predictions. Since corporate board members are more likely to assess job candidates

²⁵ As expected, the first stage regression results for each firm-layer, given in Table A.3 in the Appendix, are similar to those presented in Table 2.

²⁶ The negative impact of FSB_{jt} on the gender wage gap is due to the average wage of women increasing, rather than that of men declining, as shown in column 1 of Table A.4 in the Appendix.

²⁷ The positive impact of FSB_{jt} on the share of women employed in layers 2 and 3 is not due to a rise in the share of men exiting the firm, as shown in column 2 of Table A.4 in the Appendix.

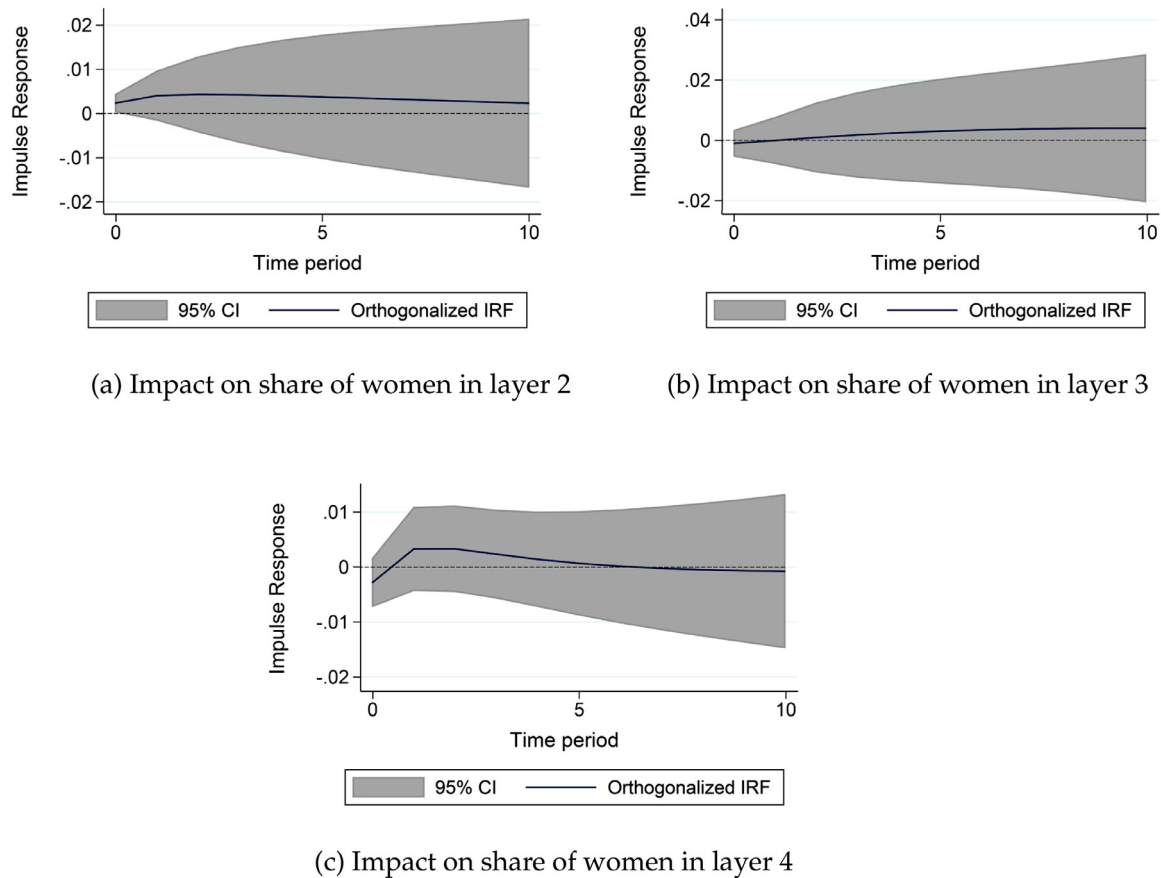


Fig. 2. Per-period effect of a one-time, one standard deviation increase in FSB_{jt} on the share of women employed in lower layers.

Table 4
Effect of female board share on gender wage and representation gaps.

	Gender wage gap		Share of women employed	
	OLS	IV	OLS	IV
<i>Layer 2 (Senior executives & professionals)</i>				
FSB_{jt}	-0.132	-0.236**	0.024	0.239**
	(0.067)	(0.075)	(0.035)	(0.092)
Observations	948	948	1080	1080
<i>Layer 3 (Middle management & technicians)</i>				
FSB_{jt}	-0.393	-0.301	0.021	0.251**
	(0.336)	(0.333)	(0.129)	(0.093)
Observations	699	699	851	851
<i>Layer 4 (Administrative, sales, security & blue-collar workers)</i>				
FSB_{jt}	0.015	-0.091	0.440	0.204
	(0.232)	(0.502)	(0.303)	(0.336)
Observations	797	797	888	888

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. All regressions include industry-year effects and firm fixed effects. Standard errors in parentheses are clustered at the firm and year levels. Gender wage gap regressions also include as controls the average female and male age and age-squared. The gender wage gap is given by the log ratio of average male to female real hourly wages. The share of women employed is given by the total number of women divided by the total number of people employed. Observations are weighted by total employment in the firm-layer. The sample for columns 1 and 2 consists of firm-layers with at least one male and female employee. The sample for columns 3 and 4 includes firm-layers with at least one employee, regardless of gender.

employed and a lower wage gap for these layers. In contrast, gender gaps in the lowest layers are not affected in the short-run, since these layers are not directly managed by those at the corporate board level and are only indirectly affected by trickle-down effects over time.

Lastly, our IV results are valid only if our instrument is plausibly exogenous. So far, we have assumed that firms' pre-reform female board shares in 2010 do not have an impact on current gender gaps, apart from via current female board shares. To support our assumption, we check for different pre-reform trends in gender gaps among firms with varying levels of female board share in 2010. Table A.5 presents the results from regressions of the gender wage and representation gaps on firms' female board shares in 2010 interacted with a time trend for the five-year period prior to the reform. From the coefficients on the interaction term $FSB_{j2010} \times Year$, we find that pre-reform trends in gender gaps for all 3 layers did not vary significantly with firms' female board shares in 2010. This therefore suggests that prior to the reform, firms with higher female board share were not on different linear time trends compared to those with lower female board share.

In short, we find that a greater share of women on corporate boards narrows gender wage and representation gaps at upper layers of the firm and has limited impact on lower layers. These results contrast with those presented in Bertrand et al. (2018) and (Maida and Weber, 2019), who find that corporate gender board quotas had minimal impact on firm-level gender gaps in Norway and Italy. Our framework suggests that the firm-level effect of the gender board quota belies the differing impacts along each layer of firm hierarchy. Indeed, corporate board members may not be involved in the hiring and assessment of all workers in the firm; they are more likely to interact with and manage workers closer to them in the firm hierarchy. As such, our framework adds to these earlier findings, by showing how a greater share of women on corporate boards does significantly narrow gender wage and employment gaps,

for upper layers of the firm hierarchy, a greater share of women on the board implies a more favourable assessment of female candidates for promotion and for hire. This then leads to a greater share of women

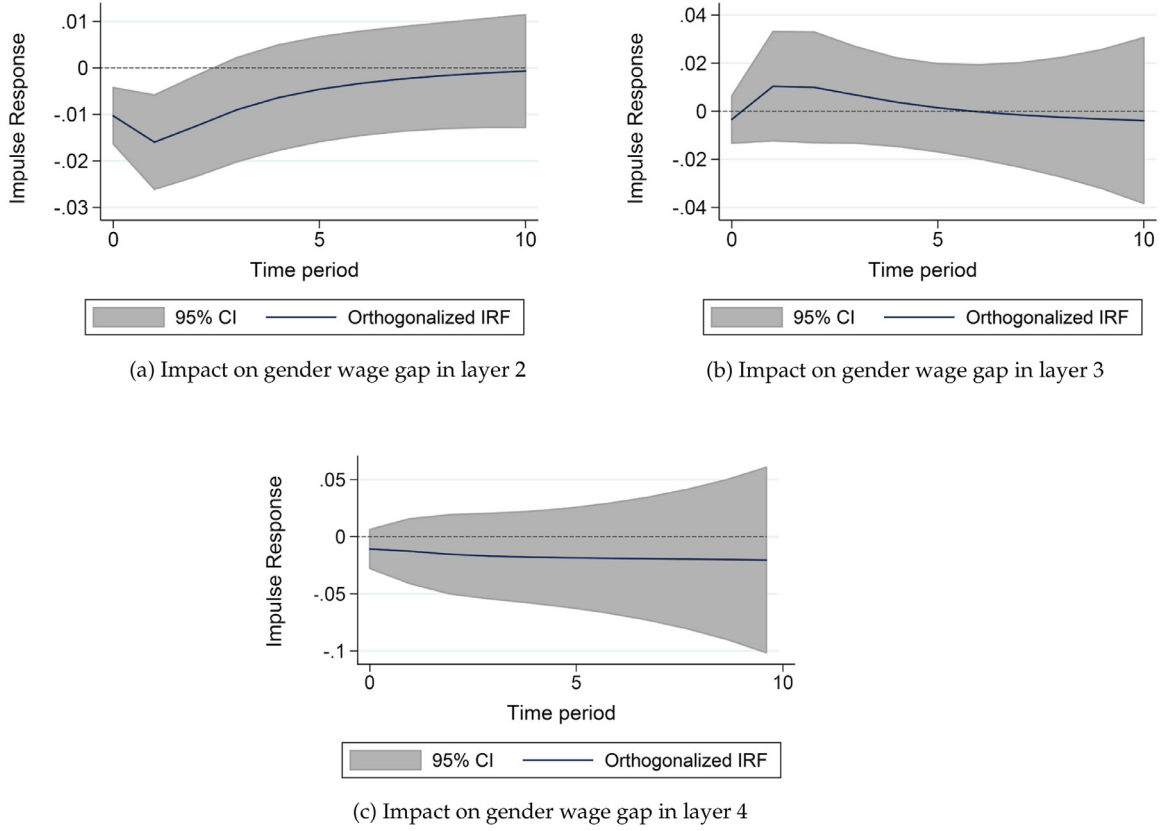


Fig. 3. Per-period effect of a one-time, one standard deviation increase in FSB_{jt} on the gender wage gap in lower layers.

albeit only for the upper layers of the firm hierarchy. This is not inconsequential, given that gender gaps at these layers are on average the highest, as previously shown in Table 1.

5.3. Long-run trickle-down effects of gender composition

Our findings in the Section 5.2 suggest that over the reform period, the increase in the share of women on corporate boards had limited impact on gender gaps in lower layers of the firm hierarchy. From Eqs. (10) and (11) in our model, a rise in the share of women in a given layer $k - 1$ first narrows gender wage and employment gaps in layer k , then layers $k + 1$ and so on, where the extent of trickle-down effects depends on the turnover and promotion rate of jobs. To ascertain whether an increase in female representation at the top of the firm eventually trickles down to lower layers, we go beyond the reform period and exploit the long time dimension of our administrative data. Specifically, we run the following panel VAR (P-VAR) from 1999 to 2016 at the firm level:

$$Y_{jt} = A_0 + A_1 Y_{jt-1} + \zeta_j + CU_{jt} \quad (15)$$

where

$$Y_{jt} = (y_{1jt}, y_{2jt}, y_{3jt}, y_{4jt})^T$$

Here, y_{kjt} refers to the gender gap outcomes in each layer $k \in \{1, 2, 3, 4\}$, in firm j at time t . The terms ζ_j and CU_{jt} refer to firm fixed effects and the error term respectively. In line with our model and empirical exercise in Section 5.2, the gender gap outcomes denoted by y_{kjt} are as follows:

$$y_{1jt} = FSB_{jt}$$

$$y_{kjt} = (WG_{kjt}, FS_{kjt})^T \text{ for } k > 1$$

As before, FSB_{jt} is the share of female board members (layer 1) in firm j at time t . FS_{kjt} and WG_{kjt} refer to the share of women and the log ratio

of average male to female real hourly wage respectively in layer k , firm j at time t , for $k \in \{2, 3, 4\}$.²⁸

We restrict C to be a lower triangular matrix, such that CU_{jt} can be expressed as:

$$u_{1jt} = \epsilon_{1jt}$$

$$u_{2jt} = c_{21}u_{1jt} + \epsilon_{2jt}$$

$$u_{3jt} = c_{31}u_{1jt} + c_{32}u_{2jt} + \epsilon_{3jt}$$

$$u_{4jt} = c_{41}u_{1jt} + c_{42}u_{1jt} + c_{43}u_{2jt} + \epsilon_{4jt}$$

Our assumption that CU_{jt} is a lower triangular matrix hence implies that while shocks to gender gaps in a given layer can contemporaneously affect gender gaps in layers below, the reverse is not true. This reflects our model's assumption that workers are managed by those above them in the firm hierarchy.

To find the optimal lag order, we minimize the Akaike information criterion (AIC) by choosing a lag order of 1. No restriction is imposed on A_1 . To remove firm-level fixed effects ζ_j , we use a forward orthogonal deviation (FOD) transformation. We then estimate A_1 of the transformed model

$$Y_{jt}^* = A_1 Y_{jt-1}^* + CU_{jt}^* \quad (16)$$

where the asterisk denotes the FOD transformation of the variable. Since the P-VAR analysis can only be conducted with firms containing all four layers of firm hierarchy and with non-zero employment of men and women in each layer, we are left with 756 firm-level observations. Using the estimates of A_1 , we then compute the orthogonalised impulse response functions (IRF). The IRFs show the per-period impact of a one-

²⁸ Since the compensation structure for corporate board members is typically more complex than for salaried employees of firms, we do not model the gender wage gap at the corporate board level (layer 1).

off one standard deviation increase in FSB_{jt} on gender gaps across firm layers over ten years.

Fig. 2 displays the impact of a positive shock to the share of female board members on the share of women employed in each firm layer. From the figure, the initial shock to FSB_{jt} increases the share of women in layer 2 significantly in the first period, by around 0.3 percentage points. However, the subsequent per-period effects quickly become insignificant beyond the initial shock. In contrast, no significant impact on layers 3 and 4 is observed over a ten-year period. As such, there is no significant evidence of trickle-down effects on the share of women in these layers, even in the long run.


Likewise, the per-period impact of a one-time, one standard deviation shock to FSB_{jt} on gender wage gaps in each firm layer is given in Fig. 3. From the figure, the shock to FSB_{jt} has a significantly negative impact on the gender wage gap in layer 2, lowering the ratio of male to female wages by approximately 1.6 percent in the first year. In the second year following the shock, the wage gap in layer 2 continues to decline significantly, by around 1 percent. While the per-period effect on layer 2 remains negative beyond the second up to the eighth year following the shock, it ceases to be statistically significant around the third year. In contrast, no significant impact is observed for layers 3 and 4 over a ten-year period. Hence, just as for the share of women, the shock to FSB_{jt} has no significant long-term trickle-down effects on the gender wage gap in these lower layers.

Given the limited trickle-down effects of greater female board representation on gender gaps in lower layers of the firm even in the long-run, we now ask whether an alternative quota reform would have a greater impact on the relative outcomes of women at lower layers of the firm. From our theoretical model, raising the share of middle managers (layer 3) would have a more direct impact on gender gaps in the lowest layers of the firm. As such, we investigate whether a one-time positive shock to the share of women in layer 3, FS_{3jt} , has any impact on gender gaps in layer 4 over a ten-year period.

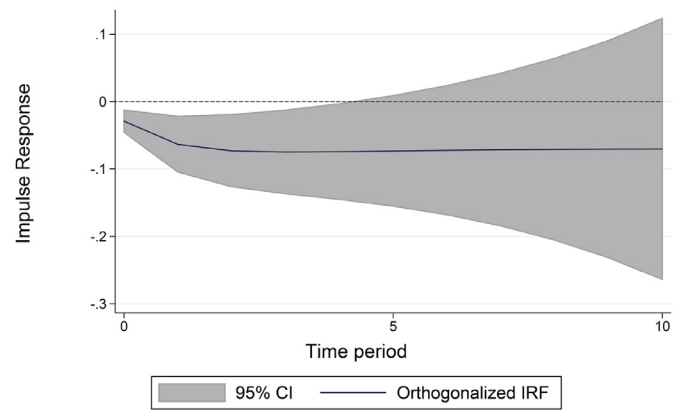
As previously shown in Table 1, while the gender wage gap is narrowest in layer 4, it stands at 7 percent in 2016, which is still substantial. Moreover, given that most women are employed in layer 4, a narrowing of the gender wage gap would impact a wider number of women, thereby lowering the gender wage gap in aggregate. We use the estimates from our P-VAR regression in Eq. (16) to conduct this counterfactual exercise. Fig. 4 presents the IRFs from a one-time, one standard deviation increase in the share of women in layer 3 on gender gaps in layer 4 over a ten-year period. From Fig. 4a, the shock to FS_{3jt} has no significant positive impact on the share of women employed in layer 4. This is perhaps unsurprising, since we have seen from Table 1 that women already made up 34 percent of layer 4 in 1999 and 56 percent in 2016. However, as shown in Fig. 4b, the shock to FS_{3jt} does significantly narrow the gender wage gap by around 6 percent in layer 4 in the first year following the shock. The per-period impacts remain significantly negative till the fourth year. Beyond the fourth year, the per-period effects are negative but insignificant.

Overall, the P-VAR results suggest that a one-off positive shock to female representation on corporate boards has positive long-run effects on gender gaps on the next hierarchical layer of firms, namely layer 2 (professionals and managers). However, the predicted trickle-down effects on layers 3 (intermediate professions and middle management) and 4 (administrative and blue-collar workers) are insignificant even over a ten year period. Our theoretical model suggests that to directly improve the relative outcomes of women in lower layers, raising the share of women in the layer just above would be more effective. Our counterfactual exercise confirms this, as we find that raising the share of women in layer 3 would have a greater impact on the gender wage gap in layer 4.²⁹

²⁹ An alternative policy that would directly impact a broader segment of the firm would instead impose gender-balanced hiring committees for all layers of



(a) Impact on share of women in layer 4



(b) Impact on gender wage gap in layer 4

Fig. 4. Per-period effect of a one-time, one standard deviation increase in FS_{3jt} on the gender gaps in lower layers.

Together, the IV regression and P-VAR results suggest that imposing quotas on female leadership in firms has significant positive effects on female employment and wages in the top firm layers. However, these effects do not necessarily trickle down the firm hierarchy in a significant way. Given that gender gaps are larger in the upper layers of the firm hierarchy, the impact of greater female board membership is not inconsequential. However, other policies beyond corporate board quotas should also be considered if one's goal is to improve labour market outcomes for a broader segment of women.

6. Conclusion

In this paper, we have considered how an increase in top female representation impacts the relative labour market outcomes of women across the firm. To do so, we construct a stylised theoretical model of statistical discrimination with testable empirical implications. In our

the firm. Specifically, if the hiring committees had less than 50% women, imposing a 50-50 gender composition on all hiring committees would increase the likelihood of women being hired and promoted throughout the firm. Note however that with a fixed 50 percent share of women on hiring committees, there would be no further trickle-down effects given that the gender composition of the hiring committee for layer k is now independent of the share of women in layer $k - 1$.

model, workers in a given firm layer manage those in the layer just below. We assume that female managers observe the skill of other women with relatively greater precision than men. Under plausible assumptions on the profit function, the model predicts that an increase in the share of women at the top layer will have a direct impact on gender wage and employment gaps in the second layer, which may trickle down the firm hierarchy over time. The extent of the trickle-down effect will vary, depending on the turnover and promotion rates within the firm.

In order to test our model's predictions, we exploit a recent reform in France seeking to increase female representation in corporate boards of directors using explicit gender quotas. This reform induced a plausibly exogenous variation that allowed us to estimate the impact of an increase in the share of women on corporate boards on gender wage and employment gaps across firm hierarchical layers. As predicted by our model, we find that greater female representation on corporate boards significantly increases female employment and narrows the gender wage gap at the upper layers of the firm, but not at the lower firm layers over the reform period. A panel VAR analysis suggests that a positive shock to the share of women on corporate boards would have limited trickle-down effects to the lowest firm layer, even over a ten-year period.

As such, while increasing female representation at the corporate board level does improve the labour market outcomes of women relative to men in upper layers of the firm, where gender gaps are widest, other policies should be explored if the goal is to narrow gender gaps for a broader segment of the labour market. This is particularly relevant, given that a series of similar corporate board gender quotas have been imposed across several countries in Europe as well as the state of California in the US with the aim of promoting gender equality in the workforce. Indeed, our counterfactual P-VAR exercise suggests that increasing the share of women in middle management would do more to significantly shrink gender wage gap in the lowest layer of the firm.

Lastly, our paper presents a number of unexplored avenues for future research. First, the interactions between workers in different layers of the firm are likely more complex than that presented in our model. For instance, there may be a threshold female share that has to be attained before greater impact on subordinate layers can be observed. Second, workers of different genders within the same hierarchical layer may also interact differently, which would also have ramifications on gender gaps in the firm. Finally, it may be interesting to account for how managers learn about their workers' true abilities over time. Exploring these complexities within the firm is important in guiding future policies to promote gender equality at the workplace and would be a promising next step for the literature.

Appendix A. Additional Empirical Results

Table A.1
Probability of firms de-listing.

	Probability of de-listing	
	OLS	Probit
FSB _{J2010}	0.108 (0.241)	0.449 (0.921)
Observations	241	232

Note: ****p* < 0.01, ***p* < 0.5, **p* < 0.1. All regressions above are at the firm level and include industry fixed effects. Probability of exit refers to the probability of firms previously listed in 2010 that de-listed during the period 2011 to 2016.

Table A.2
Layer transitions among incumbent workers.

Layer at <i>t</i>	Layer at <i>t</i> -1				Total
	Layer 1	Layer 2	Layer 3	Layer 4	
Layer 1	2426	455	30	27	2938
Layer 2	385	715,266	45,495	19,573	780,719
Layer 3	18	30,611	348,452	28,018	407,099
Layer 4	59	20,578	29,319	245,443	295,399
Total	2888	766,910	423,296	293,061	1,486,155

Table A.3
Effect of pre-reform female board share on current female board share (FSB_{*it*}) by firm layer.

	Layer 2	Layer 3	Layer 4
2012 × FSB _{J2010}	-0.434*** (0.092)	-0.513*** (0.109)	-0.515*** (0.082)
2013 × FSB _{J2010}	-0.476*** (0.080)	-0.501*** (0.089)	-0.489*** (0.073)
2014 × FSB _{J2010}	-0.372*** (0.069)	-0.362*** (0.063)	-0.439*** (0.077)
2015 × FSB _{J2010}	-0.312*** (0.084)	-0.391*** (0.071)	-0.379*** (0.084)
2016 × FSB _{J2010}	-0.449*** (0.078)	-0.526*** (0.082)	-0.482*** (0.112)
Observations	948	699	797

Note: ****p* < 0.01, ***p* < 0.5, **p* < 0.1. All regressions include industry-year effects and firm fixed effects. Standard errors in parentheses are clustered at the firm and year levels. The base year is 2011. Observations are weighted by total employment in the firm-layer.

Table A.4
Effect of female board share on male outcomes.

	Log male wage	Share of male exits
	<i>Layer 2 (Senior executives & professionals)</i>	
FSB _{<i>it</i>}	0.002 (0.177)	0.445 (2.875)
Observations	948	926
	<i>Layer 3 (Middle management & technicians)</i>	
FSB _{<i>it</i>}	1.285 (1.265)	-2.579 (1.495)
Observations	699	744
	<i>Layer 4 (Administrative, sales, security & blue-collar workers)</i>	
FSB _{<i>it</i>}	1.075 (0.619)	1.895* (0.922)
Observations	797	809


Note: ****p* < 0.01, ***p* < 0.5, **p* < 0.1. All regressions include industry-year effects and firm fixed effects. Standard errors in parentheses are clustered at the firm and year levels. Log male wage refers to the log of average real hourly wage of men. Share of male exits is given by the number of male exits divided by the total number of employee exits between *t* - 1 and *t*. Sample for column 1 consists of firms with at least one male and one female employee. Sample for column 2 consists of firms with at least one worker exit between *t* - 1 and *t*. Observations are weighted by total employment in the firm-layer.



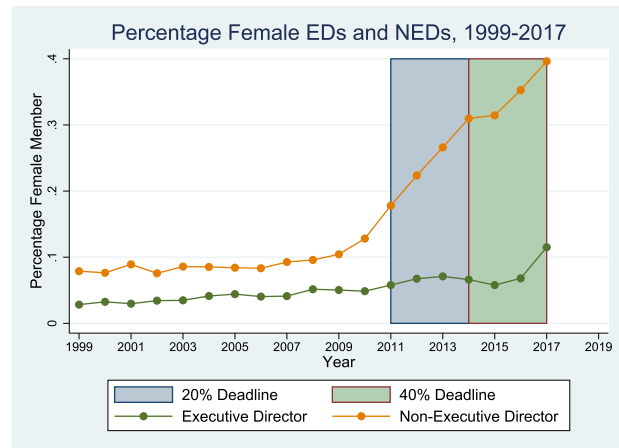
(a) CEO (Intent-to-Treat)



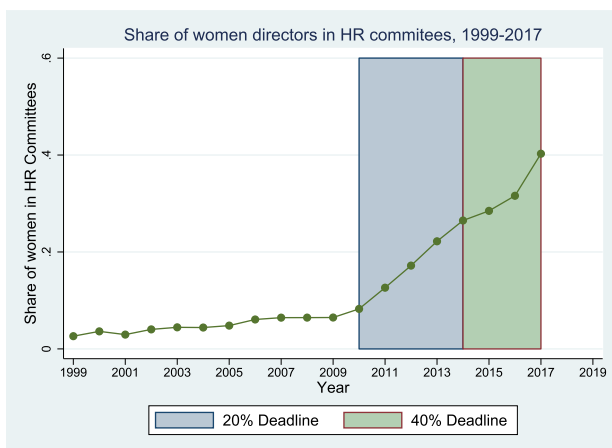
(b) CEO (Treated)



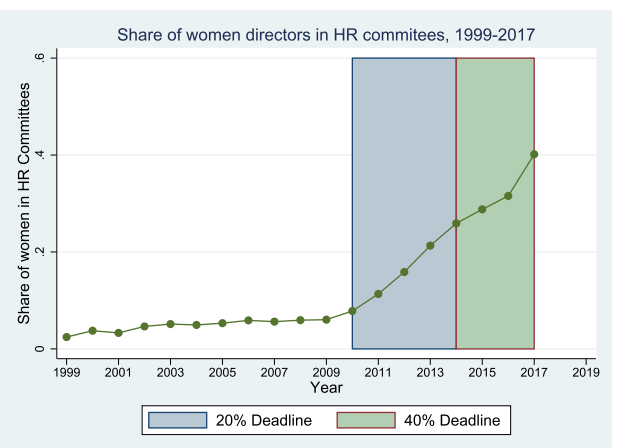
(c) Executives or not (Intent-to-Treat)



(d) Executives or not (Treated)

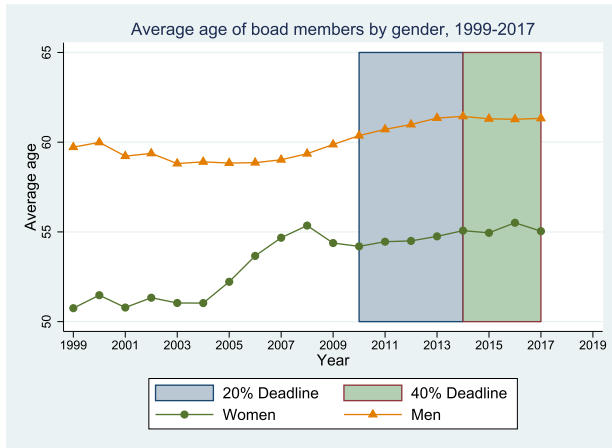


(e) HR committee (Intent-to-Treat)




(f) HR committee (Treated)


Fig. A.1. Average share of women in various BoD roles.




(a) Age (Intent-to-Treat)




(b) Age (Treated)




(c) No. of current boards (Intent-to-Treat)



(d) No. of current boards (Treated)

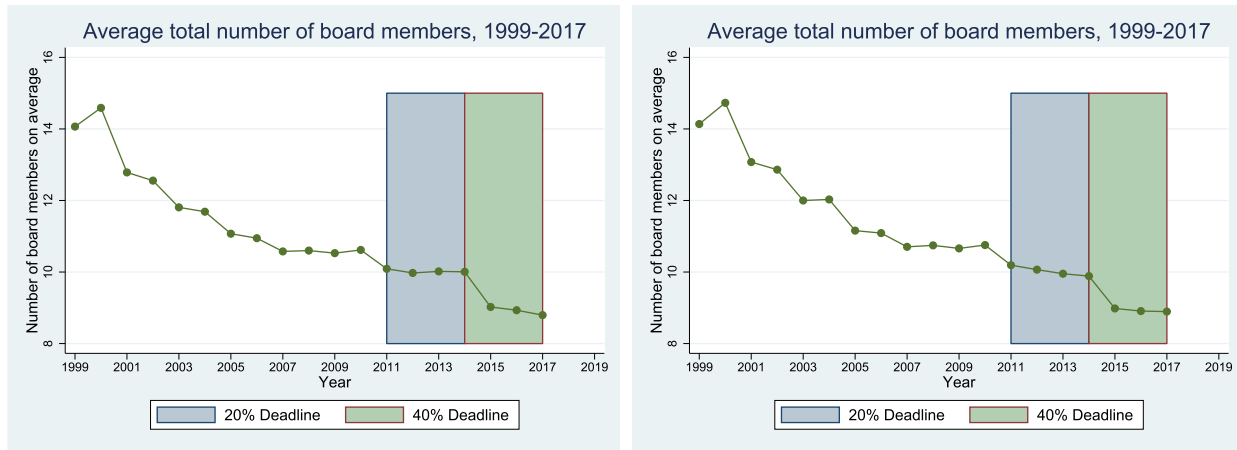


(e) No. of years on boards (Intent-to-Treat)



(f) No. of years on boards (Treated)

Fig. A.2. No. of boards, age and experience of BoD members by gender.



(a) Intent-to-Treat

(b) Treated

Fig. A.3. Average size of boards.

Table A.5
Evidence of pre-reform trends (2005–2010).

	Gender wage gap	Share of women employed
	<i>Layer 2 (Senior executives & professionals)</i>	
FSB ₂₀₁₀ × year	0.015 (0.060)	-0.028 (0.030)
Observations	686	791
	<i>Layer 3 (Middle management & technicians)</i>	
FSB ₂₀₁₀ × year	0.024 (0.051)	0.159 (0.111)
Observations	549	656
	<i>Layer 4 (Administrative, sales, security & blue-collar workers)</i>	
FSB ₂₀₁₀ × year	0.262* (0.104)	-0.047 (0.038)
Observations	515	622

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. All regressions include industry-year effects and firm fixed effects. Standard errors in parentheses are clustered at the firm and year levels. The gender wage gap is given by the log ratio of average male to female real hourly wages. The share of women employed is given by the total number of women divided by the total number of people employed. Observations are weighted by total employment in the firm-layer. The sample for column 1 consists of firm-layers with at least one male and female employee. The sample for column 2 includes firm-layers with at least one employee, regardless of gender.

Appendix B. Value of a Match and Variance of Unobserved Skills

To see how the expected retraining cost depends on the perceived dispersion of unobserved skills let us define $x_i = \frac{\epsilon_i}{\sigma}$, the standardized transformation of ϵ_i where σ denotes the standard deviation of ϵ_i . Accordingly, x_i is distributed as a standard normal. Using the fact that $q_i = s_i + \epsilon_i$, we can thus rewrite the expected retraining cost conditional on a given level of observed job-specific skills as $E[\chi_k(s_i + \sigma x_i)|s_i]$. Taking the derivative of the value of a match with respect to σ conditional on observing s_i thus gives

$$\frac{\partial \Pi(\omega|s_i)}{\partial \sigma} = -\omega \frac{\partial E[\chi_k(s_i + \sigma x_i)|s_i]}{\partial \sigma} = -\omega \frac{\partial}{\partial \sigma} \int \chi[s_i + \sigma x] \phi(x|s_i) dx \quad (17)$$

where ϕ is the standard normal pdf. Using Leibniz's rule one obtains

$$\frac{\partial \Pi(\omega|s_i)}{\partial \sigma} = -\omega \int \chi'(s_i + \sigma x) x \phi(x|s_i) dx \quad (18)$$

using the properties of the standard normal distribution and applying integration by parts one obtains

$$\frac{\partial \Pi(\omega|s_i)}{\partial \sigma} = -\omega \sigma \int \chi''(s_i + \sigma x) \phi(x|s_i) dx = -\omega \sigma E[\chi''(s_i + \sigma x)|s_i] \quad (19)$$

By the definition of a strictly convex function, $\chi''(q) > 0$ and thus $E[\chi''(s_i + \sigma x)|s_i] > 0$ as long as $\chi(q)$ is strictly convex. Accordingly, the value of a match is a decreasing function of σ as long as $\omega > 0$.

Appendix C. Derivations

To derive Eq. (3) in the main text, let's start by considering the expression for the value of a match once the productivity penalty is paid, which we define as $\tilde{\Pi}_k(\omega)$. Let us also define the value of a match of type ω in layer k from the perspective of the current manager conditional on the worker being internally promoted to a job of higher complexity as $V_k(\omega)$. From perspective of the current manager, the continuation value of the match is given by the following expression

$$\tilde{\Pi}_k(\omega) = \pi_k(\omega) + \beta(1 - \delta) \left[\delta^p (\Theta V_k(\omega) + (1 - \Theta) \tilde{\Pi}_k(\omega)) + (1 - \delta^p) \tilde{\Pi}_k(\omega) \right] \quad (20)$$

where Θ denotes the probability that the internal candidate is promoted if she is considered for a promotion. Eq. (20) says that the continuation value of the match consists of the per-period profits today and expected discounted value of the match tomorrow.

In our model we assume that the manger receives a transfer equal to the continuation value of the match when the worker is promoted to a job of higher complexity. In other words, we assume that $V_k(\omega) = \tilde{\Pi}_k(\omega)$. Under this assumption, the term in square brackets in Eq. (20) simplifies to $\tilde{\Pi}_k(\omega)$ and the equation can be rewritten as

$$\tilde{\Pi}_k(\omega) = \pi_k(\omega) + \beta(1 - \delta) \tilde{\Pi}_k(\omega) \quad (21)$$

Solving for the value of the match, (21) gives the following expression

$$\tilde{\Pi}_k(\omega) = \frac{1}{1 - \beta(1 - \delta)} \pi_k(\omega) \quad (22)$$

Finally, subtracting the one-time expected retraining cost paid by the firm conditional on observing s_i , $\omega E[\chi_k(q_i)|s_i]$, from $\tilde{\Pi}_k(\omega)$ we get the expression for the value of a match of type ω conditional on s_i , $\Pi_k(\omega|s_i)$, as reported in Eq. (3).

Appendix D. Sample Selection

Our samples are constructed as follows:

1. From DADS Postes, we keep only private sector firms. We then drop workers that do not have a full-time contract and that do not have a reported occupation code. Also, we remove workers that report zero hours worked and/or zero net or gross yearly salary, as well as

those aged below 26 or above 64. We then aggregate the data to the firm-layer level for each year. Firms that have less than one full-time equivalent employee are also eliminated.

2. Next, using the firm identifier (SIREN), we merge the aggregated and trimmed DADS Postes from Step 1 with the BoardEx data. As previously mentioned, the Board-Ex data includes firms that have ever been publicly listed. Merging the two data sets gives us the sample of French firms that have ever been publicly listed between 1999 and 2016.
3. To conduct the IV regressions in Section 5.2, we keep the sample of firms from Step 2 that are publicly-listed in 2010 and remain so till 2016.
4. Concerning the P-VAR regressions in Section 5.3, we keep all firms from Step 2 with all four layers of firm hierarchy.

References

Adams, R., Ferreira, D., 2009. Women in the boardroom and their impact on governance and performance. *J. Financ. Econ.* 94, 291–309.

Adams, R.B., Funk, P., 2012. Beyond the glass ceiling: does gender matter? *Manage. Sci.* 58 (2), 219–235.

Ahern, K., Dittmar, A., 2012. The changing of the boards: the impact on firm valuation of mandated female board representation. *Q. J. Econ.* 127, 137–197.

Angier, M., Axelrod, B., 2014. Realizing the power of talented women. *McKinsey Q.* 3 (107–115).

Bagues, M., Sylos-Labini, M., Zinovyeva, N., 2017. Does the gender composition of scientific committees matter? *Am. Econ. Rev.* 107 (4), 1207–1238.

Beaman, L., Chattopadhyay, R., Duflo, E., Pande, R., Topalova, P., 2009. Powerful women: female leadership and gender bias. *Q. J. Econ.* 124 (4), 1497–1540.

Becker, S., 1971. *The Economics of Discrimination*. Univ. Chicago Press.

Bertrand, M., Black, S., Jensen, S., Lleras-Muney, A., 2018. Breaking the glass ceiling? The effect of board quotas on female labor market outcomes in Norway. *Rev. Econ. Stud.* 1–48.

Bhide, A., 2019. Do Female Executives Reduce Gender Gaps?. Stanford Mimeo.

Blau, F., Kahn, L., 2016. *The Gender Wage Gap: Extent, Trends, Explanations*.

Caliendo, L., Monte, F., Rossi-Hansberg, E., 2015. The anatomy of french production hierarchies. *J. Pol. Economy* 123 (4), 809–852.

Canary, D.J., Dindia, K., Canary, D.J., 2009. *Sex Differences and Similarities in Communication*. Routledge.

Cardoso, A., Winter-Ebmer, R., 2010. Female-led firms and gender wage policies. *ILR Rev.* 64 (1), 143–163.

Cornell, B., Welch, I., 1996. Culture, information, and screening discrimination. *J. Pol. Economy* 104 (3), 542–571.

Cullen, F., Perez-Truglia, R., 2018. How Much Does Your Boss Make? The Effects of Salary Comparisons. Technical Report. National Bureau of Economic Research.

Deschamps, P., 2018. Gender Quotas in Hiring Committees: A Boon or a Bane for women? LIEPP Working Paper.

Drechsel-Grau, M., Holub, F., et al., 2020. Gender Gaps and the Role of Bosses. Technical Report. University of Bonn and University of Mannheim, Germany.

Ellison, S.F., Mullin, W.P., 2014. Diversity, social goods provision, and performance in the firm. *J. Econ. Manage. Strategy* 23 (2), 465–481.

Flabbi, L., Macis, M., Moro, A., Schivardi, F., 2019. Do female executives make a difference? The impact of female leadership on gender gaps and firm performance. *Econ. J.* 129 (622), 2390–2423.

Gagliarducci, S., Paserman, D., 2012. Gender interactions within hierarchies: evidence from the political arena. *Rev. Econ. Stud.* 79 (3), 1021–1052.

Garicano, L., 2000. Hierarchies and the organization of knowledge in production. *J. Pol. Economy* 108 (5), 874–904.

Garicano, L., Van Zandt, T., 2012. Hierarchies and the division of labor. In: *The Handbook of Organizational Economics*. Princeton University Press, pp. 604–654. chapter 15

Goldin, C., Kerr, S., Olivetti, C., Barth, E., 2017. The expanding gender earnings gap: evidence from the LEHD 2000 census. *Am. Econ. Rev.* 107 (5), 110–114.

Kunze, A., Miller, A.R., 2017. Women helping women? Evidence from private sector data on workplace hierarchies. *Rev. Econ. Stat.* 99 (5), 769–775.

Maida, A., Weber, A., 2019. Female Leadership and Gender Gap within Firms: Evidence from an Italian Board Reform. CEPR Discussion Paper No. DP13476.

Matsa, D., Miller, A., 2011. Chipping away at the glass ceiling: gender spillovers in corporate leadership. *Am. Econ. Rev.* 101 (3), 635–639.

Mulligan, C., Rubinstein, Y., 2008. Selection, investment, and women's relative wages over time. *Q. J. Econ.* 123 (3), 1061–1110.

Olivetti, C., 2006. Changes in women's hours of market work: the role of returns to experience. *Rev. Econ. Dyn.* 9, 557–587.

Pastorino, E., 2015. Job matching within and across firms. *Int. Econ. Rev.* 56 (2), 647–671.

Scollon, R., Scollon, S.W., Jones, R.H., 2011. *Intercultural Communication: A Discourse Approach*. John Wiley & Sons.