ESSAYS IN LABOR ECONOMICS AND CORPORATE GOVERNANCE

VALERIA FERRARO

A dissertation

submitted to the Faculty of

the department of Economics

in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Boston College
Morrissey College of Arts and Sciences
Graduate School

May 2022

Essays in Labor Economics and Corporate Governance

Valeria Ferraro

Advised by Ph.D. Claudia Olivetti, Ph.D. Arthur Lewbel, Ph.D. Hanno Foerster

Abstract

The goal of this dissertation is to understand the absence of women from executive and high-earning positions, with a special focus on the corporate environment. In the first chapter, I analyze the role of news media towards explaining why women in top executive roles in the United States face more unstable appointments relative to their male counterparts. To improve female representation at the top of the firm, several European countries mandated gender quotas on corporate boards. In the second chapter, I analyze a board gender quota mandated on Italian listed companies and its effects on the composition of the board and firm performance. Family responsibilities are among the most important factors that prevent women from reaching high-earning positions. In the third chapter, I broaden the scope of my investigation to high-skill women in the United States, and provide explanations

for the very large increase in childcare hours spent on young children by high-skill mothers of the recent generations.

The first chapter, "Media Focus, Executive Turnover, and Female Leadership", analyzes how the tendency of news media to focus on negative events affects executive turnover in publicly listed firms in the U.S., and to what extent negative media focus explains the relatively higher incidence of turnover for women in top executive roles. Negative media focus implies that news reporting decisions can produce downward-biased public beliefs on firm performance. From the standpoint of a rational board, pessimistic public beliefs on firm performance may affect the expected benefit of retaining a CEO, and in turn, turnover decisions. Linking CEO positions to firm-level news, I provide evidence that the negative focus is higher when a company is led by a woman or an outsider CEO. Counterfactual simulations from a model of executive turnover with event-dependent media focus show that the higher negative focus explains around 15% of the differential turnover rate in female-led firms, even when women are as effective at managing the firm as their male counterparts.

In the second chapter, "Do Board Gender Quotas Matter? Selection, Perfor-

mance, and Stock Market Effects", co-authored with Giulia Ferrari, Paola Profeta, and Chiara Pronzato, I analyze the effects of a gender quota mandated on corporate boards of Italian listed companies in 2013. Exploiting staggered board elections, we find that quotas are associated with a new selection of board members – characterized by higher education and lower age – and no significant costs, neither on firm performance nor on the stock market.

In the third chapter, "Revisiting the Childcare Gap Between High- and Low-Skill mothers", I show that information diffusion on the importance of early child development has been growing fast starting from the mid-1990s. At the same time, childcare hours have increased, especially for mothers of very young children and the high-educated. I argue that information diffusion on the importance of early investments coupled with increasing income inequality plays an important role towards rationalizing some of the trends in childcare time and the widening of the education gradient in childcare hours at different ages of the child's lifecycle.

Contents

List of Tables

List of Tables

List of Tables

List of Figures

List of Figures

List of Figures

Chapter 1

Media Focus, Executive Turnover, and Female Leadership

1.1 Introduction

Women are underrepresented in leadership positions. As of 2017, women held 19.9% of board seats in Fortune 500 companies and covered 5.8% of CEO positions in the same companies. Even when they do make it to the top, women tend to face more unstable appointments (Gayle et al., 2012; Gupta et al., 2020; Keller et al., 2020). In public companies in the U.S., women in executive positions leave the

firm or are replaced at a higher rate than similar men, even when the company is thriving (Keller et al., 2020; Gupta et al., 2020). The general lack of diversity in corporate leadership may come with costs for society as a whole. If talent is equally distributed across groups in the population, narrowing the talent pool may lead to lower economic growth (Hsieh et al., 2019).

This paper asks whether editorial decisions made by news media affect executive replacement in U.S. companies, and to what extent such phenomenon explains the higher incidence of turnover for women in executive roles. My paper focuses on the particular case of Chief Executive Officers (CEOs). There are reasons to expect news coverage to matter for how CEOs are retained and replaced. In the short run, news coverage can affect stock price performance. In the long run, it can affect consumer demand and the firm's ability to attract and retain workers (Graf-Vlachy et al., 2020). Previous research shows that the appointment of a female CEO receives much larger media attention than the appointment of a male CEO (Gaughan and Smith, 2016). Can such higher degree of public attention play any role towards explaining the relative instability of female appointments in highly visible positions?

A company's CEO is at the top of the corporate hierarchy and the main responsible of firm performance. Previous literature finds that female CEOs are more likely

to be replaced than their male counterparts, but the difference is not explained by firm performance (Gupta et al., 2020). In the literature, the link between firm performance and managerial turnover is well established. Prolonged periods of low performance provide a signal of low managerial ability, leading the board to replace the company's CEO (Weisbach, 1988; Gibbons and Murphy, 1990; Murphy, 1999; Taylor, 2010; Jenter and Lewellen, 2019). In my paper, news media act as information intermediaries that monitor firm performance and disseminate information to the public (Nimark and Pitschner, 2019). While boards learn CEO ability over time by relying on actual firm performance, news coverage can influence public beliefs on CEO ability. If public beliefs matter to the board, differences in turnover can arise from differences in news coverage.

I show that news media tend to focus on *negative* performance events. In the long run, such systematic negative focus can produce downward-biased public beliefs on the ability of a CEO. If the performance of a particular CEO – a female CEO – is more newsworthy, a negative news on a female CEO will be more likely to be covered. Over time, public beliefs on the ability of a female CEO can become more pessimistic relative to the case of a less-covered CEO, even for the same realized performance. To offset the bias in public beliefs, the board may require female CEOs to be of higher ability than their male counterparts, and higher turnover rates can arise for

equally qualified female CEOs.

Making quantitative statements on the importance of systematic coverage decisions for executive replacement is challenging. News coverage tends to increase in times of worse firm performance. At any given point in time, turnover decisions depend on past realized firm performance and past news coverage, and the two are not independent from one another. Given such complications, I impose structure on the relationship between firm performance, news coverage, and turnover, and complement the use of reduced form methods with a structural approach.

I start by showing empirical patterns of news coverage, and evidence on the importance of news for CEO replacement. I link firm-level financial news from RavenPack News Analytics to public firms in Compustat and executive positions in BoardEx over the period 2000-2017. I aggregate news data at the news event level, and construct a measure of coverage for each news event. A news event corresponds to a particular happening reported in news media. Coverage for a news event is measured as the number of articles it generates. I document coverage patterns at the news event level, a strategy motivated by the nature of firm-level media coverage. News events are often disclosed by the company itself. Documenting media coverage at the news event level mitigates the risk of confounding media choices with firms' choices. The strategy is also motivated by the online nature of

firm-level news: online news move fast, and for the same news event there is little variation in content across outlets (Cagé et al., 2019).

Two key facts on firm-level media coverage turn out to be relevant. First, a news event with negative sentiment receives more coverage relative to a positive event. The result points to the existence of a negative selection bias in news media, a finding in line with empirical regularities of media behavior (Harrington, 1989; Soroka, 2012; Harcup and O'Neill, 2017). Second, the severity of the negative selection bias is related to the identity of the firm's leadership. A news event with negative sentiment receives 33% more coverage when a firm is led by a female CEO. The differential is not explained by firm performance, nor by systematic differences across firms. Previous research finds that the appointment of a female CEO attracts more media attention relative to the appointment of an average CEO (Gaughan and Smith, 2016). I provide evidence that female-led firms receive more coverage throughout the entirety of the CEO appointment. And because news coverage is mostly negative, the difference is driven by negative events. I argue that such differential treatment in the media is related to women's outsider status in the executive labor market.

Executives are a highly homogeneous group, and entry barriers in the occupation are high (Terviö, 2009). Given their minority status, female CEOs may be

perceived as challenging the status quo, and information on their performance may be more valuable to investors. I investigate such possibility by constructing measures that correlate with the outsider status of a CEO. Such measures include, for example, CEOs at their first appointment, CEOs at the beginning of their tenure, and founders. Similarly to women, negative performance events for CEOs with outsider status are more likely to be covered. However, outsider status does not fully account for the female differential, and female outsiders receive more negative coverage than male outsiders.

A key empirical question is whether news coverage matters to the board when making the turnover decision. When evaluating CEO performance, the board is likely to rely on private information not available in news media. However, to the extent that negative news coverage may come with reputational costs, the board may be reluctant to retain an executive who performs poorly in news media. To the best of my knowledge, only Farrell and Whidbee (2002) shows empirically that negative performance news published in the Wall Street Journal are related with the incidence of CEO turnover. I complement the evidence in Farrell and Whidbee (2002) by using news data from a wide range of news outlets, and by looking at the effects of both positive and negative coverage. My empirical strategy exploits variation in the timing of news releases across firms. I find that the number of negative

news articles released in a quarter is predictive of an appointment ending the following quarter. The marginal effect of a negative article is sizable, and corresponds to 2.3% relative to the sample mean. Positive news, instead, seem to have no effect.

To better isolate the effect of news releases from other confounders and provide a more transparent test of anticipation effects, I complement the previous strategy with an event study approach. I define a sharp news release shock as a quarter in which the firm experiences a number of negative (positive) performance articles greater than the 95th percentile in the firm-specific distribution. The turnover probability jumps discontinuously at the time of a negative news shock, with no clear pattern of following a positive news shock. The evidence suggests that boards respond to news coverage when making the turnover decision, and especially negative news, supporting the hypothesis that reputational concerns may be important.

Having established empirical patterns of news coverage and on the importance of news for CEO replacement, I formalize the editorial role of news media in a model of executive turnover. The model provides a tractable framework to link firm performance, media focus, and firms' replacement decisions. The model's tractability makes it possible to take it to the data in a parsimonious way, and quantify the effects of news coverage decisions.

The model features a forward-looking board taking the turnover decision in every

period. The key difference with traditional models of managerial turnover – such as Taylor (2010) – is that both the private beliefs of the board and public beliefs matter for the turnover decision. Whereas the board's beliefs rely solely on firm performance, public beliefs are updated using published performance. The mapping between realized and published performance is provided by news media (Nimark, 2014).

In line with the empirical evidence, news media are more likely to publish negative performance information. Over time, the private and public information sets can diverge, and the public can become more pessimistic relative to the board. For the same firm performance, more pessimistic public beliefs lead the board to set a higher ability threshold for the CEO to be retained. The model can explain higher replacement probabilities in female-led firms, which cannot be rationalized by differences in actual firm performance. A clear prediction of the model is that female CEOs will be more positively selected around the time of replacement. The prediction finds support in the data.

I derive two additional testable implications that validate the model. Both implications follow from public learning on the ability of a CEO over time. I show that, conditional on CEO tenure, the turnover probability increases the longer the history of negative coverage for a CEO. Conditional on CEO tenure, a long history

of negative coverage also increases the marginal effect of a negative news on the replacement probability.

I then calibrate the model's parameters to quantify how much the bias in news selection contributes towards explaining firms' replacement decisions, and the higher incidence of turnover for female appointments. I set the baseline parameters of the turnover model to match data moments for the sample of male CEOs using simulated method of moments. The key moments pinning down the baseline parameters of the model are the turnover hazard and firm profitability over the first 15 years of CEO tenure.

Using the calibrated parameters, I run two quantitative exercises. First, I remove the negative selection bias, thus making public beliefs aligned with the board's beliefs. Removing the negative selection bias decreases the incidence of turnover by 9.7%. Second, I feed in the higher negative bias estimated for women in the first part of the paper. Holding everything else constant, negative media focus accounts for around 15% of the average turnover differential measured for female appointments relative to male appointments.

My results show that media focus can represent an important source of distortion for the advancement of women and other minorities in leadership roles. The result can have important policy implications. In Europe and more recently in the

U.S., gender quotas have been proposed as a policy to improve female representation in top corporate positions. My results suggest that quotas likely need to be implemented for a long period of time for their beneficial effects to fully arise. As a critical mass of women is reached, the negative focus may become less meaningful, suggesting that in the long run policies mandating equal representation can effectively overcome differential treatment in news media.

Related Literature My paper is at the intersection of several literatures. First, it relates to a large body of research on executive turnover. Several papers analyze executive turnover in relation to firm performance (Weisbach, 1988; Gibbons and Murphy, 1990; Murphy, 1999; Taylor, 2010; Jenter and Lewellen, 2019). Weisbach (1988), Laux (2008), and Kaplan and Minton (2012) study the role of board independence. Fich and White (2003) analyzes board interlocks. Goyal and Park (2002) and Taylor (2010) study the role of CEO entrenchment. In particular, Taylor (2010) is the only paper estimating a structural model of executive turnover, as I also do in my paper.

I contribute to the vast literature on executive turnover by documenting the link between firm performance and news coverage, and formalizing how news coverage affects replacement decisions. The closest paper to mine is Farrell and Whid-

bee (2002), which studies the effect of earnings news reported in the Wall Street Journal on CEO replacement. In line with my results, Farrell and Whidbee (2002) shows that negative earnings news are predictive of CEO replacement. Differently from Farrell and Whidbee (2002), I document novel evidence on coverage patterns in the financial press, and structurally estimate a model of executive turnover to quantify the consequences of such patterns.

Second, my work relates to the literature on the glass ceiling and the barriers to career advancement that women face in top positions. In the executive labor market, gender differences in pay or career advancement have been widely documented. Bertrand and Hallock (2001), for example, show that the gender gap in executive compensation is due to the higher chance of women to be employed in smaller firms and cover lower-ranked positions. Albanesi, Olivetti, and Prados (2015) find that the compensation of female executives is more exposed to declines in firm value and less sensitive to increases in firm value than the compensation of similar males. Recent work shows that women in corporate executive roles exit the occupation at higher rates than men (Gayle, Golan, and Miller, 2012), are more likely to be fired (Gupta et al., 2020), and leave the company at higher rates than comparable men (Keller et al., 2020). The evidence in Gupta et al. (2020) shows that, for similar levels of firm performance, women CEOs are more likely to be dis-

missed. I show that media focus can contribute towards rationalizing the turnover gap: news selection decisions can generate higher turnover rates for female CEOs even when they are as effective at managing the firm as their male counterparts.

My paper is the first one to study the role of media focus to explain the low representation of women in leadership positions. Anedoctally, the idea that women in executive roles may attract public scrutiny is known.[1] The same idea has also been proposed in the corporate finance literature, where the role of media attention is usually analyzed in connection with executive appointment. Gaughan and Smith (2016), for example, shows that the announcement of the appointment of a female CEO receives more than three times higher media attention than the announcement of a male CEO. Lee and James (2007) shows evidence that news articles about the appointment of a female CEO are more likely to emphasize gender.

Finally, my paper relates to research on media bias. In political economy, media bias or refers to the choice to publish biased or inaccurate information (Gentzkow and Shapiro, 2006; Petrova, 2008; Besley and Prat, 2006). In my paper, I refer to a selection bias in the choice of which information to cover. I build on the notion of news selection first introduced by Nimark and Pitschner (2019). The main idea behind modeling news selection decisions is to acknowledge that only a subset of events is considered newsworthy in news media. The notion of news selection

[1] See for example Financial Review, November 2017.

has been applied to partial information models in the macroeconomics literature to show how information reported in news media can shape agents' expectations and drive business cycles (Nimark, 2014; Chahrour et al., 2019). In my model, I use a similar information structure as in Nimark (2014) and Chahrour et al. (2019). Similarly to Nimark (2014) and Chahrour et al. (2019), the results in my paper show that editorial decisions made by news media can represent an important source of distortion for agents' decisions.

The rest of the paper is structured as follows. Section 1.2 introduces the data and the main estimating sample. Section 1.3 briefly describes the institutional environment in which firm-level coverage operates. Section 1.4 documents the key motivating facts. Section 1.5 introduces the model, and section 1.6 presents the model calibration and counterfactuals. Section 1.7 concludes.

1.2 Data

1.2.1 Datasets and sample selection

CEOs BoardEx provides detailed data on executives in large companies around the world, including demographic characteristics, education, employment history, board interlocks, and network data. I select CEO positions in publicly listed US

companies that started between 2000 and 2017. I exclude from the sample CEOs that cover dual positions or are also the company's President or CFO, and I exclude CEO Emeritus positions.

Companies I link CEOs to firm-level data using Compustat and CRSP, and obtain quarterly performance measures and stock price data. I use the firm-level files of BoardEx to obtain characteristics of the board and the the firm's management.

News News data are obtained from RavenPack News Analytics, a database that uses machine learning tools to organize unstructured content from news articles into structured data. RavenPack is a private company that tracks news released by both press and web sources all around the world. The database is used by private investors and across a broad range of academic research on the effects of media on financial markets.[2] The sources tracked by RavenPack include The Wall Street Journal, Dow Jones Newswires, Barron's, MarketWatch, as well as a large number of industry and business publishers, national and local news, and blog sites (roughly 19,000 sources). Most of the firm-level news included in the analysis are full articles (60% of the sample). The remainder of the sample is represented by news flashes, namely news articles composed of a headline and no body text. Relative

[2]For more information on RavenPack, see https://www.ravenpack.com/

to other media databases – such as Factiva – RavenPack does not allow the user to directly access the content of an article, and every news entry is associated with variables containing structured information provided by the algorithm. Moreover, observations in RavenPack are at the entity-level, so that there may be multiple entries for the same news article, depending on the numbers of entities involved in a news story. Although it provides the user with less flexibility than Factiva, RavenPack is particularly suitable for studying the effects of informational flows rather than specific events, and is often used to analyze media sentiment around specific entities or events.

Every news observation in the dataset is categorized by an "event taxonomy", which allows understanding the broad content of an article, and an entity tag, which allows identifying the main entities involved in a news story. I only match news that are very strongly related to the entity mentioned, that is, I only match news in which the "relevance score" of a given entity is equal to 100.[3] For every entity-news entry, the database also provides a "sentiment score", which allows determining the sentiment content of the news article from the point of view of the entity mentioned. The score is derived from a collection of surveys in which financial experts rate entity-specific events as conveying either positive or negative

[3]For any news story that mentions an entity, the data provide a relevance score that indicates how strongly related the entity is to the underlying news story. A score of 0 means the entity was passively mentioned while a score of 100 means the entity was prominent in the news story.

sentiment, and to what extent. The analysts' ratings are then included in an algorithm that generates a score ranging from 0 to 100, where 50 indicates neutral sentiment, values above 50 indicate positive sentiment, and values below 50 indicate negative sentiment. I define a news as either "positive" or "negative" based on the sentiment score distribution. In the remainder of the paper, negative news will usually have sentiment score below the 10th percentile of the distribution, whereas positive news will have sentiment score above the 90th percentile.

News event and coverage for a news event I make an important distinction in news data. A *news event* is equivalent to a news story, and represents a particular happening for a firm at a given point in time. A distinctive feature of news data is that the same news event can be reported by multiple articles. *Coverage for a news event* is defined as the number of articles reporting the same news event. I link articles reporting the same news event using RavenPack's "novelty" and a "similarity" scores, which allow determining how new or similar news articles are, and grouping news articles by similarity. Section 1.3 discusses the motivation behind such choice in greater detail.

1.2.2 Sample description

I match news data to BoardEx using unique ISIN identifiers. Out of 3,126 positions in BoardEx, I am able to match 3,026 positions, 129 of which are covered by women.[4] Only 18% of these are CEO positions in very large companies (such as S&P 500, S&P MID CAP, and S&P SMALL CAP).[5] Larger firms are more likely to appear in news media: as shown by Table A.3.1 in the Appendix, the firms that I am able to match to news data are on average larger than unmatched firms, and have larger boards. Table 1.1 shows the average characteristics of CEOs in the news-CEOs matched sample, separately by gender. Men and women CEOs are a homogeneous group in terms of observable characteristics such as age and education. Women tend to be appointed after longer tenures in the company, and their appointments are shorter on average, although these differences are not statistically significant. Women who make it to the top of the corporate ladder also have larger networks, with a difference of 156 connections on average. More significant differences appear when comparing firms that appoint male and female CEOs. Out of 2,043 companies in the sample, only 105 ever appoint a woman CEO; the sample size decreases even more if I focus on companies in which there is variation in

[4] 20 positions are unmatched due to missing appointment dates in BoardEx, and 80 are unmatched due to missing news data.

[5] The rest of the positions are covered in public companies belonging to the Russell 3000 Index. Women are underrepresented in S&P 500, S&P MID CAP, and S&P SMALL CAP firms (12% of female positions, versus 18% for the full sample of CEOs).

gender across appointments – only 53 companies. Consistent with previous literature, Table 1.1 shows that women tend to become CEOs in smaller firms, and are more likely to be appointed in firms operating in the consumer and service sector (Bertrand and Hallock, 2001; Gayle et al., 2012).

Figure 1.1 shows the sentiment score distribution for the sample of matched news events, with the two vertical bars representing the 10th and 90th percentiles of the distribution. Over 30% of the news events reported have neutral sentiment, but there is substantial variation across news events. Table 1.2 presents descriptive statistics for the same data. I present news events by broad topic, separately by event sentiment. The most common news reported in the media include performance news and analysts' ratings. Negative news events also involve legal and regulatory issues, whereas the top positive events are represented by the release of new products and services. There is substantial variation in the number of articles across positive and negative events, with much smaller variation when looking at the number of days over which an event is reported: the overwhelming majority of news are "short lived" and are reported in the media for one day at most.

1.3 Firm-level media coverage

Media outlets monitor firm performance and deliver information that can be easily accessed by shareholders, investors, and the general public. Although it is possible for investors to monitor performance directly, for example through the company's website or social media, most of the public will rely on processed information available in the news.

The U.S. Securities and Exchange Commission (SEC) requires firms to release quarterly earnings information, which is then made publicly available in the SEC database. Starting from October 23, 2000, the U.S. Securities and Exchange Commission also requires all publicly traded companies to disclose "material nonpublic information" to all investors at the same time. The measure intends to regulate the practice of disclosing information to a select groups of investors and prevent insider trading. Examples of "material information" include: (1) earnings information; (2) mergers and acquisitions; (3) new products or discoveries; (4) changes in control or in management; (5) change in auditors; (6) events regarding the company's securities; and (7) bankruptcies or receiverships. To disclose information, firms can file a Form 8-K with the SEC, distribute a press release through a news or wire service, or through any other non-exclusionary method of disclosure, such as a public conference. Starting from April 2013, companies have the possibility to use social media

platforms such as Facebook or Twitter.

Relative to general interest news, such as those reported by the *New York Times* or the *Washington Post*, firm-level news are often disclosed by the company itself, are verifiable, and arew published online. Once a company releases information, newswires can edit the content of the press release and disseminate it. Appendix Figure A.2.3 shows the number of press releases and news articles released for General Motors over the quarter January-March 2009. Press releases, full articles, and news flashes closely follow each other. However, conditional on a company reporting an event, there is large variation in the choice of news media over which events to emphasize. Online news tend to move fast and be short-lived, thus creating strong incentives for media outlets copy content from those moving first (Cagé, Hervé, and Viaud, 2019). Appendix Figure A.2.3 shows that most of the news articles released for a news event tend to be made available within the first 5-10 minutes following the first breaking news. The speed of diffusion of online news suggests that there is little scope for editing the content of an article relative to the first mover (Cagé, Hervé, and Viaud, 2019). Relative to more traditional means, such as newspapers, the online format also makes space constraints less binding for media outlets. Firms' obligation to disclose and the online format of news motivate the choice to document coverage patterns at the news event level. Given that

an event was disclosed, how much coverage does it generate in news media? The following section provides an answer to such question and documents empirical regularities in firm-level coverage.

1.4 Motivating facts

1.4.1 Empirical regularities in firm-level news coverage

When monitoring states of the world, news media make decisions as to which events are more newsworthy. Editorial decisions made by the media can be thought of as a *selection function* that maps states of the world to newsworthy events. The notion of news selection function was first introduced by Nimark and Pitschner (2019). The idea behind news selection functions is to provide a flexible way to model editorial decisions, without imposing structure on the mechanisms that drive those decisions (Chahrour et al., 2019). The journalism literature has identified empirical regularities on the features of news selection functions. In this section, I document empirical regularities of media behavior in firm-level coverage, and establish a number of key facts that will be useful in the reminder of the paper.

Unusual events are more covered In general, news media tend to emphasize extreme events rather than mundane events (Shoemaker and Vos, 2009). Such feature of media reporting has already been explored in the macroeconomics literature: Nimark (2014), for example, shows how extreme events published in the media can shape agents' expectations and drive business cycles.[6] Even in my sample of public companies, a few very rare events receive much more coverage than more commonplace events. Examples of rare events include, for example, antitrust investigations, insider trading stories, and product recalls. In Figure 1.2, I plot the raw average number of articles generated by each event, with the number of occurrences for an event reported on top of each bar. Out of more than 6 million events in the dataset, antitrust investigations, insider trading, and product recalls are extremely rare occurrences, and receive up to six times the coverage received by more commonplace news such as equity operations, technical analyses, and stock price events.[7]

Negative events are more covered Another empirical regularity in news reporting is that negative events are more likely to be covered in the media (Har-

[6] Nimark (2014) developed a terminology for public signals provided by the media. A story on a dog biting a man would not be published, whereas a man biting a dog would most likely be published. Therefore, he labels public signals provided by the media as "man-bite-dog" signals.

[7] I use the full sample of matched company-news data to have a more representative idea of the universe of published news events in the period 2000-2017.

rington, 1989; Soroka, 2012; Harcup and O'Neill, 2017). For example, Harrington (1989) documents that network television news overemphasize bad economic news. Similarly, Soroka (2012) documents that the *New York Times* is more likely to report bad news about unemployment, inflation, and interest rates rather than good news about the same variables. In Figure 1.3, I show the number of articles for events of different quartiles of the sentiment distribution. Events at the bottom receive 12% more coverage than events at the top, after adjusting for firm performance, firm fixed effects, time fixed effects, and event category fixed effects.

1.4.2 News coverage and leadership

Having established general patterns of media coverage, I turn to understanding how the characteristics of the firm and the leadership correlate with media focus. I document that female CEOs receive more media coverage on average. When looking at news coverage by event sentiment, the difference is fully driven by events at the bottom of the sentiment distribution. Higher coverage being driven by negative events is consistent with negative events being inherently more newsworthy. In the next subsection, I investigate possible reasons behind such finding, and argue that it may be related to the outsider status of women in the executive labor market.

In order to show that a negative news event receives more media coverage when a company is led by a female CEO, I estimate the following equation:

$$\text{Coverage}_e = \alpha + \text{Sent}_e\,\gamma + \text{CEO char}_{f(e)t(e)}\,\delta + \text{Perf}_{f(e)t(e)}\,\eta + \phi_{c(e)} + \phi_{f(e)} + \tau_{t(e)} + \nu_e \quad (1.1)$$

where Coverage_e is the number of articles for event e, Sent_e is the sentiment score of the event, $\text{CEO char}_{f(e)t(e)}$ is a vector of characteristics of the CEO in firm f at time t, and $\text{Perf}_{f(e)t(e)}$ is a vector of performance measures, such as sales and assets. $\phi_{c(e)}$ is an event category fixed effect, $\phi_{f(e)}$ is a firm fixed effect, and $\tau_{t(e)}$ is a time fixed effect. I estimate Equation 1.1 for the full sample of events and separately for events with different sentiment. The results are shown in Table 1.3. First, the coefficient on the sentiment score in Table 1.3 is large and highly significant for news events at the tails of the sentiment distribution, and insignificant in the middle. This is consistent with the first fact documented in the previous section: extreme events are more likely to be covered in the media. Moreover, the size of the coefficient on the sentiment score is almost double in absolute value for very negative news events (i.e. in the bottom 10%) relative to very positive events (i.e. in the top 90%), consistently with negative news being more likely to be covered – the second fact in the previous section. When looking at the characteristics of the CEO, female CEOs receive more media coverage on average, but only for negative events. For the full

sample of events, a news event generates 0.5 additional articles for a female-headed firm; the difference is sizable, and corresponds to 24% relative to the sample mean. When looking at the results by event sentiment, the difference in average news coverage is entirely driven by negative news events. The coefficient on the female indicator is plotted in Figure 1.4. The coefficient is economically and statistically significant at the bottom of the sentiment distribution, with a gap between 30% and 37% relative to the sample mean.

It is possible that women CEOs get appointed in times of worse firm performance, which in turn would result in worse media coverage. Controlling for quarterly sales and assets mitigates this concern, but further robustness checks are presented in the next section.

Why is media coverage higher for female-led firms?

Female-led firms are more monitored by news media. But because negative events are more likely to be covered, the difference is larger for negative events. In this section, I investigate possible reasons why women are more covered by the media, especially for negative outcomes. As it will become clear, the explanations I propose are all related to the diversity status of women in the executive labor market, and thus may all co-exist at the same time.

Executives are a highly homogeneous group, and entry barriers in the occupation are high (Terviö, 2009). Women are still a minority in top leadership roles: as of 2017, female CEOs held only 4.8% of positions in Fortune 500 companies. Female CEOs may be perceived as outsiders and challenging the "status quo", and investors and shareholders may demand more information on their performance. To investigate this possibility, I construct measures that should capture outsider status and check if they fully absorb the female differential, and how they correlate with media coverage. I construct an indicator for CEOs at their first appointment, CEOs at the beginning of their tenure, namely in the first year of their appointment, and founders. The idea is that these CEOs should be less likely to be part of the known pool of CEOs and less connected to the appointing firm, thus attracting investors' interest. In Tables 1.4 and 1.5 I show how outsider status correlates with media coverage for negative and positive events. Similarly to women, CEOs with outsider status are more likely to be covered. The difference is driven by negative performance events. In Figure 1.5, I plot news coverage by event sentiment for outsider CEOs. Although the differences are not as high as those observed for female CEOs, the pattern is similar: for negative events, the difference in news coverage is sizable, while it is almost absent for more positive events. Information may be more valuable when the firm is led by an outsider CEO, especially for negative events. A

key issue to address is whether such high demand for information under a female appointment – and that of outsiders – is associated with higher firm-level uncertainty, which in turn would affect firm performance. The data do not reveal any difference in firm-level uncertainty after the appointment of a female CEO relative to a male CEO, neither when looking at stock prices nor when considering analysts' expectations. I address this point in detail in Appendix A.4.

Another interpretation is that higher media coverage for women could be the result of persistency in coverage decisions following appointment. CEO appointment is a crucial event for a firm, and usually attracts high media coverage, which may persist for some time following the appointment event. The fact that CEOs with low tenures are usually more covered may be related to the fact that lower tenures are closer to appointment. I verify in my data that the appointment of a female CEO generate more coverage relative to the appointment of other CEOs. The difference is 20% (p-value: 0.058) after controlling for firm size and removing outliers (that is, appointment events with coverage above the 99th percentile). Similarly, CEOs at their first appointment in a company generate more coverage: relative to other CEOs, the difference is 23% (p-value: 0.023) after controlling for firm size and removing outliers.

When documenting differences in coverage between male- and female-headed

firms, the key challenge is understanding whether such differential treatment is driven by systematic differences in firm performance. I perform a number of robustness tests to corroborate the fact that my results are not explained by differences in firm performance or heterogeneity across firms. First, I check whether there are any significant differences in firm performance between male- and female-headed firms. If firms were more likely to appoint women in difficult times (the "glass cliff" hypothesis), then negative news for female-headed firms would be worse news, and thus would be more likely to be covered in the media. I plot the distribution of sales and stock prices separately for male- and female-headed firms in Figure A.2.4, and run OLS and quantile regressions in Table A.3.4. While there seems to be virtually no difference in stock price returns between male- and female-led firms, the distribution of sales looks less dispersed for women. If anything, results from the quantile regressions show a slightly positive difference in sales for female-headed firms relative to other firms.[8]

Even if there seem to be little or no differences in observable firm performance for firms that appoint a female CEO, one might still be concerned that the results on news coverage reflect unobservable circumstances that coincide with the appointment of female CEOs. This could be the case, for example, if the company was

[8]These differences are small, given that the 25th, 50th, and 75th percentiles of the log-sale distribution for male-led firms correspond to 3.08, 4.6, and 6.08.

undergoing a change in firm strategy and the board wanted to signal the change by appointing a woman. If this was the case, then the results in the previous section would just be reflecting a spurious correlation due to company circumstances that may have relatively little to do with the characteristics of the leadership. In order to check whether this is the case, I match to CEOs news articles that specifically mention the CEO as the main individual involved in a news story. I extend my main sample to include lower-ranked executives, and link news stories that mention the CFO, COO, and other lower-ranked executives. Again, I only match news articles an executive has maximum relevance. The results are presented in Appendix Tables A.3.2 and A.3.3. I standardize the dependent variable into z-scores to make the results comparable across executives. The results show that female executives are more likely to attract media attention relative to male executives in the same position, with the largest relative effects observed for CEOs. The results assure, at least partially, against the concern that changing company characteristics fully drive the results, and support the conjecture that women in top leadership positions may attract more media interest *per se*.

1.4.3 News coverage and turnover

In this section, I document whether media exposure has any effect on executive turnover. The board is likely to rely on private information not available in news media. Therefore, performance in news media may carry little weight towards the assessment of the quality of CEO performance. However, to the extent that firm reputation affects stock price performance, consumers' demand, and the ability of the firm to attract workers, the board may be reluctant to retain an executive who performs poorly in the news (Graf-Vlachy et al., 2020).

Empirical strategy

To document whether media exposure has any effect on CEO turnover, I exploit variation in the precise timing of news releases across firms. In order to have a homogenous sample of news events, I focus on performance-related news events only, and exclude events related to acquisition and mergers, legal and labor issues, and products and services. I also exclude all performance events related to bankruptcy. I estimate the equation:

$$P(\text{End of CEO app.}_{it}) = \alpha + \delta_1 \text{Negative articles}_{i,t-1} + \delta_2 \text{Positive articles}_{i,t-1} + \theta \text{Number of articles}_{i,t-1} + g(x_{it}) + \text{Perf.}_{i,t-1}\gamma + Z_i'\eta + \phi_{f(it)} + \tau_t + \nu_{it} \quad (1.2)$$

where the dependent variable, $P(\text{End of CEO app.}_{it})$, is the probability that the appointment of CEO i ends in quarter t. Such event can be CEO turnover or move to another firm, as I explain below. I regress the turnover indicator on a number of lagged variables, including the number of negative and positive articles, the total number of articles released, and CEO and firm characteristics. In particular, $\textit{Negative articles}_{i,t-1}$ represents the number of news articles with sentiment below the 10th percentile of the sentiment score distribution released in quarter $t-1$ for CEO i, and $\textit{Positive articles}_{i,t-1}$ represents the number of news articles with sentiment above the 90th percentile released in quarter $t-1$ for CEO i. $g(x_{it})$ is a second order polynomial in tenure, and $\text{Perf.}_{i,t-1}$ is firm performance at time $t-1$, namely quarterly ROA. [9] [10] Z_i is a vector of time-invariant characteristics at the CEO level, including a female indicator, network size, year of appointment fixed effects, and a quadratic function of age at the time of appointment. Finally, $\phi_{f(it)}$ and τ_t represent firm and time fixed effects. I cluster standard errors at the firm level.

My empirical strategy corresponds to a difference-in-differences specification with continuous treatment. However, the timing of news releases is not random. Although companies have an obligation to disclose information as soon as it becomes available, high-ability CEOs will still be able to manipulate the timing of news dif-

[9] The results are unchanged if I control for tenure non-parametrically, or if I allow tenure to have a differential effect by gender.

[10] The results are unchanged if I control for quarterly sales or stock price returns instead of ROA.

fusion so that negative news on the company are released in more favorable times. Under such scenario, my estimates would be downward biased. Moreover, news reports rarely come as a shock, and firms may anticipate the effect of media coverage. Such scenario would correspond to a violation of the parallel trends assumption, and would also bias to my estimates downward. At the same time, the intensity of firm-level coverage may correlate with firm-specific shocks unobservable to the econometrician. If this was the case, I would be confounding the effect of news releases with the effect of unobserved firm-specific shocks, and my estimates would be biased upwards.

In order to provide a more transparent test of pre-trends and better isolate the effect of news releases from other confounders, I complement the strategy in Equation 2.9 with an event study approach. The idea is to inspect how news releases matter for the average firing behavior of firms, before and after a negative news release "shock". I define such shock as a quarter in which the firm experiences a number of negative (positive) performance articles greater than the 95th percentile in the firm-specific distribution over the period 2000-2017. I estimate the equation:

$$P(\text{End of CEO app.})_{ft} = \alpha + \sum_{q \neq -1} \beta_q \mathbf{I}[q = t] + \sum_j \gamma_j \mathbf{I}[j = y(it)] + \sum_k \delta_k \mathbf{I}[k = s(it)] + \phi_f + u_{ft} \tag{1.3}$$

$P(\text{End of CEO app.})_{ft}$ is the probability of CEO appointment ending in firm f and quarter t. $\mathbf{I}[q = t]$ is an indicator variable for whether the event is experienced q quarters from quarter t. $\mathbf{I}[j = y(it)]$ is an indicator variable for calendar year y, and $\mathbf{I}[k = s(it)]$ is an indicator for quarter of tenure $s(it)$, measured at the time of the event. ϕ_f represents firm fixed effects. The coefficients δ_k are CEO-specific and identified in presence of firm fixed effects because I observe multiple CEO appointments within the same firm. Standard errors are clustered at the firm level.[11]

Results

I start by considering the results of the difference-in-differences specification in Equation 2.9. The results are shown in Table 1.6. Differently from previous work in the corporate finance literature, I do not attempt to classify the nature of turnover as due to resignation, retirement, or firing.[12] Instead, I compare three definitions of turnover. The first one defines turnover as any quarter in which I observe a CEO appointment ending. Former CEOs are often retained as lower-ranked executives, consultants, or board members. The second definition indicates whether the CEO is no longer retained in the company under any job title. Finally, I also look at

[11] I do not control for firm performance, because executive turnover is a determinant of firm performance.

[12] In an influential paper, for example, Parrino (1997) provides a method for classifying CEO turnover as due to firing, resignation, or retirement.

whether the CEO moves to a private company or a company with smaller sales relative to the departing company, or whether information on the following job move is missing.[13] In order to avoid measurement error coming from the fact that for 10% of the quarter-position observations there are no news releases, in Panel A of Table 1.6 I focus on a subset of companies that is frequently covered in news media. High coverage firms include firms for which the median number of articles in a quarter is above the median across all firms (which corresponds to 4 articles per quarter). Table 1.6 shows that an additional negative article increases the probability of an appointment ending the following quarter by 2.3% relative to the sample mean. The number of negative articles released in a quarter is also strongly associated with the probability of being dismissed from all job appointments (column 2), and the probability of moving to a private or smaller firm (column 3). As for the number of positive news articles, the effect is small and insignificant in all specifications. Because the estimates are stable across high-coverage firms and the full sample of firms, measurement error is likely not to be the driver of the smaller coefficient on the number of positive articles.

Column 2 of Table 1.6 implies that an increase in the number of negative news

[13] In order to rank companies in terms of size, I divide companies into two-digit SIC sectors and obtain deciles of yearly sales in a given sector-year. I define a company as "smaller" if the difference in yearly sales with the departing company is greater than two deciles in the fiscal year preceding the job move.

from the 25th to to 75th percentile is associated with an increase in the replacement probability of about 0.2 percentage points, from 2.7% to 2.9% – which corresponds to a 7.5% increase. Such number is difficult to compare to previous work on CEO turnover, given that the corporate finance literature typically focuses on determinants of turnover other than the media. For example, Jenter and Lewellen (2019) find that turnover probability increases from 3.3% to 6.68% as stock market performance decreases from the 70th to the 20th percentile of the distribution in the preceding year. Jenter and Kanaan (2015) find that forced turnover probability increases from 2.05% to 4.14% as industry performance falls from the 90th to the 10th percentile in the preceding year. Relative to previous work on performance, my estimates for the effects of news are smaller. The discrepancy may be due to several reasons. Previous literature looks at turnover decisions made at a one-year horizon, whereas my estimates are at a quarterly frequency. Second, my results do not intend to reflect a causal effect, and may be downward-biased.

In order to assess the relevance of pre-trends, and to better isolate the effect of news releases from other confounders, I turn to the results of the event study analysis. The results for negative news releases are shown in Panel A of Figure 1.6. Relative to one quarter before the event, the probability of turnover jumps discontinuously at the time of the event, and peaks to 3.3 percentage points one quarter after. Panel

A of Figure 1.6 also suggests the absence of anticipation effects in any of the five quarters leading to the event. Similarly to the results in Table 1.6, Panel B of Figure 1.6 shows that there is no clear pattern of replacement decisions following positive news releases. Although not necessarily causal, the results presented in this section suggest that news releases, and in particular negative news, are highly predictive of CEO replacement. Positive news seem to have very little effect. The asymmetry may be due to several reasons. First, when making hiring and firing decisions firms may seek to screen out particularly poor candidates in order to avoid very bad outcomes, rather than selecting the very top ones (Bergman et al., 2020). Moreover, people tend to put more weight on negative relative to positive information, a pattern that is known in psychology as negativity bias (Trussler and Soroka, 2014). Such interpretation could be even more relevant when disclosing information may harm a firm's reputation. As summarized by a famous quote: "It takes many good deeds to build a good reputation, and only one bad one to lose it". [14]

[14] Benjamin Franklin.

1.5 A model of executive turnover with news selection

1.5.1 Model

I build a model of CEO turnover with event-dependent news reporting decisions. The model serves two purposes. First, it provides a framework for understanding how negative media focus affects firms' replacement decisions. I will then take the model to the data, and quantify how much the bias in news selection matters for CEO turnover, and for turnover in female-led firms. The model builds on classic models of employer's learning in the spirit of Jovanovic (1979). In every period, the firm observes current and past signals of firm performance and makes one decision: whether keeping or dismissing the CEO. Media outlets monitor performance realizations and decide which realizations to cover. News selection is event-dependent: worse performance events are more likely to be covered. Public beliefs on CEO ability are informed by the news, and taken into account by the board of directors when making the turnover decision.

Model set-up

Turnover In every period, the firm decides whether to keep or dismiss the CEO. The turnover decision d_t maximizes expected utility:

$$V(\mathbf{x_t}) = \max_{d_t, d_{t+1}, \ldots} E_t \left(\sum_{s=t}^{\infty} \delta^{s-t} u_s(d_s, \mathbf{x_s}) | d_t, \mathbf{x_t} \right)$$

where $\mathbf{x_t}$ is the vector of state variables. The optimization problem can be written recursively as a Bellman equation:

$$V(\mathbf{x_t}) = \max_{d_t} E_t(u_t(d_t|\mathbf{x_t})) + \delta V_{t+1}(\mathbf{x_{t+1}}|d_t, \mathbf{x_t})$$

The intra-period utility from keeping the CEO is a function of firm performance, q_t, the public reputation of the CEO, $\hat{q}_t$, and an idiosyncratic shock ϵ_t^K. As it is standard in the discrete choice literature, ϵ_t^K is distributed with a Type 1 Extreme Value distribution with scale parameter τ:

$$u_t(1, \mathbf{x_t}) = \kappa_1 q_t + \kappa_2 \hat{q}_t + \epsilon_t^K$$

If instead the firm dismisses its CEO, it pays the dismissal cost c and obtains a

random utility shock ϵ_t^D:

$$u_t(0, \mathbf{x_t}) = \kappa_1 q_t + \kappa_2 \hat{q}_t - c + \epsilon_t^D$$

Let $V_t^K = E_t(u_t(1, \mathbf{x_t})) + \delta V_{t+1}((\mathbf{x_{t+1}})|1, \mathbf{x_t})$ and $V_t^D = E_t(u_t(0, \mathbf{x_t})) + \delta V_{t+1}((\mathbf{x_{t+1}})|0, \mathbf{x_t})$ be the choice-specific value functions for keeping and dismissing the CEO. These correspond to:

$$V_t^K(\mathbf{x_t}) = \kappa_1 \cdot E_t(q_t|\mathbf{x_t}) + \kappa_2 \cdot E_t(\hat{q}_t|\mathbf{x_t}) + \delta V_{t+1}((\mathbf{x_{t+1}})|1, \mathbf{x_t}) + \epsilon_t^K = \bar{V}_t^K + \epsilon_t^K \tag{1.4}$$

$$V_t^D(\mathbf{x_t}) = -c + V_0(\mathbf{x_0}) + \epsilon_t^D = \bar{V}^D + \epsilon_t^D \tag{1.5}$$

$V_0(\mathbf{x_0})$ in Equation 1.5 represents the utility from hiring a new CEO: if the board dismisses its CEO, the problem "reverts" to time $t = 0$, when the information set is given by the board's priors. The expectations $E_t(q_t|\mathbf{x_t})$ and $E_t(\hat{q}_t|\mathbf{x_t})$ in Equation 1.4 come from the fact that at the time of making the turnover decision, the board has not yet observed current CEO performance. CEO performance is a function of CEO ability. The board learns about CEO ability over time, as more and more performance signals are observed. Suppose learning is complete after T time periods.

Then the asymptotic choice-specific value functions are:

$$V^K(\mathbf{x_T}) = E_T(\kappa_1 q_T|\mathbf{x_T}) + E_T(\kappa_2 \hat{q}_T|\mathbf{x_T}) + \delta V_{T+1}(\mathbf{x_{T+1}}|1, \mathbf{x_T}) + \epsilon^S = \bar{V}^K + \epsilon^K$$

$$V_T^D(\mathbf{x_T}) = -c + V_0(\mathbf{x_0}) + \epsilon^D = \bar{V}^D + \epsilon^D$$

and the optimization problem is $V(\mathbf{x}) = \max_{d\in\{0,1\}} \left(V^K(\mathbf{x}), V^D(\mathbf{x})\right)$

Learning environment

Private learning At the time of CEO appointment ($t = 0$), the board of directors has a normally distributed prior belief on CEO ability:

$$\alpha \sim N\left(\alpha_0, \sigma_0^2\right),\ \sigma_0^2 > 0$$

In every period of CEO tenure t, firm performance q_t is realized. q_t is a function of CEO ability and a random shock ϵ_t^q:

$$q_t = \alpha + \epsilon_t^q$$

$$\epsilon_t^q \sim N\left(0, \sigma_q^2\right),\ \sigma_q^2 > 0$$

The first expectation in Equation 1.4, $E_t(q_t|\mathbf{x_t})$, is given by:

$$E_t(q_t|\mathbf{x_t}) = E_t(q_t|q_1,...,q_{t-1})$$

which is calculated by the board using Bayes' rule, based on its prior and the history of performance signals up to $t-1$.

Public learning The media monitor performance realizations q_t and decide which realizations to make public. Publishing decisions are represented by the random variable S_t: when the media decide to publish event q_t, $S_t = 1$ is realized, and the signal q_t is made available to the public. Since the publication decision S_t is publicly observable, the board can calculate the second expectation in Equation 1.4, $E_t(\hat{q}_t|\mathbf{x_t})$:

$$E_t(\hat{q}_t|\mathbf{x_t}) = E_t(E_t(q_t|q_1,...,q_{t-1},S_1,...,S_{t-1})|\mathbf{x_t}) = E_t(q_t|q_1,...,q_{t-1},S_1,...,S_{t-1}) = \hat{q}_t$$

Note that even if private and public learning are about the same object – firm performance q_t – the two posterior beliefs $E_t(q_t|q_1,...,q_{t-1})$ and $E_t(q_t|q_1,...,q_{t-1},S_1,...,S_{t-1})$ are allowed to differ, depending on the sequence of random variables $S_1,...,S_{t-1}$.

News selection

The availability of the public signal ($S_t = 1$) depends on the realized event: the key assumption on the publication rule is that *negative* performance events are considered more newsworthy. The assumption is in line with the empirical evidence presented in the previous sections, and is an empirical regularity when looking at news reporting decisions.

Definition 1. Negative events are considered more newsworthy if the odds ratio of a publication conditional on the realization q_t, $\frac{P(S_t=1|q_t)}{P(S_t=0|q_t)}$, is decreasing in q_t.

Definition 1 introduces a selection bias in the way news media report information: worse performance realizations are more likely to be reported. I present a number of propositions that should help clarify how the selection bias introduced by Definition 1 produces downward-biased beliefs relative to a situation where information is reported in an unbiased manner.

First, it is possible to show that under the publication rule in Definition 1, the distribution of unpublished events first order stochastically dominates the distribution of published events.

Proposition 1. *If* $\frac{P(S_t=1|q_t)}{P(S_t=0|q_t)}$ *is decreasing in* q_t*, then* $P(q_t \leq q|S_t = 0) \leq P(q_t \leq q|S_t = 1)$.

Proof. In the Appendix.

The proposition states that published events come from a "worse" distribution relative to unpublished events. In fact, realizations on the left tail of the unconditional distribution of firm performance, $P(q_t)$, are more likely to be published. From first order stochastic dominance, it follows that the mean of published events is lower than the mean of unpublished events: $E(q_t|S_t = 1) \leq E(q_t|S_t = 0)$: on average, the value of firm performance is lower when it is made public relatively to when it is not. The next proposition states that the mean of published events is also lower than the unconditional mean of all events:

Proposition 2. *The mean of published events is lower than the unconditional mean of all events, that is:* $E(q_t|S_t = 1) \leq E(q_t)$.

Proof. In the Appendix.

Figure 1.7 helps visualizing these results. Figure 1.7 plots the unconditional distribution of firm performance $P(q_t)$, the conditional probability of publication $P(S_t = 1|q_t)$, and the distribution of published firm performance, $P(q_t|S_t = 1)$. The unconditional distribution $P(q_t)$ – the blue solid line – is centered around zero. The conditional probability of an event being reported, $P(S_t = 1|q_t)$, increases monotonically as q_t decreases, and approaches 1 for very low values of q_t. The distribution of reported events $P(q_t|S_t = 1)$ – the blue dashed line in Figure 1.7 – is shifted to the left relative to the unconditional distribution $P(q_t)$: the average event published

by the media is a "worse" event relative to the average event in the true underlying distribution.

The fact that a publication is more likely to be available for negative performance realizations has implications for how public beliefs are updated. Public beliefs are more likely to be updated with negative performance information, and therefore are likely to be downward-biased. In Figure 1.8, I simulate the evolution of private and public beliefs over time for a draw of 100 CEOs. Simulating the evolution of private and public beliefs requires making assumptions on the distributions' parameters such that the publication rule in Definition 1 is satisfied. In Appendix B I describe how the distributional assumptions on CEO ability and firm performance q_t allow characterizing the family of conditional distributions $P(q_t|S_t = 0)$ and $P(q_t|S_t = 1)$ such that the publication rule in Definition 1 is satisfied. While the two learning processes in Figure 1.8 start from the same prior, they diverge over time, with public beliefs converging to a lower value in the long run. The result is due to the bias introduced by news selection, which is such that low realizations of firm performance are more likely to be published.

Consequences for turnover

Given the state variables up to time $t-1$, at every point in time t the board compares the expected benefit of keeping a CEO with the value of dismissing the CEO. News selection biased towards the negative performance states has two opposite effects on turnover. On the one hand, the value of dismissing the CEO decreases. Since the board is forward-looking, the negative selection bias will decrease the value of a hire to the firm, thus lowering the value of the firm's outside option. Everything else constant, decreasing the value of the firm's outside option decreases turnover. On the other hand, the negative selection bias decreases the value of keeping a CEO, especially as time moves on and private and public beliefs start diverging. Holding everything else constant, decreasing the value of keeping a CEO increases turnover. Theoretically, it is ambiguous which of the two effects will prevail. In practice, the second effect will turn out to be much stronger than the first one. Biased news selection affects the value of keeping a CEO at higher tenure levels: private and public beliefs diverge at higher tenures, as information becomes abundant and uncertainty decreases. But because the value of a hire is in present discounted terms, the value of keeping a CEO at higher tenure levels is more discounted relative to lower tenures. Therefore, the decrease in the value of a hire will not be enough to offset the loss in utility caused by biased public beliefs, and turnover will increase

relative to the case with no selection bias.

1.5.2 Model's testable implications

The model delivers testable implications on the sensitivity of turnover to the arrival of negative news. At every point in time, public beliefs on CEO ability are summarized by the average value of the signal up to time t:

$$\hat{h}(t) = \frac{\sum_{s=1}^{t-1}(q_s|S_s = 1)}{t-1}$$

where q_s is firm performance at time s and S_s is the publication indicator. Two testable implications delivered by the model are:

(i) Conditional on tenure t, the probability of turnover increases as $\hat{h}(t)$ decreases.

(ii) Conditional on tenure t, the sensitivity of turnover to the arrival of a negative news increases as $\hat{h}(t)$ decreases.

The two predictions follow from Bayesian updating. In the model, an executive is dismissed at tenure time t if the board's posterior belief on the ability of the executive falls below the endogenous replacement threshold set by the board. Prediction (i) states that the probability of turnover increases the closer beliefs get to

the replacement threshold. Prediction (ii) states that, the closer beliefs get to the replacement threshold, the higher the marginal impact of a negative news on the probability of crossing such threshold. The two predictions are proved formally in Appendix A.1.

1.5.3 Empirical tests of model's implications

In order to take the model's predictions to the data, I map the average value of the signal up to time t, $\hat{h}(t)$, to the data as follows. I divide histories in two groups. In the first group of histories, the CEO experienced a negative publication in more than 50% of the tenure-quarters up to tenure-quarter $t-1$, where a negative publication is a news with sentiment at the bottom 10% of the distribution. Because such first group experienced negative news for a larger fraction of tenure time, I define it as the long history group. The second group includes all other histories, and is defined as the short history sample. I only consider high-coverage firms to avoid noise coming from sparse news data. In order to keep tenure constant, I focus on tenure-quarters below 2.5 years of tenure. According to prediction (i), the probability of turnover should be higher in the long history sample. The empirical test is presented in Table 1.7. Conditional on tenure, on average the turnover probability is higher for long histories. The results are robust to the addition of sector fixed effects and

performance controls. According to prediction (ii), the marginal effect of a negative news on turnover should be higher in the long history sample. The prediction is tested in Table 1.8. I present results from the same specification as in Equation 2.9, separately for the two subsamples and replacing firm fixed effects with sector fixed effects. In line with the prediction, the coefficient on the number of negative news is positive and significant for the subsample of long histories, and much smaller for the subsample of short histories. When the probability of crossing the firing threshold is high, an additional negative news is likely to lead to the replacement of an executive.[15]

1.6 Implications for turnover in female-led firms

1.6.1 Female-led firms

I now turn the case of female-led firms. Consider two types of firms: female- and male-headed firms ($g = F, M$). The two types are identical in terms of prior distribution of CEO ability and unconditional distribution of firm performance, but differ with respect to one feature: the media are more likely to publish a low per-

[15] I perform additional checks to isolate the effect of news from the effect of firm performance. First, I control for average realized ROA up to tenure-quarter t. Second, I cut the sample so as to focus on tenure-quarters with ROA below the median across the entire sample. The results are robust in both specifications and available in the replication files.

formance realization for a female-headed firm relative to a male-headed firm. The assumptions on female- and male-headed firms are:

(i) The prior ability distribution is the same in the two firms: $\alpha^F \sim \alpha^M \sim N\left(\alpha_0, \frac{1}{\tau_0}\right)$

(ii) The unconditional distribution of firm performance is the same in the two firms: $P^F(q_t) \sim P^M(q_t)$;

(iii) There exists a performance threshold q^* such that $P^F(S_t = 1|q_t) > P^M(S_t = 1|q_t)$ for every $q_t < q^*$.

The model's assumptions are supported by empirical evidence. I discuss Assumption (i) and present corroborating evidence in Appendix section A.4.2. Assumption (ii) has been discussed in section 1.4.2, where I verify that there is no significant difference in performance between the two types of firms, neither when looking at sales or stock price returns. Assumption (iii) has been discussed in the first part of the paper, in section 1.4.2.

Given the assumptions, the intuition from the homogeneous case carries through the case of heterogeneous firms. When performance is low, the public is *more likely* to observe the public signal for female-led firms relative to male-led firms. For the same firm performance distribution, at any point in time public beliefs on female-led firms are likely to be more pessimistic relative to public beliefs for an average firm

(see Figure 1.8), and the incidence of turnover will tend to increase for a female-led firm relative to an average firm. In other words, for the same firm performance the board requires female CEOs to be of a higher ability relative to male CEOs, so as to offset the loss in utility caused by more pessimistic public beliefs. A clear prediction of the model is that female CEOs will be more positively selected than their male counterparts around the time of replacement. As shown by Figure A.2.6 in the Appendix, the prediction is supported in the data. Right before the dismissal of a male CEO, firm performance follows a clear downward trend. No such trend is visible for female CEOs, for which firm performance is a much weaker predictor of dismissal.

1.6.2 Model calibration

I solve the dynamic programming problem numerically through value function iteration and obtain the board's optimal dismissal policy. The Appendix provides a detailed description of the model's solution, and the simplifying assumptions I make in order to deal with state space dimensionality.

The goal of the calibration is to obtain the model's parameters for the sample of male CEOs, and then feed in the differential media coverage measured in the data for women to run counterfactual simulations.

A period t in the model corresponds to a tenure year in the data. I require male CEOs to be observed at least 4 years to be included in the sample. I drop positions that lasted less than a quarter, and positions with incomplete news or performance data. The final sample includes 1,624 male CEOs.

I measure firm performance q_t as industry-adjusted ROA. I choose industry-adjusted ROA as opposed to sales or stock prices for several reasons. Relative to ROA, sales confound profitability with firm size. Stock prices typically react to news information as soon as it becomes available. Moreover, using ROA makes the results comparable with previous research (Bertrand and Schoar, 2003; Taylor, 2010). I measure CEO dismissal as an appointment ending, and the CEO not being appointed in the same company under any job title in the following quarter (that is, the indicator in the second column of Table 1.6). In order to calibrate the model, I divide parameters into three blocks.

Pre-set parameters The first block of parameters is set outside of the model. I set the discount factor δ to 0.9 to match the annual discount rate in Taylor (2010). Because utility is defined up to a scale, the scale parameter τ of the taste shock distribution is not identified and normalized to 1.

News selection The key insight for mapping news data to the model is that news coverage for an event in the data mirrors a selection probability in the model. Therefore, the parameters governing news selection are set to match the coverage bias in the data. First, I fix $\mu_{q|S=1}$, the mean of published performance events. I proceed as follows. For every parameter search, I simulate the probability distribution of performance events $f(q)$, and re-weight quartiles of $f(q)$ so as to match news coverage for events of different sentiment quartiles in Figure 1.3. I set $\mu_{q|S=1}$ equal to the mean of the re-weighted probability distribution of performance events. I then search over a grid of possible values for $\mu_{q|S=0}$ – the mean of unpublished events – and select a value so as to match the slope in Figure 1.3, namely such that an event at the bottom 25% of the performance distribution has a 12-percent higher chance of being selected by the news relative to an event at the top 75%:

$$\frac{P(S=1|q<q_{25})-P(S=1|q>q_{75})}{P(S=1|q>q_{75})}=0.12$$

where q_j is the j-th percentile of the performance distribution. Note that for values α_0, $\mu_{q|S=0}$, and $\mu_{q|S=1}$ the unconditional probability of publication ω is fixed, because the relationship $\alpha_0=\omega\cdot\mu_{q|S=1}+(1-\omega)\cdot\mu_{q|S=0}$ has to hold.

Simulated method of moments The rest of the parameters are pinned down

by moments in the data using simulated method of moments. The target moments and their value are described in Table 1.9. I explain the simplifying assumptions I make in order to obtain a fully identified model, and how each moment is informative of different parameters. First, I run a $AR(1)$ regression for firm profitability: $q_{it} = \lambda_0 + \lambda_1 q_{it-1} + \epsilon_{it}$. The profitability intercept λ_0 is informative about the average skill across CEOs, and helps pin down the mean of the prior distribution of CEO ability, α_0. λ_1 captures how persistent firm performance is within a firm-CEO, and is informative about the within-CEO dispersion in firm performance σ_q. Since λ_1 is high, implying a low within-CEO variance in firm performance, σ_q is likely to be low. A low σ_q would imply that the board learns CEO ability quickly, a statement that does not fit the data. Moreover, firm performance is only an imperfect predictor of CEO turnover, and the board of director's assessment of the CEO relies on several unobserved factors outside of the model. Therefore, I assume that the board's perceived dispersion of firm performance is $\tilde{\sigma}_q$.[16] $\tilde{\sigma}_q$ is pinned down by mean performance by tenure time. In the data, mean performance increases with tenure: in the model this is due to the changing composition in the pool of CEOs, as

[16] In order to better fit the data, Taylor (2010) assumes that the board relies on a noisy private signal in addition to firm performance. Although the assumption in Taylor (2010) is slightly different, the purpose is the same: firm performance is only an imperfect predictor of CEO turnover, and many other factors outside of the model contribute towards explaining turnover. Note that by assuming that the board's perceived standard deviation in firm profitability is $\tilde{\sigma}_q \neq \sigma_q$ I am imposing a departure from rational expectations. In a different context, the same assumption is made by Hoffman and Burks (2020).

the less able are dismissed and the more able remain in office. Because $\tilde{\sigma}_q$ governs how good the board is at detecting high-ability CEOs, mean productivity by tenure has to increase slower as $\tilde{\sigma}_q$ increases.

I discretize the news history variable – the average share of negative news over total news up to time $t-1$ – into three categories, corresponding to terciles of the distribution. In the model, such categories map to a "publication state", where states with more negative publications are associated to more pessimistic public beliefs. I then run the regression: $q_{it} = \delta_0 + \delta_1 pub2_{it-1} + \delta_2 pub3_{it-1} + \epsilon_{it}$, where q_{it} is industry-adjusted ROA for firm i in quarter t, and $pub2_{it}$ and $pub3_{it}$ are two indicators for whether the history of negative publications in firm i and quarter t belong to the second or third tercile (the omitted category is $pub1_{it}$, corresponding to the first tercile). δ_0, δ_1, and δ_2 capture average firm profitability by publication state, where worse publication states are associated with lower firm profitability. The three coefficients are informative about the standard deviation of CEO skill, σ_0: the further apart the three publication states, the higher the dispersion in CEO skill. The survival rate at lower tenure levels and mean profitability over time help pin down the utility parameter κ_1, the board's utils per dollar of firm profits. In order to have a fully identified model, I set κ_2 – the board's utils per dollar of firm profits as perceived by public beliefs – equal to κ_1. The assumption is needed because true

firm performance and public performance are highly correlated by construction, and intra-period utility is linear in both components. Therefore, it is hard to find a data moment that shifts κ_1 without affecting κ_2, so that the two parameters can be separately identified. The assumption implies that true and public performance have equal weight in the board's intra-period utility. Although it imposes a further restriction, the assumption may not be too strong, as the data show that both true and public performance – as proxied by performance news – are predictive of CEO turnover. Finally, the cost of dismissal c is pinned down by the survival rate at higher tenure levels: as tenure increases, learning converges and the firing cost plays a larger role. Note that the cost of dismissal c is in board's utils. c represents the board's perceived cost from dismissing a CEO, which includes not only monetary costs – such as severance payments – but also costs in terms of reputation and shareholders' satisfaction with the board's operations.

1.6.3 Model parameters

The estimation results for the model's parameters are presented in Table 1.10. In Figure A.2.7, I show the model fit for the target moments. The model fits the target moments fairly well. In Figure 1.9, I also show the model fit of the survival rate and average mean profitability over the first 15 years of tenure. The model

fits the data quite well, in particular when considering the turnover hazard. In the data, the average turnover hazard over the first 15 years of tenure is 4.01%. In the model, the average turnover hazard over the first 15 years is 4.23%. The numbers are close to those estimated by previous literature: Taylor (2010), for example, finds the incidence of turnover to range between 3.45% and 4.04% over the period 2000-2006. When looking at mean profitability by tenure time, the model overpredicts profitability in the first tenure period, and underpredicts profitability in the last tenure period. To further assess the sensitivity of my results, I compare my estimates with previous literature. To the best of my knowledge, Taylor (2010) is the only paper structurally estimating a model of CEO turnover, so I will mostly compare my estimates to Taylor (2010), although he analyzes an earlier time period (1990-2006) relative to my sample. First, the prior mean CEO ability in my model is higher than in Taylor (2010). The prior mean CEO ability is 2.06% of assets in my model, and 1.24% in Taylor (2010). The difference is possibly due to the high profitability intercept in my data (see Figure 1.9). The prior variance of CEO ability is equal to 4.84%, thus being within the range of previous estimates: 2.72% in Taylor (2010) and 7% in Bertrand and Schoar (2003). The within-CEO variance of firm performance, σ_q, is 2.28% in my model, and 3.61% in Taylor (2010). The difference is due to different modeling assumptions. In his model, Taylor (2010)

assumes that firm profitability follows a $AR(1)$ process, and thus 3.61% represents the residual profitability variance after accounting for persistence. Because I do not have the $AR(1)$ assumption in my model, the variance of firm profitability has to be relatively low in my model in order to fit the high persistence of profitability within a firm-CEO. The perceived within-CEO variance of firm performance – $\tilde{\sigma}_q$ – is high, and equal to 9.65%. As explained in the previous subsection, the assumption is needed in order to slow down board's learning, which most likely relies on additional factors outside of the model when making the turnover decision. To fit the same feature of the data, Taylor (2010) assumes that the board relies on an additional private signal of firm performance, whose variance is also large and close to my estimate (9.51%). Finally, the cost of dismissal is 3.46% in my model, which is close to the estimate in Taylor (2010) (3.95%). Given the average value of firm assets in my estimating sample, a cost of 3.46% implies that the board of the average firm behaves as if dismissing a CEO costed $347 million to the firm.

1.6.4 Counterfactual simulations

Having estimated the structural parameters of the turnover model, I can run counterfactual simulations and quantify of much the bias in news selection is able to account for differential turnover in female-led firms. In practice, given the parame-

ters of the ability and profitability distributions and the board's utility parameters, I change the parameters governing news selection so as to match the differential bias measured empirically for women. Row B of Table 1.11 shows the news selection bias and implied hazard for the baseline model. As explained in Section 1.6.2, in the baseline version of the model news selection bias is defined as the differential selection probability of an event at the bottom of the profitability distribution relative to the top, and is set to match the slope in Figure 1.3. The implied turnover hazard averaged over the first 15 years in office of the CEO is 0.0423. Removing the selection bias in Panel A of Table 1.11 decreases the turnover hazard by 9.7% relative to the baseline model. Removing the selection bias implies that a performance realization at the bottom of the distribution has the same publication probability than an event at the top, and makes public beliefs aligned with the board's beliefs. The absence of the selection bias creates two opposite effects on turnover relative to the baseline model. On the one hand, the firm's outside option increases, because the absence of the selection bias will increase the value of a hire to the firm. Everything else constant, increasing the firm's outside option increases turnover. On the other hand, the absence of the selection bias increases the value of keeping a CEO, especially for CEOs with higher tenures, when beliefs are less volatile and both private and public beliefs converge to their long-run value. Holding everything else

constant, increasing the value of keeping a CEO decreases turnover. The second effect turns out to be much stronger than the first one.

Row C of Table 1.11 sets news selection to match the evidence for women. For women, a performance event at the bottom 10% of the sentiment distribution generates 41% additional coverage relative to an event at the top 10% (Table 1.3). Increasing the selection bias from 12% to 41% increases turnover by about 3%. Given that the differential turnover for female CEOs is around 20%, the difference in news selection explains around 15% of the differential turnover observed for women.[17] Because in the baseline version of the model all CEOs are homogeneous, the counterfactual in row C assumes that a female CEO will always be replaced by another female CEO. I run an additional counterfactual assuming that the firm's outside option is a male. In practice, I replace the value of a hire implied by the model with the value of a male hire as implied by the baseline model. Because the value of hiring a male is higher, the turnover hazard increases a little, but the difference is negligible. Such small difference is due to the fact that the bias in news selection matters the most for high values of tenure, as public and private beliefs diverge. Since the value of a hire is in present discounted terms, high tenure values

[17] Differences in turnover by gender are shown in Appendix Table A.3.5. Keller et al. (2020) finds that women executives are more likely to leave the firm by 2 percentage points relative to comparable men. Their estimates include CEOs and lower-ranked executives as well. Gupta et al. (2020) finds that female CEOs are 45% more likely to be fired relative to comparable men. Their definition of turnover uses the classification introduced by Parrino (1997).

are more discounted by the board. The model implies that the value of hiring a female CEO is almost the same as the value of hiring a male ex ante, but not ex post: as tenure increases, female CEOs will generate less value to the firm relative to their male counterparts. Finally, in row D I simulate the model feeding in the news selection bias estimated for women at their first appointment. The estimates in Tables 1.4 and 1.5 imply that a for women at their first appointment a performance event at the bottom 10% of the sentiment distribution generates 67% additional coverage relative to an event at the top 10%. Under such counterfactual scenario, the news selection bias increases turnover by around 4.7% relative to the baseline model, thus accounting for about 24% of the gap in the turnover hazard measured for female appointments relative to male appointments.

1.7 Conclusions

I show that negative media focus can affect firms' replacement decisions in public companies, especially in firms led by executives who may be perceived as outsiders, such as women and newly appointed CEOs. My results suggest that between 15% and 20% of the differential turnover measured in female-led firms may be excessive and would not take place in a counterfactual scenario where news media behave as for the average CEO. In several contexts, information disclosure has beneficial

effects, especially when it exposes negative outcomes. Public attention improves firms' accountability and mitigates the agency problem between the management and stakeholders. My results are not in contradiction with this view, and show that adverse effects of information disclosure can arise when the negative focus is systematic and more severe for leaders for which there is less information to begin with. From the standpoint of policymakers, my results are positive rather than normative, and highlight one additional hurdle for improving access to top leadership positions. Policymakers have recognized such goal as a priority: in Europe and more recently in California, for example, gender quotas on corporate boards have been mandated to improve female representation in the boardroom. My results could also point to a shortcoming of quotas. Because quotas are a tool to promote outsiders, unintended effects may arise in the short run if outsiders are penalized in news media. My results suggests that quotas likely need to be implemented for a long period of time in order to fully reap their beneficial effects.

My work tackles a specific mechanism that can apply to an extraordinarily special group of workers: CEOs. More research is needed in order to understand how to promote the career advancement of women in professional environments and at the top echelons of the earnings distribution, a goal that has been shown to improve efficiency (Hsieh et al., 2019). As argued by Terviö (2009), public information plays

a crucial role in highly-paid professions in which performance on the job is publicly observable. Further research is needed in order to understand more broadly how the media influence the executive labor market – for example, through executive compensation.

My paper is concerned with studying the consequences of media focus rather than the reasons behind specific editorial decisions. Understanding the reasons behind editorial decisions would be an important question to answer in order to understand the sources of inefficiencies, and better guide policymaking. I leave the answer to such important question to future research.

1.8 Figures

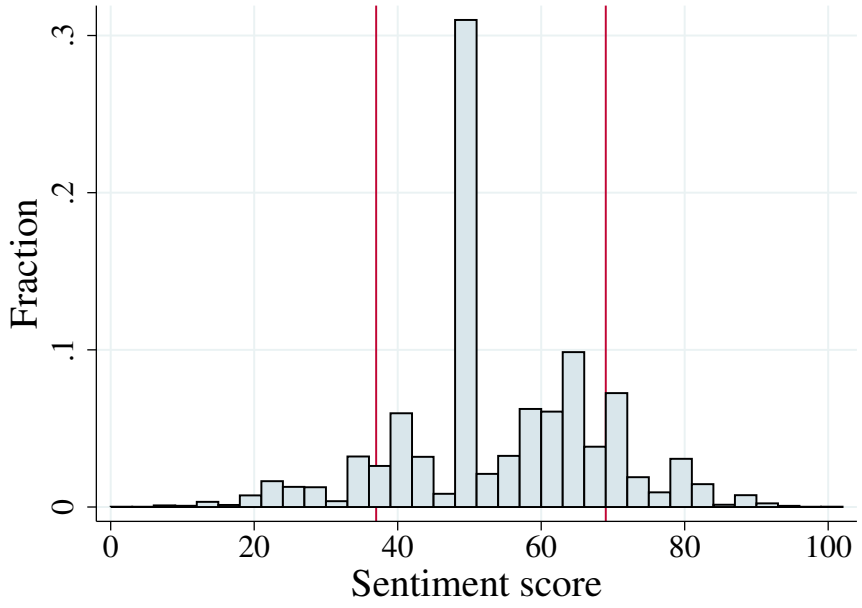

Figure 1.1: Sentiment score distribution

Notes: Sentiment score distribution of news events. The vertical bar on the left represents the 10th percentile of the distribution (score = 37), the vertical bar on the right represents the 90th percentile of the distribution (score = 69).

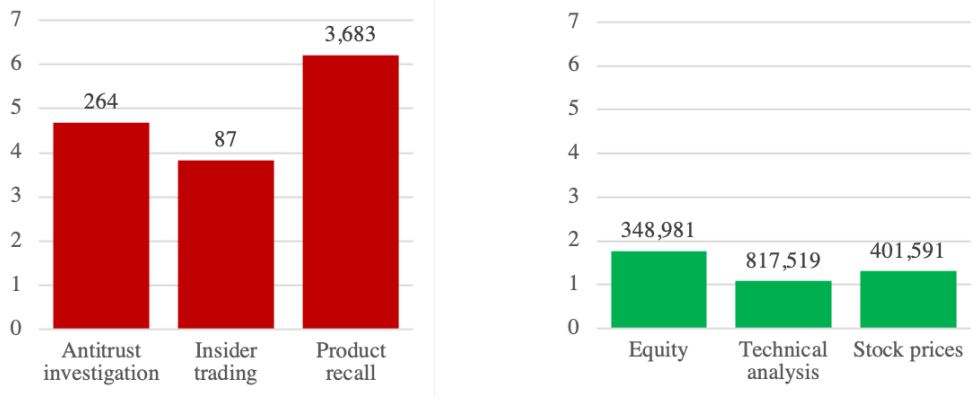

(a) Unusual news events

(b) Common news events

Figure 1.2: Unusual events are more covered

Notes: Average number of articles for different categories of news events. The number on top shows the number of events of each category in the dataset. The total number of events in the dataset is 6,923,931.

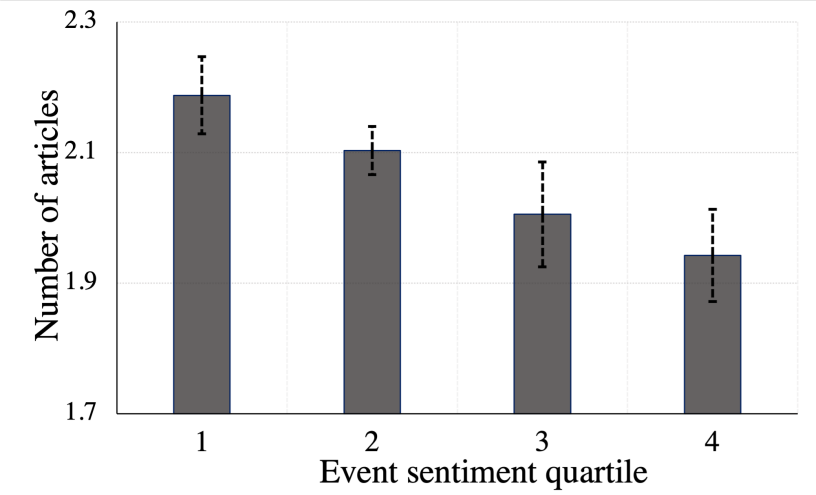

Figure 1.3: Negative events are more covered

Notes: Linear prediction from a regression of the number of articles for an event on sentiment quartiles, log(sales), event category fixed effects, (35 categories), firm and time fixed effects. The dotted bars show the 90% confidence interval.

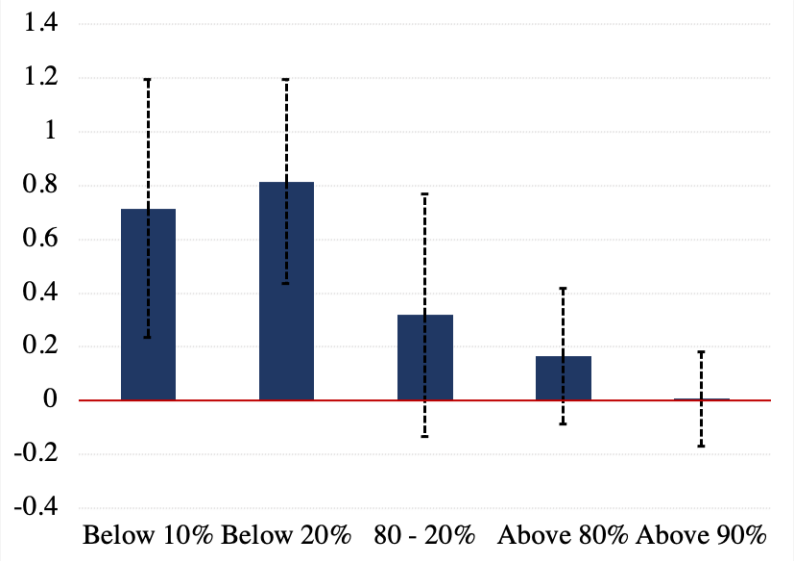

Figure 1.4: Coverage for a news event, female differential

Notes: The graph shows the coefficient on the female indicator from a regression of the number of articles for a news event on news event sentiment, CEO characteristics, firm characteristics, firm and time fixed effects. Every bar corresponds to the coefficient from a different regression. The y-axis unit is number of articles for a news event. The x-axis shows the news sentiment distribution corresponding to each subsample of news events. The plotted coefficients are shown in Table 1.3. The dotted bars show the 90% confidence interval.

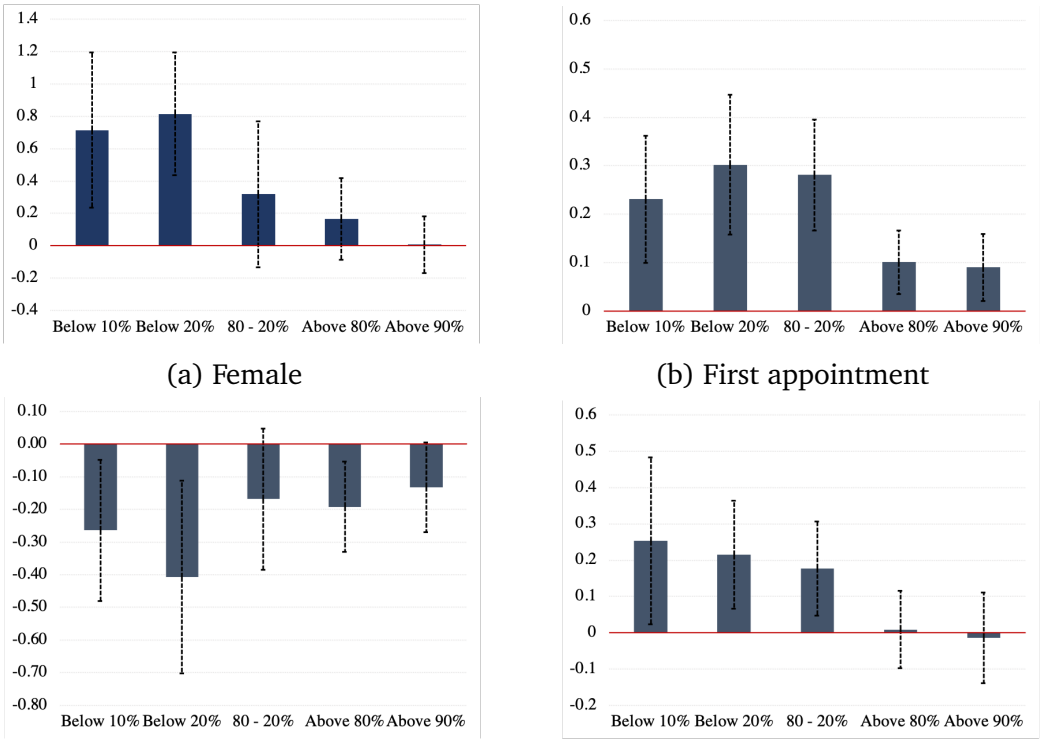

(a) Female (b) First appointment

(c) Tenure (d) First year

Figure 1.5: Coverage for a news event, outsider CEOs

Notes: The graph shows the coefficient on the female indicator (Figure a) and measures indicating outsider status (Figures b,c, and d) from a regression where the dependent variable is represented by the number of articles for a news event. Every regression controls for news event sentiment, CEO characteristics, firm characteristics, firm and time fixed effects. Every bar corresponds to the coefficient from a different regression. The y-axis unit is number of articles for a news event. The x-axis shows the news sentiment distribution corresponding to each subsample of news events. The plotted coefficients are reported in Tables 1.3, 1.5, and 1.4. The dotted bars show the 90% confidence interval.

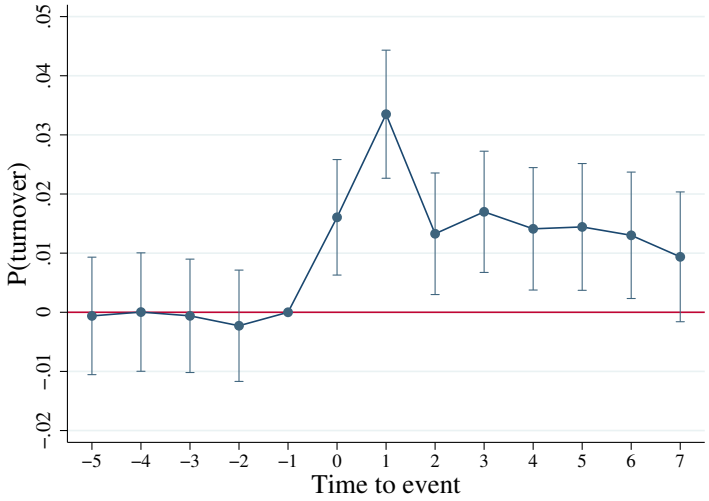

(a) Effect of negative news on turnover

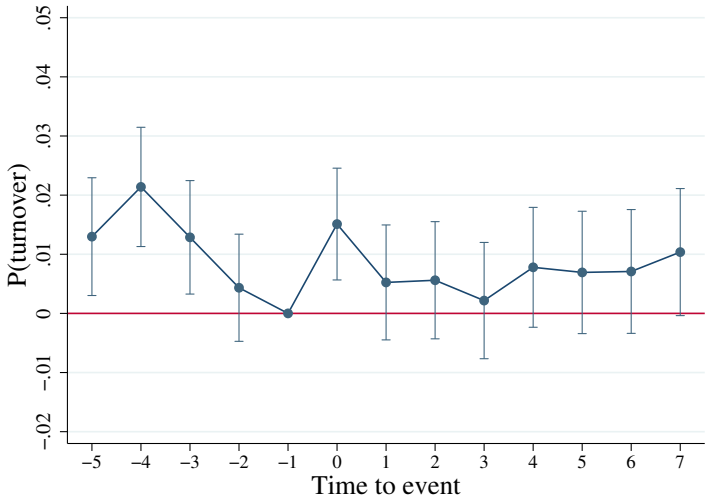

(b) Effect of positive news on turnover

Figure 1.6: Effect of news on turnover

Notes: The graph shows the results from the event study for the effect of negative news releases on firms' replacement decisions. Quarterly observations between 2000 and 2017. A negative news release event is defined as a quarter in which the firm experiences a number of negative (positive) performance articles greater than the 95th percentile of the firm-specific distribution over the period 2000-2017. The dependent variable is an indicator for whether the CEO is in the first quarter of tenure. The omitted time period corresponds to the quarter preceding the event. Standard errors are clustered at the firm level. The vertical bars represent 95% confidence intervals.

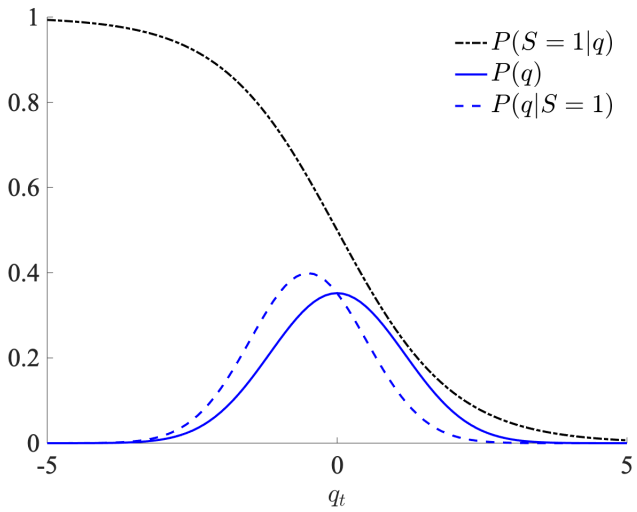

Figure 1.7: News selection

Notes: The graph shows how the unconditional distribution $P(q)$ and the conditional distribution $P(q|S=1)$ map to a conditional selection probability $P(S=1|q)$. The blue solid distribution represents the unconditional distribution of firm performance $P(q)$, and the blue dotted distribution is the distribution of published firm performance $P(q|S=1)$. The dotted probability represents the conditional publication probability (or news selection function) $P(S=1|q)$ of firm performance.

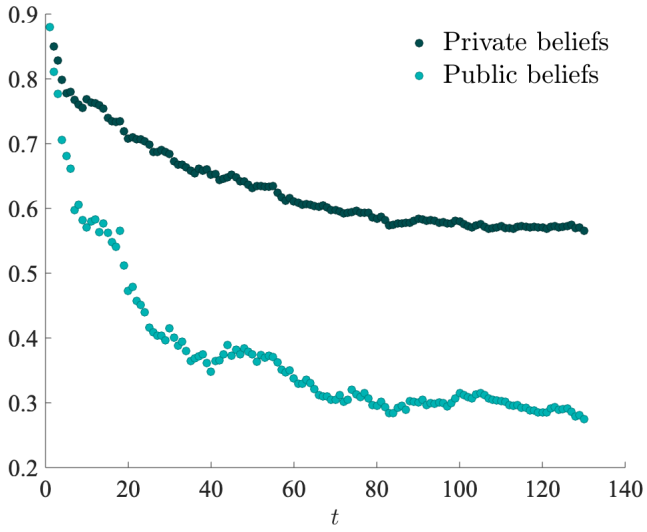

Figure 1.8: Private and public beliefs

Notes: Simulation of private and public beliefs over the long run for a draw of 100 CEOs from the distribution $\alpha \sim N(\alpha_0 = 0.88, \sigma_0 = 2.42)$ (from Taylor, 2010). The dark series on top represents private beliefs, whereas the lighter series at the bottom represents public beliefs.

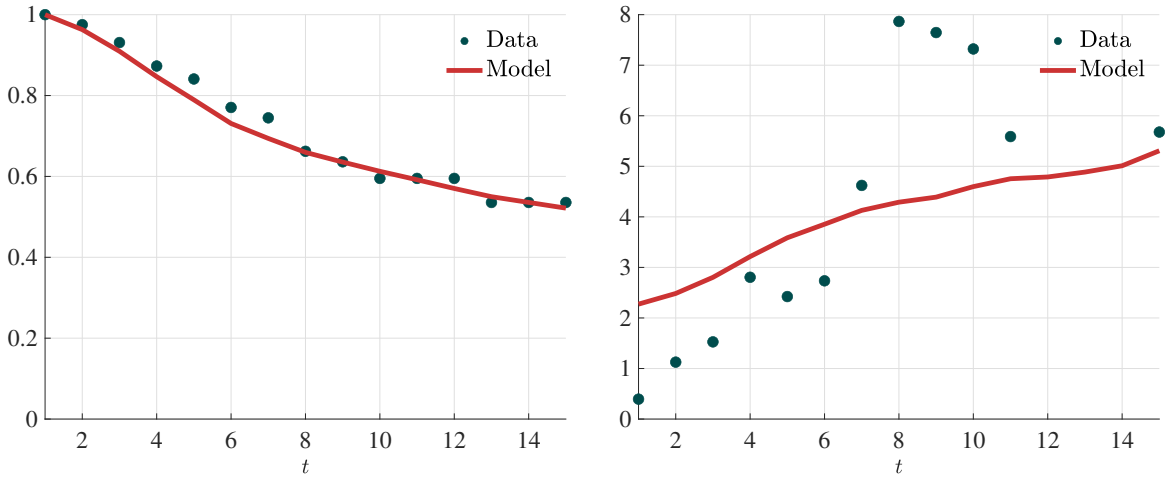

(a) Survival function (b) Average profitability

Figure 1.9: Model fit

Notes: Survival function in Panel (a) and average profitability by tenure year in Panel (b). The model is simulated for 1,624 CEOs using the parameters in Table 1.10.

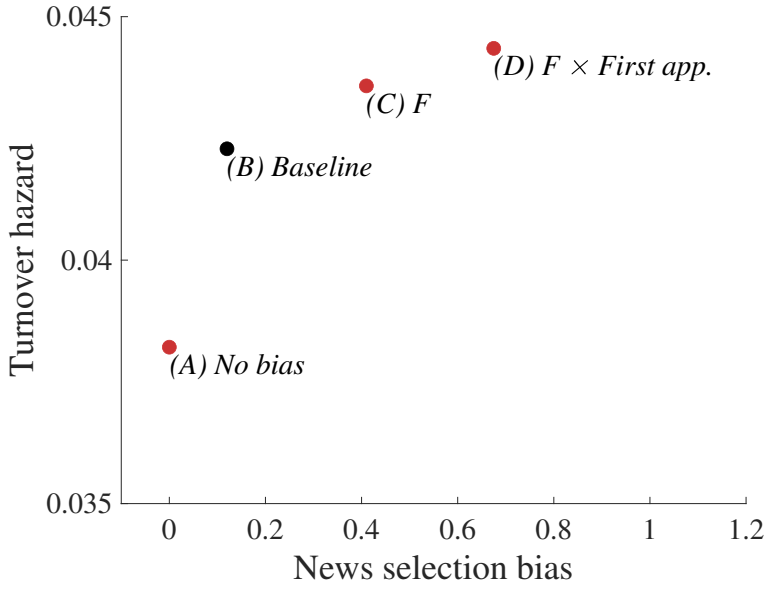

Figure 1.10: Baseline model and counterfactuals

Notes: *(A) No bias*. Counterfactual simulation, obtained by removing the news selection bias. *(B) Baseline*. Baseline model. The parameters for the baseline model are in Table 1.10. *(C) F*. Counterfactual simulation, obtained by simulating the model feeding in the news selection bias estimated for women CEOs. *(D) F × First app*. Counterfactual simulation, obtained by simulating the model feeding in the news selection bias estimated for women CEOs at their first appointment.

1.9 Tables

Table 1.1: CEOs, by gender

	Women		Men		Difference	p
	Mean	SD	Mean	SD		
Panel A1. Individual characteristics						
Age	52.59	7.06	52.60	8.22	-0.01	0.993
Born in the US	0.94	0.24	0.92	0.28	0.02	0.574
Bachelor's degree	0.35	0.48	0.29	0.45	0.07	0.115
Master's/MBA/Prof. degree	0.36	0.48	0.45	0.50	-0.09	0.055
Doctorate degree	0.16	0.36	0.15	0.36	0.00	0.914
Number of qualifications	1.89	1.20	1.92	1.09	-0.03	0.801
Appointment dur. (days)	650.42	730.26	697.77	765.54	-47.35	0.514
Tenure in company (years)	7.32	9.37	6.53	8.29	0.79	0.374
Network size	1,325.24	1,617.72	1,169.26	1,420.68	155.97	0.229
Total number of boards	2.01	1.61	1.93	1.65	0.08	0.662
Panel A2. After end of appointment:						
End of all appointments	0.29	0.45	0.21	0.41	0.08	0.077
Private or smaller firm/ missing move	0.22	0.42	0.17	0.37	0.05	0.264
Panel B. Board characteristics						
Gender ratio	0.76	0.11	0.91	0.10	-0.15	0.000
Number of directors	8.23	2.06	8.46	2.51	-0.24	0.378
Panel C. Firm characteristics						
Assets	5,214.55	20,908.14	8,123.41	73,910.11	-2908.87	0.686
Employees	9.70	29.61	8.37	28.74	1.32	0.644
Sales	3,523.68	16,343.86	2,555.31	9,929.19	968.37	0.343
Gross profits	921.21	3,071.22	842.44	3,397.51	78.77	0.815
Market value	2,889.77	8,973.50	3,698.14	16,623.82	-808.37	0.623
Primary sector	0.03	0.18	0.15	0.35	-0.11	0.000
Consumer sector	0.26	0.44	0.15	0.35	0.12	0.000
Service sector	0.71	0.46	0.71	0.46	0.00	0.946
Number of positions	129		2,897			
Number of firms	105		1,938			

Notes: Source: Panel A and B: BoardEx, 2000-2017, Panel C: Compustat, 2000-2017. Data for the sample of matched news-firm-CEOs. Individual and board characteristics are measured in the year of the appointment (except *Appointment duration*), whereas firm characteristics are measured in the year before the appointment.

Table 1.2: News events, by sentiment

	Number of events published	Share of total	Sentiment score: Mean	Sentiment score: SD	Articles per event: Mean	Articles per event: SD	Days per event: Mean	Days per event: SD
Panel A. Negative events (< 10th ptile)								
earnings	29,466	0.30	26.94	7.89	2.02	4.47	1.06	0.26
analyst-ratings	18,577	0.49	31.50	6.37	1.32	1.13	1.03	0.18
order-imbalances	13,757	0.64	32.96	0.51	1.37	0.77	1.12	0.39
legal	12,109	0.76	22.10	1.98	5.09	11.89	1.30	0.70
revenues	4,626	0.81	24.84	6.73	2.81	9.40	1.11	0.38
regulatory	3,582	0.84	22.30	0.71	3.21	5.87	1.22	0.58
price-targets	3,464	0.88	25.87	7.32	1.18	0.71	1.02	0.15
products-services	3,153	0.91	28.87	5.84	4.47	14.86	1.23	0.71
credit-ratings	2,366	0.94	29.52	4.94	2.13	1.79	1.03	0.18
Panel B. Positive events (> 90th ptile)								
products-services	82,220	0.20	66.31	5.14	3.95	24.51	1.20	1.25
earnings	54,526	0.32	72.00	8.80	2.20	4.40	1.06	0.25
technical-analysis	46,004	0.43	58.96	1.65	1.09	0.41	1.04	0.24
analyst-ratings	37,148	0.52	71.21	10.85	1.19	0.62	1.02	0.14
stock-prices	34,630	0.60	63.00	0.00	2.45	6.55	1.14	0.43
acquisitions-mergers	28,369	0.67	66.46	7.10	2.26	6.43	1.10	0.35
partnerships	23,371	0.73	61.04	0.19	2.97	5.53	1.12	0.42
equity-actions	20,373	0.78	64.35	6.67	1.98	4.31	1.07	0.29
revenues	18,351	0.82	66.70	11.31	2.15	3.88	1.06	0.30

Notes: Source: RavenPack News analytics, 2000-2017. Data for the sample of matched news-firm-CEOs. Negative events in Panel (A) are events at the bottom 10% of the sentiment distribution. Positive events in Panel (B) are events at the top 90% of the sentiment distribution.

Table 1.3: News coverage for an event, and firm and CEO characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
		By event sentiment:				
	All	*Below 10%*	*Below 20%*	*20% – 80%*	*Above 80%*	*Above 90%*
Female	0.539***	0.715**	0.815***	0.318	0.166	0.007
	(0.150)	(0.292)	(0.231)	(0.276)	(0.154)	(0.107)
Network size	-0.000***	-0.000***	-0.000***	-0.000	-0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Born in the US	-0.023	-0.051	-0.045	-0.079	0.152***	0.098
	(0.065)	(0.096)	(0.096)	(0.065)	(0.056)	(0.061)
Number of qual.	0.046	-0.052	0.081*	0.044	-0.009	0.022
	(0.031)	(0.045)	(0.048)	(0.031)	(0.024)	(0.026)
Age	-0.047	-0.032	-0.223**	0.025	0.018	-0.002
	(0.059)	(0.065)	(0.102)	(0.059)	(0.042)	(0.052)
Age sq.	0.001	0.000	0.002**	-0.000	-0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)
Tenure	-0.030***	-0.047**	-0.066***	-0.015	-0.032**	-0.019*
	(0.008)	(0.021)	(0.021)	(0.012)	(0.013)	(0.011)
Tenure sq.	0.000***	0.000***	0.000**	0.000***	0.000**	0.000*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Sentiment score	0.002	-0.050***	-0.018***	0.022	0.031***	0.032***
	(0.001)	(0.008)	(0.007)	(0.019)	(0.005)	(0.005)
Log(sales)	-0.037*	0.011	-0.075*	-0.049**	0.000	0.006
	(0.022)	(0.037)	(0.040)	(0.022)	(0.040)	(0.044)
Log(assets)	0.015	0.202**	0.281***	-0.048	-0.038	-0.059
	(0.038)	(0.079)	(0.068)	(0.041)	(0.050)	(0.064)
Quarter FE	Y	Y	Y	Y	Y	Y
Year of app. FE	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	
N	591,257	62,384	123,344	351,837	116,047	93,128
Mean	2.239	2.361	2.189	2.292	2.131	2.179

Notes: Observations are news stories released between 2000 and 2017 in the full sample of matched news-CEO firms. The dependent variable is represented by the total number of articles for a news event. The estimating specification is equation 1.1 in the text. Standard errors are clustered at the position level. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

Table 1.4: News coverage for a negative event and outsider CEOs

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	A. Sentiment below 10%						
Female	0.695**	0.694**	0.688**	0.687**	-0.079	-0.085	0.698*
	(0.329)	(0.322)	(0.322)	(0.327)	(0.310)	(0.284)	(0.356)
First appointment		0.204**	0.190**		0.148*		
		(0.082)	(0.079)		(0.078)		
First year			0.253*			0.162	
			(0.139)			(0.113)	
Founder				0.186**			0.170*
				(0.093)			(0.095)
F × First app.					1.291**		
					(0.521)		
F × First year						1.718*	
						(0.911)	
F × Founder							0.393
							(0.658)
	B. Sentiment below 20%						
Female	0.469	0.469	0.456	0.458	-0.091	-0.126	0.514
	(0.334)	(0.320)	(0.319)	(0.331)	(0.310)	(0.215)	(0.364)
First appointment		0.263***	0.261***		0.225***		
		(0.085)	(0.084)		(0.084)		
First year			0.219**			0.148**	
			(0.088)			(0.075)	
Founder				0.288			0.297
				(0.189)			(0.196)
F × First app.					0.977*		
					(0.523)		
F × First year						1.401**	
						(0.642)	
F × Founder							-0.469
							(0.600)
CEO char.	Y	Y	Y	Y	Y	Y	Y
Tenure quadratic	Y	Y	N	Y	Y	N	Y
Quarter FE	Y	Y	Y	Y	Y	Y	Y
Sector FE	Y	Y	Y	Y	Y	Y	Y

Notes: Observations are news events released between 2000 and 2017 in the full sample of matched news-CEO firms. The dependent variable is represented by the total number of articles for a news event. The estimating specification is Equation 1.1 in the text, in which firm fixed effects are replaced with sector fixed effects. CEO characteristics include network size, a dummy for whether the CEO was born in the US, the number of qualifications, a quadratic in age, and year of appointment fixed effects. The number of observations is 62,384 in Panel A and 123,344 in Panel B. Standard errors are clustered at the position level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 1.5: News coverage for a positive event and outsider CEOs

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	A. Sentiment above 90%						
Female	-0.053	-0.058	-0.058	-0.058	-0.172*	-0.073	-0.024
	(0.131)	(0.128)	(0.129)	(0.131)	(0.102)	(0.216)	(0.144)
First appointment		0.091**	0.085**		0.084*		
		(0.042)	(0.041)		(0.043)		
First year			-0.013			-0.019	
			(0.075)			(0.076)	
Founder				0.108			0.116
				(0.088)			(0.090)
F × First appointment					0.185		
					(0.211)		
F × First year						0.043	
						(0.230)	
F × Founder							-0.300
							(0.300)
	B. Sentiment above 80%						
Female	-0.047	-0.050	-0.052	-0.051	-0.224*	-0.114	-0.017
	(0.156)	(0.152)	(0.152)	(0.156)	(0.118)	(0.214)	(0.171)
First appointment		0.102***	0.099**		0.091**		
		(0.039)	(0.039)		(0.040)		
First year			0.008			-0.003	
			(0.064)			(0.065)	
Founder				0.120			0.128
				(0.083)			(0.085)
F × First appointment					0.290		
					(0.246)		
F × First year						0.143	
						(0.178)	
F × Founder							-0.332
							(0.307)
CEO char.	Y	Y	Y	Y	Y	Y	Y
Tenure quadratic	Y	Y	N	Y	Y	N	Y
Quarter FE	Y	Y	Y	Y	Y	Y	Y
Sector FE	Y	Y	Y	Y	Y	Y	Y

Notes: Observations are news events released between 2000 and 2017 in the full sample of matched news-CEO firms. The dependent variable is represented by the total number of articles for a news event. The estimating specification is Equation 1.1 in the text, in which firm fixed effects are replaced with sector fixed effects. CEO characteristics include network size, a dummy for whether the CEO was born in the US, the number of qualifications, a quadratic in age, and year of appointment fixed effects. The number of observations is 93,318 in Panel A and 116,047 in Panel B. Standard errors are clustered at the position level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 1.6: Turnover and news

	(1)	(2)	(3)	(4)	(5)	(6)
	End of:		*CEO moves to:*	*End of:*		*CEO moves to:*
	CEO app.	All app.	Private or smaller firm, missing move	CEO app.	All app.	Private or smaller firm, missing move
	A. High-coverage firms			*B. All firms*		
Negative articles	0.0016**	0.0011**	0.0007**	0.0016***	0.0012***	0.0010**
	(0.001)	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)
Positive articles	0.0002	-0.0002	-0.0003	-0.0000	-0.0003	-0.0002
	(0.001)	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)
CEO and firm controls	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
N	9,541	9,541	9,541	15,668	15,668	15,668
Number of clusters	751	751	751	1,250	1,250	1,250
Mean of dep. var.	0.0695	0.0300	0.0252	0.0722	0.0311	0.0272

Notes: Quarterly observations between 2000 and 2017. High coverage firms (Panel A) include firms for which the median number of quarterly articles is above the median across all firms. CEO controls include network size, a dummy for whether the CEO was born in the US, the number of qualifications, a quadratic in age, and year of appointment fixed effects. Firm controls include quarterly ROA. All regressions include controls for the total number of articles released in a quarter. Standard errors are clustered at the firm level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 1.7: Average turnover probability, by history

		Short history of neg. coverage		*Long history of neg. coverage*			
		Mean	*SD*	*Mean*	*SD*	*Diff.*	*p-value*
End of:	CEO app.	0.074	0.262	0.084	0.278	-0.01	0.092
	All app.	0.025	0.155	0.039	0.193	-0.014	0.000
CEO moves to:	Private/smaller firm, missing move	0.021	0.143	0.031	0.173	-0.01	0.004

Notes: The sample includes high-coverage firms only, namely firms for which the median number of articles in a quarter is above the median across all firms, and tenure-quarters of CEO appointment between 2 and 10 quarters of tenure. Long histories (first panel) are tenure-quarters where a negative publication was available in more than 50% of the tenure-quarters up to tenure-quarter $t-1$, where a negative publication is a news with sentiment at the bottom 10% of the distribution. include tenure-quarters where a negative publication was available in less than 50% of the tenure-quarters up to tenure-quarter $t-1$.

Table 1.8: Turnover and news, by history

	(1)	(2)	(3)	(4)	(5)	(6)
	End of:		*CEO moves to:*	*End of:*		*CEO moves to:*
	CEO app.	All app.	Private or smaller firm, missing move	CEO app.	All app.	Private or smaller firm, missing move
	A. Long history of neg. coverage			*B. Short history of neg. coverage*		
Negative articles	0.0022**	0.0016**	0.0011*	-0.0016	0.0003	-0.0005
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
Positive articles	-0.0008	-0.0003	-0.0000	-0.0029	-0.0013	-0.0015*
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
Sector FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
N	2,659	2,659	2,659	2,683	2,683	2,683

Notes: Quarterly observations between 2000 and 2017. The sample includes 751 high-coverage firms, that is firms for which the median number of articles in a quarter is above the median across all firms, and tenure-quarters of CEO appointment between 2 and 10 quarters of tenure. Long histories (Panel A) are tenure-quarters where a negative publication was available in more than 50% of the tenure-quarters up to tenure-quarter $t-1$, where a negative publication is a news with sentiment at the bottom 10% of the distribution. Short histories (Panel B) include tenure-quarters where a negative publication was available in less than 50% of the tenure-quarters up to tenure-quarter $t-1$. CEO controls include a quadratic in tenure, a quadratic in age, and a female indicator. Firm controls include quarterly ROA. All regressions include controls for the total number of articles released in a quarter. Standard errors are clustered at the sector level. $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

Table 1.9: Target moments

Moment	Description		Value
	(a) Average firm profitability by publication state:		
δ_0			6.156
δ_1	$\tilde{q}_{it} = \delta_0 + \delta_1 pub2_{it-1} + \delta_2 pub3_{it-1} + \epsilon_{it}$		−3.089
δ_2			−7.751
	(b) Firm profitability AR(1):		
λ_0	$\tilde{q}_{it} = \lambda_0 + \lambda_1 \tilde{q}_{it-1} + \epsilon_{it}$		0.238
λ_1			0.968
	(c) Survival function:		
$Surviv_j$	Survival function at $t = j$:	$j = 2$	0.931
		$j = 6$	0.745
		$j = 10$	0.595
		$j = 14$	0.535
	(d) Firm profitability by tenure:		
$Avgperf_j$	Average firm performance at $t = j$:	$j = 2$	1.526
		$j = 6$	4.621
		$j = 10$	5.588
		$j = 14$	5.677

Notes: Target moments used for the parameter's estimation through method of simulated moments. $\tilde{q}_{it}$ represents ROA for firm-CEO i at time t in excess of industry performance. $pub2_{it}$ and $pub3_{it}$ are dummies for the second and third tercile in the share of negative news at time t. Time t is in years.

Table 1.10: Model parameters: calibration

	(1) Pre-set parameters	
δ	Discount factor	0.9
τ	Scale of taste shock	1
	(2) Simulated Method of Moments	
	Distributions	
α_0	Prior mean of CEO ability	2.06
σ_0	Prior st. deviation of CEO ability	4.48
σ_q	Within-CEO st. deviation of firm performance	2.28
$\tilde{\sigma}_q$	Perceived within-CEO st. deviation of firm performance	9.65
	Utility	
κ_1	Utils per unit of firm performance	0.50
κ_2	Utils per unit of public firm performance	0.50
c	Dismissal cost	3.46
	(3) Calibrated to match evidence	
	News selection	
$\mu_{q\|S=1}$	Mean of published firm performance	1.84
ω	Unconditional probability of publication	0.96

Notes: Implied model's parameters. The first block of parameters is pre-set: δ matches Taylor (2010) and τ is normalized to 1. The second block is obtained through simulated method of moments using the moments in Table 1.9 as targets. The third block is calibrated to match the slope in news coverage for events with different sentiment in Figure 1.3.

Table 1.11: Counterfactuals

(1) *Model*	(2) *News selection bias*	(3) *Implied hazard*	(4) *Diff. with baseline*	(5) *Explained gap in turnover*
(A) No bias	0%	0.0382	– 9.69%	
(B) Baseline	12%	0.0423	0%	
(C) F	41%	0.0436	3.07%	15.37%
(D) F × First app.	67%	0.0443	4.73%	23.64%

Notes: *(A) No bias*. Counterfactual simulation, obtained by removing the news selection bias. *(B) Baseline model*. The parameters for the baseline model are in Table 1.10. *(C) F*. Counterfactual simulation, obtained by simulating the model feeding in the news selection bias estimated for women CEOs. The news selection bias is obtained from Table 1.3. *(D) F × First app*. Counterfactual simulation, obtained by simulating the model feeding in the news selection bias estimated for women CEOs at their first appointment. The news selection bias is obtained from Tables 1.4 and 1.5.

Chapter 2

Do Board Gender Quotas Matter? Selection, Performance and Stock Market Effects

2.1 Introduction

Women are underrepresented among top leadership positions. The *glass ceiling* – the invisible barriers which prevent women from reaching upper-level positions – is still a dominant phenomenon. Even in countries in which women participate more

in the labor market, only a minority makes it to the highest positions. According to the World Economic Forum (2020), only 58% of the gender gap in economic opportunities has been closed around the world. At the slow speed experienced over the period 2006-2020, it will take another 260 years to vanish completely.

Gender quotas have been proposed to accelerate the process towards economic gender equality and to promote women's empowerment.[1] Norway pioneered the introduction of gender quotas for boards of directors in 2005. Italy, France, and Germany, among others, followed. Three European Directives on gender quotas have been proposed and are currently under consideration, while the debate is open in many other countries. In September 2018, California was the first US state to approve a bill for the inclusion of women on the boards of directors of public companies. The approval of the bill came after a harsh debate between promoters and opponents. In fact, gender quotas are controversial. They have been widely advocated for achieving a gender-balanced representation in top positions, a crucial goal for achieving economic gender equality (see OECD, 2012; IMF, 2014). Yet, opponents argue that they violate meritocracy, with costly consequences. By equalizing outcomes rather than opportunities, quotas come with the risk of promoting less-qualified individuals, who are likely to perform poorly (Holzer and Neumark,

[1] In parallel, gender quotas have been introduced to reduce political gender gaps, the other crucial dimension of gender inequalities (see Section 2 for more references).

2006). For instance, if highly qualified women cannot be found, board gender quotas may produce negative effects on the performance of companies and negative stock market reactions. Are these negative consequences the unavoidable cost of achieving more gender-balanced representation?

What we know so far about the effects of board gender quotas on the economy is based on the Norwegian experience. In late 2003, a law was approved in Norway mandating at least 40% representation of each gender on the board of companies listed on its stock market (existing firms had to comply by January 2008, while new firms by January 2006). The Norwegian law imposed a dramatic and rapid transformation of the composition of boards of directors (Enjolras and Sivesind, 2012; Huse and Seierstad, 2013). Research has shown that the Norwegian law has been effective at increasing the number of women at the very top of the earnings distribution, but it has not been able to reduce gender gaps overall (Bertrand et al., 2019). Matsa and Miller (2013) found a decrease in operating profitability of firms after quotas. An early study by Ahern and Dittmar (2012) shows that the increase in the number of women on boards in Norway imposed a significant cost on firm value and stock market returns. However, new evidence supports the non-significant result. Eckbo et al. (2021) discuss the validity of the result in Ahern and Dittmar (2012) and show that, by using a more robust specification, the negative market

reaction in Norway becomes insignificant and the negative effect on operating profitability is only a short-term effect related to the financial crisis. Yet, Norway is a very peculiar case, being a top performer country in gender statistics worldwide. In a different context, the soft quota introduced in Spain has been analyzed (De Cabo et al., 2019): the voluntary approach adopted by Spain with an economic incentive to comply (compliant firms may receive a preference for the tendering of public contracts) but without sanctions to firms not in compliance with the recommended target, did not reach the goal of promoting gender-balanced boards. Thus, the analysis of a different case is needed to assess the effects of board gender quotas outside Norway and in a general perspective.

This paper provides new evidence based on the introduction in July 2011 of board gender quotas in Italian listed companies. The so-called "Golfo-Mosca" (by the names of the two proposers) law mandates gender-balanced representation on the board of directors and statutory auditors of publicly listed companies. Unlike in Norway, in Italy quotas are temporary, and the measure will be in place only for three consecutive board elections. The required target of representation of either gender is set for all companies subject to the law (independently on the size of the board) at 1/5 for the first election after August 2012, to be increased to 1/3 for the following two board elections. In December 2019 the law is extended for additional

three elections with an increase of the quota up to 40%. In our analysis, we focus on the first target and thus on the short-term effects of quotas.

The Italian case is a unique and innovative opportunity to assess the economic effects of board gender quotas. On the methodological side, we can exploit staggered board elections: the quota law does not apply to all firms at the same time, as in Italy board elections are held every three years on a date decided by each firm, and the year of board election depends on the past.[2] More generally, Italy features a very conservative gender culture, and ranks poorly in Europe in almost all gender statistics (Profeta et al., 2014): in the last ten years, women's participation in the labor force has remained stable at around 47%, the lowest value in Europe, if we exclude Malta. In this context, the quota policy was perceived as the only possible way to start the process towards gender equality. But at which cost? A country with no economic growth certainly cannot afford to bear substantial economic costs. In this paper we find no evidence of significant costs, neither for firms' performance nor for stock market returns, associated with the introduction of board gender quotas in Italy.

We focus on the short-term impact of the board gender quota law and consider

[2]Given the staggered board elections across Italian firms, we can use the reform period as an instrument, rather than being forced to rely on the pre-quota percentage of female directors interacted with year dummies – which may raise endogeneity concerns – to make the instrument vary across firms, as in Ahern and Dittmar (2012) and Bertrand et al. (2019).

the period 2007-2014. To perform our analysis, we manually collected individual data on all members of the boards of Italian listed companies in the period 2007-2014 (4,732 unique individuals), as well as firm-level data on relevant outcomes of these companies (243 companies) and stock-market prices. With these data, we are able to address three fundamental questions that allow us to evaluate the effectiveness of quotas in the process of promoting women's empowerment *vis-à-vis* their possible costs: Do the composition of the boards and the characteristics of board members change after the introduction of quotas? Do firms' outcomes, such as economic performance and variability of stock prices, change after the introduction of gender quotas? How does the stock market react to the approval and implementation of board gender quotas?

Our results can be summarized as follows. First, we exploit the staggered compliance of Italian firms with the gender quota law to study how the boards change following the appointment of women directors. We consider several characteristics of board members, such as gender, age, and education. We find that, when gender quotas are enforced, firms show a higher share of women directors (well above the required threshold), higher average education levels of all members of the board, and fewer older members than before the quota. Our results suggest that gender quotas change the selection process of the entire board. Despite having to select

more women, we do not find an increase in female board members belonging to the firm owner's family, nor a clear increase in the average number of positions held by each woman.

Second, to address endogeneity and move closer to the causal effect of gender quotas on firm performance, we use the reform period – a measure exogenous to firms' decisions – as an instrument for the share of female directors. We only focus on the short-term effects of the introduction of gender quotas. Our results show that quotas are not associated with significant effects on firm performance as measured by number of employees, assets, production, profits, ROA, Tobin's Q, and debts. When looking at stock market performance, we show that the presence of female directors reduces the variability of stock prices – a crucial dimension of performance for listed companies, not explored before in connection with board gender quotas.

Third, we run an event study at the date of the approval of the law and show that there is no significant difference in returns between more gender-diverse and less gender-diverse Italian firms. We also perform an event study at the date of board election, which happens on a different day for each firm. By comparing the returns of companies with a smaller share of women in the pre-reform board composition (i.e., farther from the quota target) and companies that were closer to the target, we exclude that the introduction of quotas is associated with a costly reduction of

stock market returns. If anything, investors positively reacted to the appointment of women on boards in elections that happened after the approval of the quota law.

The paper is organized as follows. The next section reviews the related literature and Section 3 describes the Italian law. Section 4 presents the data that we will use throughout the different analyses performed in the paper. The three sections that follow present analyses related to our three fundamental questions: the impact of the law on board characteristics, the impact of the law on firm performance and the variability of stock prices, and the impact of the announcement of the quota law and the appointment of female directors on stock market prices. Each of these sections presents first the empirical methodology adopted and then the results. Conclusions are in Section 8. Additional evidence is provided in the Appendix.

2.2 Related literature

Board gender quotas have been previously analyzed with reference to Norway. Several studies assess whether the increased female representation in top positions due to the Norwegian quota had any impact on firm economic performance. Matsa and Miller (2013) find that firms affected by the quota law fired fewer workers, thus causing an increase in relative labor costs and employment levels and reducing short-term profits. Ahern and Dittmar (2012) show that gender quotas caused

negative market reactions, due to the appointment of young and less-expert members. However, Nygaard (2011) shows that this effect depends on asymmetric information between independent members of the boards and the managers (Ferreira, 2015). Moreover, Eckbo et al. (2021) show that, once a more robust methodology is used, the negative result found by Ahern and Dittmar (2012) vanishes.

On cultural grounds, Italy is very far from Norway and closer to Spain, which represents an interesting comparison country to our study. Spain, however, followed a different way and introduced a soft quota, which, by being a simple recommendation without sanctions, was not able to produce a substantial increase in the share of women on boards (De Cabo et al., 2019).

More generally, our paper relates to the large literature on the effects of gender quotas, a controversial policy (Pande and Ford, 2012; Profeta et al., 2014). The main argument in favor of the adoption of gender quotas is their effectiveness as a means to equalize opportunities in specific areas where women face systematic barriers due to discrimination or persistent stereotypes (Holzer and Neumark, 2006). These policies may lead to a redistribution of jobs, positions, contracts, or parliament seats in favor of women, and thus allow for a fair distribution of rewards of good jobs. Moreover, if women who benefit from affirmative action are largely qualified to successfully perform the tasks they are appointed to, the benefits do not

remain within the group of women but spread to the entire economy (Hsieh et al., 2019). If women accumulate more human capital that raises their productivity, these policies may even increase efficiency (Conde-Ruiz et al., 2017). Quotas are an instrument (often considered the only one) to break down the masculine monopolistic power, which obviously does not lead to an equal outcome, but probably neither to an efficient one. Critics of affirmative actions, instead, share the view that the underrepresentation of women is not due to discrimination, but is merely the result of women's choices, especially related to fertility and motherhood.[3] Thus, by equalizing outcomes rather than opportunities, affirmative actions may lead to the promotion of less-qualified individuals. Not only there may be the risk of decreasing average quality if there are not enough women with the appropriate qualifications to be appointed, but a "mismatch" may occur if women are allocated to positions in which they are unable to perform successfully. Recent studies have also doubted the effectiveness of quotas in reducing gender inequalities in specific contexts (Bagues and Esteve-Volart, 2010). Bertrand et al. (2019) found that gender quotas for listed companies in Norway improved the representation of female employees at the very top of the earnings distribution within affected firms, while they had no trickle-

[3] A greater involvement of women in the economy may also have beneficial effects on cultural development. Dominant gender stereotypes and social norms have played a crucial role in generating gender gaps. Learning from other women's experience in the labor market may generate a virtuous and persistent circle of gender equality through changes to the cultural process (Fernández, 2013; Fernández et al., 2004).

down effects on gender gaps.[4] In the context of academia, Bagues et al. (2017) find that the gender composition of evaluation committees does not necessarily increase the chances for women to be promoted, thus limiting the effectiveness and desirability of gender quotas. In politics, recent studies have shown that gender quotas are not at odds with meritocracy, as they help increase the quality of representatives. In the Italian political context, gender quotas have been associated with better-quality politicians (Baltrunaite et al., 2014), measured by their level of education. In the Swedish case, the "zipper" quota requiring alternating a male and a female candidate on the party's list of candidates has increased both female representation and the competence of male politicians (Besley et al., 2017). No previous study has established a similar relationship in the business context.

A sizable literature has analyzed the relationship between female leadership and firm performance outside the context of gender quotas. It would be impossible to summarize all these studies. We emphasize two main aspects: first, it is difficult to overcome endogeneity concerns without exploiting the introduction of quotas, although some of the existing studies use instrumental variables. Second, existing results are not fully conclusive. Several studies have argued that having both men and women in top positions of a company may have positive consequences on performance. In a heterogeneous context, perspectives are enlarged, the pool of talent

[4] See also Wang and Kelan (2013).

and qualification are diversified, and shareholders are better represented (Hoogendoorn et al., 2013; Rose, 2007; Van der Walt and Ingley, 2003). The female style of leadership, including higher levels of risk aversion, may also improve performance (Bertrand, 2011). These results are, however, challenged by other studies. Adams and Ferreira (2009) find a negative impact of gender diversity on performance measures such as return on assets (ROA) and Tobin's Q, while Gagliarducci and Paserman (2015) find no evidence that female leadership is related to performance outcomes. The view itself that women are more risk-averse than men is challenged by Adams and Ragunathan (2017) and Adams and Funk (2012) when female directors are considered. Other studies qualify the conditions under which a positive relationship between women's empowerment and firms' performance may arise: the existence of a critical mass of women (Schwartz-Ziv, 2017), a positive interaction among female CEOs and women on boards Amore et al. (2014) or among female CEOs and female employees (Flabbi et al., 2019).

Non-conclusive results also emerge when looking at the relationship between women's empowerment and stock market returns. Wolfers (2006) finds no differences in stock price performance between female-headed firms and other firms. Dobbin and Jung (2010) argue that women on corporate boards are more likely to adversely affect stock prices. Ryan and Haslam (2005) find a significant increase in

share price following the appointment of a female director. However, women are more likely to be appointed in times of general financial downturn, and thus have a more precarious position (the so-called "glass cliff"). How the stock market reacts to the appointment of a female director is ambiguous: Chapple and Humphrey (2014) for Australia find no reaction, Adams et al. (2011) find a positive reaction, whereas Lee and James (2007) find a negative reaction.[5] Adams and Ferreira (2004) find that firms facing more variability in their stock returns have fewer women on their boards. Though not in connection with board gender quotas, stock price volatility has been previously explored in the diversity literature (see for example Adams and Ragunathan (2017) and (Adams and Ferreira, 2004)). Recently, Giannetti and Zhao (2019) find that boards with more ancestral diversity are associated with higher stock price volatility.

Finally, our paper also speaks to the corporate governance literature, which has underlined the importance of diversity for the quality and the functioning of the board (Dhir, 2015). An old yet unanswered question is whether the composition of the board matters for performance and firm value. Our results will suggest that quotas may be effective at increasing diversity and encouraging a better selection of board members. This is particularly important for countries such as Italy (Consob,

[5]For Italy see also Rossi and Cebula (2015), who, for a small sample of 100 Italian listed companies during the period 2012–2014, find a positive reaction within 20 days around the date of the announcement of the composition of the board.

2015), where the pre-quota situation was characterized by the presence of women almost exclusively being appointed on boards of family firms, and by a selection process not purely meritocratic (Bianco et al., 2015).[6]

2.3 The Italian law

Women are largely underrepresented in the Italian labor market: in the last ten years the labor force participation rate of Italian women has been stable, around only 48%, against a European average of 60%. In 2009 the average share of women on the boards of directors of publicly listed companies was 7%, one of the lowest in Europe. Despite this context, Italy introduced board gender quotas in July 2011 (Law 120/2011).

Figure 1 clarifies the timeline of implementation of the law, which is important to our analysis. The law was first proposed in May 2009 by a member of the Chamber of Deputies, Lella Golfo, of the center-right coalition; in November 2009, the draft was re-submitted by another member of the Chamber of Deputies, Alessia Mosca, of the center-left coalition. However, it was only two years later that the draft began being discussed thoroughly by the Italian Parliament. On March 15, 2011, the draft was approved by the Senate. The final draft of the law was approved by the

[6]This is also consistent with the descriptive evidence on the characteristics of board members after the quota provided by Solimene et al. (2017) for a selected sample of Italian firms.

Italian Parliament on June 28, 2011, by an overwhelming majority. The act came into force, after publication in the Official Gazette, on August 12, 2011. We analyze news coverage of the quota law on Lexis-Nexis and find that news related to the law are concentrated around the dates of March 15, 2011 and June 28, 2011. These are the official dates of the approvals. The process of approval of the law was not easy and before the official approval it was very unlikely to expect the law to be approved (Profeta et al., 2014). In fact, the social and political debate was very intense and strongly divided: on one side, many considered the law detrimental to the right of economic initiatives, the right of shareholders to own private property, and the principle of equality written in the Italian constitution. These criticisms were difficult to overcome and the debate was intense, both among the general public and in the Parliament. On the other side, proponents of the law built upon the raising awareness on the existence of large gender gaps in Italy and on the economic losses related to them. The awareness also increased as a reaction to severe political scandals that reinforced stereotypes on gender attitudes in Italy and clearly showed to the public opinion, both at the national and international level, that Italy was far from gender equality. However, though the issue of gender equality was clear, it seemed not obvious that board gender quotas were the appropriate way to promote it. During the law passage through the Parliament there

were several dissenting voices until the very last moment of the approval, including the associations of industries and banks, who thought they could stop the process at the Senate. However the night before March 15, when the law was scheduled for approval at the Senate, senators were bombed by thousands of emails of citizens, associations, women and men advocating the approval of the law. As a compromise between stakeholders against the law and the members of the Senate, who were under a strong pressure from citizens, on March 15 at the Senate some adjustments to the original draft were included, namely the introduction of a transition period, after the approval and before the implementation of the law. Firms, which did not anticipate the approval of the law, could use this short period to adjust their behavior.

Law 120/2011, also known as the "Golfo-Mosca" law, mandates that publicly listed companies should have a minimum target of either gender on their boards of directors and statutory auditors. The quota is implemented gradually: at the first board election, the required target is 1/5 and becomes 1/3 for the following two elections. The measure is temporary and remains in place for three consecutive board elections only.[7] If a firm does not comply, CONSOB (the regulatory body of

[7] In December 2019 the law was extended for additional 3 elections with an increase in the quota up to 40%.

the Italian stock exchange) warns the company, which has four months to comply. The warning system continues with a fine ranging from a minimum of EUR 100,000 to a maximum EUR 1,000,000. If the company persists in failing to comply without responding to the second warning within the following three months, the sanction culminates with the invalidation of the appointment of every board member. Under such an enforcement system, all companies have so far complied with the law. The law explicitly states that its effects become binding for listed firms one year after coming into force, specifically on August 12, 2012. In February 2012, the law was extended to state-owned companies, i.e. public companies under the control of the government, with immediate effect.[8]

The crucial features of the law are the following: time-limited nature, gradualism, sanctions. These features make the Italian law different from the Norwegian quota. In particular, the time-limited nature is consistent with the idea that gender quotas are a measure to "shock" and thus break the male-dominated status quo, and to lead the market to a new, more gender-balanced, equilibrium. Gradualism of the threshold (at least 1/5 of each gender at the first board election, and 1/3 at the second and third one) is based on the idea that, especially in conservative countries

[8]Around 4,000 state-owned companies must comply with the gender quota law. For them, the Department of Equal Opportunities at the Presidency of Council of Ministries is in charge of the monitoring and sanctioning system. It is however unfeasible to obtain detailed information on these companies. Thus, our analysis focuses on listed companies.

such as Italy, firms need some time to adapt to changes and it may be useful to set a not too ambitious first target.

Italian companies may choose among the following governance models: a one-tier governance system (*Monistico*); a dual-tier system with distinct supervisory (*Consiglio di sorveglianza*) and management (*Consiglio di gestione*) functions; or the traditional model with a decision-making board (*Consiglio di amministrazione*) and a separate board of statutory auditors (*Collegio sindacale*) with monitoring and control functions. In this last model, which is the one used by the majority of companies (96.2% of the companies listed on the main market in 2013), members of both boards are elected by shareholders. The two boards participate to the meetings deciding the strategy, main operations, and functioning of the firm. The board of directors has the decisive role on firm strategy. On average, the board of directors is made up of 10 members, and the board of auditors of 3 members (see Table 1). Boards of companies listed on the Italian stock exchange are elected every three years, on a date decided by the company, which is not the same for all companies, nor on a same date in a given year.[9].

In the period 2007-2014 under consideration we can classify boards in three,

[9]For more details on how companies are regulated, see Profeta et al. (2014)

almost equally distributed, cohorts: i) those changing their composition in 2007, 2010, 2013; ii) those changing in 2008, 2011, 2014; and iii) those changing in 2009 and 2012. Companies are exogenously assigned to the three cohorts: the date of renewal of the board depends on the past, well before the initial discussion of the gender quota law. In any case, we check that no firm changed the year of board election. We will highlight the division into cohorts in several parts of the analysis. As all companies are subject to the law and boards are elected every three years, with elections typically held between April and June, the first group of companies to be subject to the law for the first time had elections in 2013, the second one in 2014, and the third one in 2015.[10] However, boards with elections in 2015 had former elections in 2012. Elections in 2012 happened in the "phase-in" period of the reform: firms could endogenously adapt to the new rules, as they were not yet required to comply by the law. Therefore, throughout the analysis we will mostly focus on the first two cohorts of firms. Figure 2 clarifies the timeline of board elections for the three cohorts of firms.

[10]There were no board elections in the period August 2012-December 2012.

2.4 Data

The list of companies to which the law applies is found on the CONSOB's website. We compare this list with the one in Aida, the Italian branch of Amadeus (Bureau van Dijk), the database of comparable financial and business information on Europe's 500,000 largest public and private companies by assets. In 2013, there were 243 publicly listed firms in Italy. For each firm, we collected the election date (month and year) of the board of directors and board of statutory auditors by accessing the corporate governance page (*Relazione di Corporate Governance*) from the company's website. When this was not available, we searched on the website of the Milan Stock Exchange (*Borsa Italiana*). Alternatively, the election date was collected from the convocation notice of the shareholder meeting. Elections are typically held between April and June. For each firm, we collected from CONSOB the full names of the board members as of June 30 for every year from 2007 to 2014. Most of the time, the gender of each member was unambiguously identified through the person's first name; when the first name was ambiguous, we searched for a photo of the person.

We collected three categories of data: individual, firm-level, and stock market data. Information on the individual characteristics of board members is not available in an organized manner, and is sparse among the documents that each com-

pany must provide to CONSOB when a board member is appointed. We therefore manually collected the CVs of all members of the boards of directors and boards of statutory auditors appointed between 2007 and 2014. From our inspection of the 4,732 CVs of these individuals, we collected individual data for each member of the board.[11]

We aggregate individual characteristics at the board level and construct several board-level variables: (i) the share of women on the board, whether this share exceeds the first target of the law i.e. 20% (yes or no), the distance of this share from the threshold of 20%, and the presence of female presidents and CEOs; (ii) the share of board members with a college or a graduate degree, the share of board members with a college degree from a foreign university, the fields of study (economics, law, engineering, political science, and others) (all members, and female and male separately) and the Herfindahl index of field diversity in each board; (iii) the share of board members younger than 55 (all, and female and male separately); (iv) the percentage of board members belonging to the owner's family (all, and female and male separately); (v) the average number of board positions held by each

[11] Despite the efforts to have a complete dataset, for a limited number of boards we were not able to obtain information on all members. However, we checked that our results did not substantially change when excluding companies with more than 10% of missing values on the education variable, which was the most critical to obtain.

member in the same year (all, and female and male separately).[12] [13] [14] Table 1 presents summary statistics for these variables.[15]

We then collected firm-level data on the performance of each company. This information, again, was not immediately available. We relied, when available, on data from Orbis-AIDA (Bureau van Dijk), which we integrated with data from Bankscope on banks. In case of missing values, we hand-collected the corporate documents available on the website of the Milan Stock Exchange or on the official budget balance sheets published on each company's website. We also collected firm value measured by Tobin's Q (the ratio between a physical asset's market value and its replacement value) from Datastream. Since we will consider a one-year lag in the effect of women's participation on the boards on firm performance, performance measures are collected for the period 2011-2015. The final dataset contains the following performance information for each company for the period 2011-2015: number of employees, production (the value of production of the firm in thousands

[12] The share of board members with a college degree represents the proportion of board members who hold a college degree of any kind, namely Bachelor's degree, Master of Arts and Master of Science, MBA, or PhD. The share of board members with a graduate degree is the proportion of members with a Master's degree, MBA, or PhD.

[13] The Herfindahl index is widely used as a measure of diversity, under the expectation that higher heterogeneity is related to better performance, see Adams et al. (2011).

[14] The average number of board positions held by each member is also analyzed in the literature on Norway. Seierstad and Opsahl (2011) show that the introduction of gender quotas in Norway is associated with an increase in multiple positions, the so called 'golden skirt' phenomenon.

[15] We consider separately boards of directors and boards of statutory auditors (and the alternative forms of governance for the very few existing cases, as explained above). Table A1 presents summary statistics using individual level data.

of euros), profits (thousands of euros), share of short-term and long-term debts, ROA (return on assets, the standard indicator used to measure how profitable a company is relative to its total assets, i.e. how efficient a company's management is at using its assets to generate earnings), Tobin's Q and assets, (thousands of euros). The variables are measured at the end of December, when the budget is closed.

Sector data are also downloaded from Aida and harmonized to comply with the GICS classification of industrial sectors. We consider the following sectors: consumer discretionary, financial sector, industrials, and other sectors.[16] Table 2 presents summary statistics for firm performance. Note that sample sizes in Table 2 are slightly smaller than in Table 1 since some outcomes are missing for a few firms.

Finally, we downloaded from Bloomberg the daily closing stock price of all Italian publicly listed firms and stock market indices for the years 2011-2014.

To sum up, our final board-level dataset consists of 3,412 board-year observations over the years 2007-2014 including information on the gender composition of the board and aggregated characteristics of board members. Firm-level performance and financial data are collected for the period 2011-2015.

[16] According to the GICS classification of sectors, companies in the consumer discretionary sector include automobiles and components, consumer durables and apparel, consumer services, media, and retailing; firms in industrials include those producing capital goods and offering professional and commercial services; the financial sector includes banks and companies providing diversified financial services, insurance, and real estate. In our analysis, other sectors include energy, health care, IT, materials, telecommunication services, and utilities.

2.5 Part I. How boards of directors change

We start by analyzing the effects of the gender quota law on the characteristics of members of the board. Understanding how boards change after the quota is important to evaluate the “conventional wisdom" according to which gender quotas are associated with the appointment of less-qualified individuals. Our analysis focuses on the level of education as the main characteristic that proxies members' competence.[17] This is in line with Bianco et al. (2015) for the Italian context, Adams and Ragunathan (2017) for the U.S., and corresponds with the literature on the selection of politicians (Galasso and Nannicini, 2011).[18]

2.5.1 Methodology

In the identification of the law effects on the composition of the board we miss an appropriate control group, since in a given election year boards of all firms are subject to the law. Therefore, we need to understand how boards would have appeared in years subject to the reform had the reform not happened. We use three

[17] We do not consider CEO experience, not only because of the extremely low number of female directors and CEOs in listed companies before the law, but also because having more women in top leadership positions, and thus giving them the opportunity to acquire experience, is exactly the goal of the law. Indicators based on the evaluation of the directors' occupation (rather than of the education level) are also difficult to apply in this context, as board members do not all come from all professional backgrounds. Other indicators related to more detailed professional experience of each member would be difficult to compare.

[18] Note also that attracting better-educated people is considered an essential goal of firm strategy and one of the main reasons behind the promotion of gender equality (see OECD, 2012).

different models corresponding to different assumptions on what the composition of the boards would have been in absence of the reform. The first model assumes that the composition of the boards would have been the same; the second model assumes that the composition of the boards would have changed in a linear way, and the third one assumes that the composition of the boards would have been the one observed in the boards of that year but not re-electing. For the estimation of the three models, we use data from 2007 to 2014. While the first two models include only the cross-sections of election years, the third model includes also non-election years, thus allowing us to observe in any given year the composition of boards subject to the law and not. Whenever the dependent variable is binary, we estimate a logit model instead of a linear model.

In Figures 2.3, 2.4, and 2.5, we present visual evidence on the fitness of our assumptions on the counterfactual time trend in absence of the reform. In particular, we show the evolution of female representation, education, and age over time. The dots show the average level of the outcome for boards that were not subject to the reform. The average outcome for "treated" boards in 2013 and 2014 is represented by the triangles. While up to 2012 all boards were "untreated", in 2013 1/3 of the boards had to comply with the new rules. In 2014, the share of treated boards increased to 2/3. Note that in 2014 the group of untreated boards is represented

by firms that had board elections in the phase-in period (i.e. 2012), thus being able to endogenously adapt to the new rules.

Figure 2.3 shows that female representation increased linearly between 2007 and 2012. The jump in 2013 and 2014 clearly shows the magnitude of the change introduced by the reform. Similarly, the share of board members with at least a college degree in Figure 2.4 increased in a linear way between 2007 and 2012, whereas the time trend is flat when looking at the share of members with a graduate degree. Finally, Figure 2.5 shows that the share of board members younger than 55 largely followed a linear trend between 2007 and 2012, although the evidence is somewhat less smooth than in the previous graphs. Overall, the graphs show that a unique model may not fit the counterfactual trend for all of our outcomes of interest. Therefore, comparing the results from different assumptions may deliver more robust estimates of the reform effect.

In model 1, we assume no time trend in the evolution of the dependent variable between 2007 and 2014:

$$q_{it} = \gamma_0 + \gamma_1 \text{Reform}_{it} + \gamma_2 \text{Phase-in}_{it} + \gamma_3 \text{Members}_{it} + \phi_{\gamma,i} + \tilde{\gamma}_{it} \tag{2.1}$$

where q_{it} represents a characteristic of board i at the time t, Reform_{it} is a dummy

variable equal to 1 if the board is subject to law at time t and 0 otherwise, Phase-in$_{it}$ is a dummy equal to 1 if the board in that year of the election knows about the law but is not subject to it and 0 otherwise, Members$_{it}$ represents the number of board members.

As an example, consider q_{it} as the percentage of members with a graduate degree. Results are obtained under the assumption that in absence of the reform the percentage of members with a graduate degree would have remained the same as in the pre-reform period. Under this assumption, the comparison between boards subject and not subject to the reform identifies the causal effect of the quota law, γ_1.

In model 2 we assume a linear time trend: the outcome q_{it} would have grown in a linear way in absence of the reform.

$$q_{it} = \zeta_0 + \zeta_1 \text{Reform}_{it} + \zeta_2 \text{Phase-in}_{it} + \zeta_3 \text{Members}_{it} + \zeta_4 \cdot time + \phi_{\zeta,i} + \tilde{\zeta}_{it} \qquad (2.2)$$

The variable $time$ is equal to 1 in 2007, to 2 in 2008, and so forth. Therefore, any deviation of treated boards from the expected linear trend identifies the causal effect of the reform, ζ_1.

Model 3 does not impose a parametric assumption on the time trend: the percentage of members with a graduate degree, in any year, is given by the mean observed in that year for all boards, whose elections may have happened in the years before.

Since this specification includes all years – and not just the cross-sections of election years – we add a dummy variable indicating whether the board was elected in that year, Election_{it}:

$$q_{it} = \theta_0 + \theta_1 \text{Reform}_{it} + \theta_2 \text{Phase-in}_{it} + \theta_3 \text{Members}_{it} + \sum_{s=1}^{T} \theta_{4s} \cdot I(s = t) + \theta_5 \text{Election}_{it} + \phi_{\theta,i} + \tilde{\theta}_{it} \tag{2.3}$$

The deviation of treated boards from the mean identifies the causal effect θ_1. In all regression models we include board fixed effects $(\phi_{\gamma,i}, \phi_{\zeta,i}, \phi_{\theta,i})$. $\tilde{\gamma}_{it}$, $\tilde{\zeta}_{it}$, and $\tilde{\theta}_{it}$ are random errors, following a Type 1 Extreme Value distribution when q_{it} is a binary variable.

2.5.2 Results

Table 2.3 presents our results when we consider members of the board of directors. In Appendix Table B.0.2 we present results for boards of auditors. Column 1 shows the results under the assumption of no time trend, column 2 assumes a linear trend, and column 3 assumes the time trend described by the data. Not surprisingly, the reform is significantly associated with an increase in the share of women directors in all columns. The increase ranges between 11.41 and 16.32 percentage points. Moreover, the reform caused a significant increase in the share of women on boards

over the initial target of 20% in model 1 and 3.[19] The first model shows an increase in the share of female CEOs, although this result is not robust to the second and third specification. The share of female presidents on the board of directors significantly decreases in the third model, while female presidents of board of auditors significantly increase in all models (Table A2). Moving to our second group of outcomes – education – the reform significantly increases the share of members with a graduate degree in all models. The size of the increase ranges between 2.54 and 4.05 percentage points. Given that the average mean of the variable before the reform is 7.54, this is a large increase. The change seems to be due to women. There is however not a clear pattern for the undergraduate level of education. There is also a significant increase of board members who have studied abroad, in all models, and this is driven by women. There is no effect on the diversity of the fields of study. As for the fields of studies, we find a significant increase of members with a law degree in model 1 and 2.

Gender quotas are also associated with lower age of board members: the share of members younger than 55 years increases in the first and the third model, and this change does not seem to be gender-specific. Table 2.3 also shows that the gender quota reform is not associated with a clear significant change in the number of

[19] The law mandates to round up the number of women to be appointed to the nearest integer so as to reach 20% representation. The indicator variable *More than 20% women* is equal to one if the board appoints at least one additional woman above the mandated integer number of women.

board members with a family relationship with the ownership. More precisely, we find a reduction in the number of family-related women in the first model and an increase of men in the third model. A major concern for the introduction of a gender quota law relates to the risk of appointing non-competent women (low-educated) linked to the owner's family. The evidence seems to allay this concern.

Finally, we examine whether gender quotas increase the holding of multiple positions. This is another common concern: if quotas result in the appointment of the same woman in all boards, then they would not reach the goal of giving opportunities to all qualified, potentially eligible individuals. As result, quotas may lead to a reduction in the quality of corporate governance. In Table 2.3, we find a significant overall small decrease in the average number of positions in the first model (driven by men) and a small increase in the number of positions in the second model (driven by women). The increase of positions held by women appears also in the third specification.

When considering heterogeneous effects, results in Table 2.3 are not driven by either larger or smaller firms, which may face different constraints in the supply of qualified female members. In fact, when we consider firms above and below the median value of assets in 2012, and run separate regressions for the two subgroups, results do not differ significantly between the two subgroups.[20] As a robustness

[20]Results are available upon request.

check, we ensure that quotas are binding in most of our boards and that all effects remain if we exclude the few boards that already satisfied the required threshold in the pre-reform period (around 15% of the sample, mainly boards of auditors). All our significant effects related to women's empowerment, education, and age are even stronger (the coefficients are larger) if we only consider boards that had no women in the pre-reform period (50% of the sample), and thus had to implement more changes. The more substantial are the changes imposed by the quotas, the larger are the effects: what we observe are the consequences of a radical transformation of the status quo.[21]

2.5.3 Into the mechanism

What is the mechanism driving the observed changes to board characteristics after the implementation of gender quotas? We provide evidence that a possible mechanism lies in the selection process.

We focus on education and age, the two main variables where we have observed significant and sizeable changes after quotas. We split board members into three groups: retained, exiting, and new members.[22]

[21] Results are available upon request.

[22] We are aware that re-appointments may be constrained by factors that we do not consider (such as the number of previous appointments). These factors are, however, time-invariant, and thus should not bias our analysis.

We consider two cohorts of firms. Figure 2 shows the timeline of the implementation of the law for the two cohorts of firms. Recall that boards in Italy are elected every three years. The first cohort had elections in 2007, 2010 (before the quota) and 2013 (after the quota), and the second cohort in 2008, 2011 (before the quota) and 2014 (after the quota). Table 2.4 shows the average for each variable before and after the quota, separately for each cohort. In Panel A we compare for the first cohort of companies the last election before the quota (2010) and the first election after the quota (2013); in Panel B we compare for the second cohort the last election before the quota (2011) and the first election after the quota (2014). For the first cohort of companies (Panel A), in the pre-reform situation exiting members were more likely to have a college degree than retained ones. After the reform, instead, new members were significantly more educated than retained members, both in terms of college and graduate education, and have more college graduate than exiting members. For the second cohort of companies (Panel B) the pre-reform pattern is less clear: the average education level of retained members was not significantly different from that of exiting members, as retention is probably based on different criteria, while new members are more educated than retained and exiting. This is confirmed after the reform, with a very large difference observed between new and exiting or retained members for those with a graduate degree (men and

women). For both cohorts, new male members appointed after the reform are significantly more educated (measured by college degree in panel A and graduate degree in panel B) than both retained and exiting male members. If we compare new members appointed before and after the reform, for both cohorts new members after the reform tend to be more educated in terms of graduate education than before.

Note that, when considering graduate education, the increase in the average education level of board members after the reform is slightly stronger in boards with a lower level of male education before the reform (those with education below the median value).[23] This result is in line with gender quotas playing a role in improving the selection of board members, especially when there is more room for improvement. Our results suggest that the gender quota reform may be associated with more public scrutiny and more attention to selection.[24]

Age follows a slightly different pattern: there is evidence that new members were significantly younger than retained (Panel A and B) and exiting members (Panel B) even before the reform, a fact that is confirmed after the reform. However, the reform seems to have accelerated the process.

[23] Results are available upon request. For the other characteristics we do not detect significant heterogeneous effects in boards with a higher and lower level of male education before the reform.

[24] Other reforms associated with more public scrutiny may have a similar effect, but we focus on the causal impact of gender quotas. It is out of our scope to exclude that there may exist major board reforms different from gender quotas with a similar effect.

In the bottom panels of the table (Panel A4 and B4), we compare the characteristics of exiting males and new females after the reform. After all, since board size did not change (see Table 1), the reform required replacing male board members with female members.[25] For both cohorts, new female members are more educated than exiting male members (both college and graduate education in Panel A, only graduate education in Panel B). Female members are also significantly more likely to be younger than 55 than exiting males, suggesting that older male members were replaced by younger, more educated women.

Overall, the results confirm the patterns observed in Table 2.3: the reform increased the average level of education, for both men and women, whereas it favored the replacement of older members (i.e. older than 55) with younger ones.

2.6 Part II. The effects on performance

Do the changes to the boards induced by quotas translate into different performance? We turn to economic and financial outcomes and analyze the effects of gender quotas on firm performance. We consider board of directors, since such

[25] It could be the case that the quota law is associated with a change in the number of board members: companies may try to elude the law by reducing the number of directors on each board. Alternatively, they may increase the size of the board in order to keep all incumbent male members. This did not significantly happen in Italy, as the average size of the board remained stable over time.

body has the decisional power over firm strategy. Following the literature (Ahern and Dittmar, 2012), we consider the following measures of firm performance, as explained in Section 4: number of employees, assets, production, profits, ROA, Tobin's Q and short-term debts. We also analyze the riskiness of the company's stock as measured by the variability of stock market prices (Bloom, 2014; Adams and Ragunathan, 2017; Adams and Ferreira, 2004). This choice is motivated by the literature on uncertainty (Bloom, 2014). Stock price volatility is a common measure of firm-level uncertainty, which is known to negatively affect firm investment, production, and the ability to hire workers.

We are aware that the time span after the quota law is still limited, and a final assessment of the effects of the reform may need more time. We thus interpret our results as the short-term effect of the reform.

2.6.1 Methodology

We analyze the effect of a minimum period of time (one year) of women's participation on the boards on firm performance. For a firm subject to the law in 2013, the proportion of women on the board in 2013 is associated to firm outcomes in 2014, in order to allow one-year lag in the effects.[26] Thus, we use firm outcome data from 2011 to 2015 and data on the percentage of female directors on the board in the

[26]Results are similar if we consider simultaneous effects.

years 2010-2014. We start from a simple OLS regression that shows the correlation between the percentage of women directors on the board and measures of firm performance.

$$y_{jt} = \alpha + \beta \text{ Percentage women directors}_{j,t-1} + \gamma \text{Election}_{jt} + \tau_t + \phi_j + \epsilon_{jt} \qquad (2.4)$$

where y_{jt} is the firm's outcome, represented by the (log of) number of employees, assets, production, profits, ROA, Tobin's Q, short-term debts, and monthly stock price volatility (computed as the monthly standard deviation in the stock price) for each firm j at time t, where t goes from year 2011 to 2015 for all outcomes, except for stock price volatility which is measured monthly; Percentage women directors$_{j,t-1}$ is the proportion of women on the board of directors of firm j at time $t-1$. Election$_{jt}$ is a dummy variable for whether firm j changed its board in year t (month t for stock price volatility), and τ_t and ϕ_j represent time and firm fixed effects.

We then move to addressing the endogeneity concerns associated with women's presence on the boards. The endogenous percentage of women directors is instrumented by a dummy variable Reform$_{jt}$ equal to 1 if the board of directors in that year is subject to the quota law and 0 otherwise, and a dummy variable Phase-in$_{jt}$, which is equal to 1 if in year t the firm knows about the new rules but is not subject to it.

These instruments affect the outcome variable only through the percentage of women. Moreover, the validity of the exclusion restriction is guaranteed by the fact that the date of the reform is exogenous to firms' characteristics, and that not all firms comply with the reform at the same date.[27]

We also replace the percentage of female directors on the board with the instruments themselves, providing intention-to-treat estimates. The three models (OLS, ITT, IV) also include time dummies, and a dummy indicating whether the board election is held in that year, to avoid results being driven by other changes in that particular year rather than to the new composition of the board. We include firm fixed effects in all specifications.

[27] Given the staggered board elections across Italian companies, we can use the reform period as an instrument, rather than being forced to rely on the pre-quota percentage of female directors interacted with year dummies – which may raise endogeneity concerns – to make the instrument vary across firms, as in Ahern and Dittmar (2012) and Bertrand et al. (2019). More precisely, Ahern and Dittmar (2012) uses the distribution of female directors at the end-of-year 2002 as a measure of pre-quota variation in the number of female directors and finds detrimental effects of the law on firm performance. Eckbo et al. (2021), by using the distribution at the end-of-year 2001 (exogenous, by being observed before any announcement), does not find any significant results anymore. To provide direct comparison of our results with previous literature, we also use as an instrument the interaction between the reform period and the distance of the pre-reform share of women from the threshold of 20% (calculated as the difference between 20% and the percentage of women directors in 2010). Results, provided in Table A3, are broadly consistent with the ones presented in Table 5, although we find some positive effects on employment and assets that do not survive our specification in Table 5. Similar results as in Table A3 arise when using the distribution of the share of female directors in 2012 (when the Italian law was announced) and in 2011 (before the announcement) instead of 2010.

2.6.2 Results

Table 2.5 shows non-significant relationships between the percentage of women directors and the measures of firm performance (column 1–7) in the basic OLS regression, apart from an increase of asset. We thus do not find positive association of women with performance, nor the presence of a "glass cliff" effect (Ryan and Haslam, 2005), i.e. women more likely to be appointed in positions that are risky or precarious, which are associated with lower firm performance.

When we move to the ITT and to the instrumental variable estimation in columns 1–7, all the considered performance outcomes are not significantly affected by the proportion of women on the board. Table 2.6 shows the first stage regression of our instrumental variable strategy. The reform and phase-in indicators strongly predict a higher share of female directors on boards. With a value of 119, the F-statistic of our first stage regression is well above the conventional significance threshold.
We also consider the monthly volatility of stock prices as a relevant outcome (column 8). In this case, the OLS estimates are not significantly different from zero. However, when we move to the ITT and to the instrumental variable estimation of the impact of the share of women on stock price volatility, we find a negative and significant effect. This may be interpreted as more gender-balanced boards being perceived by the market as less risky. In fact, women are perceived as more

risk-averse than men (Jianakoplos and Bernasek, 1998; Sunden and Surette, 1998). More gender-balanced boards are not only in line with the quota requirements but they also align with one of the key objectives of current recommendations in terms of corporate governance.[28] This result is robust to the Bonferroni correction which we use to take into account the presence of multiple testing: the stock price standard deviation remains statistically significant at the 5 percent level.

Finally, we examine whether the negative effect of the share of female directors on the volatility of stock prices is driven by firms of particular industrial sectors, which may also have different financial performance independently of the presence of quotas. We perform separate regressions to find that the reduction in the monthly volatility is not driven by firms in the financial and consumer discretionary sectors, but rather by companies in industrial and other sectors. We also run separate regressions distinguishing between firms with assets above (large firms) and below the median value of assets (small firms) in every year, and find that the reduction in stock price volatility is significant in both groups.

[28] We instead do not find any significant effect of quotas on the variability of our performance measures. Results are available upon request.

2.7 Part III. Stock market reactions

As we consider listed companies, a natural way of evaluating the effects of the reform is to analyze the reaction of the stock market prices. We first focus on the stock market reactions to the announcement of the introduction of the quota law and then to the stock market reaction at the board election date. To assess the effects of the reform on the stock market returns, we consider the distance of the pre-reform share of women from the required threshold (20%) as a proxy for the magnitude of the changes to the composition of the board due to gender quotas.[29]

2.7.1 The announcement

We run an event study at the date of approval of the quota law on June 28, 2011, and at the date of the approval of the draft of the law by the Italian Senate on March 15, 2011. These two dates were chosen after checking the news coverage of the quota law on Lexis-Nexis. As explained in section 3, the Italian public opinion was confronted with the concrete possibility of the enforcement of board gender quotas for the first time on March 15, 2011. Similarly, on June 28, 2011 the final approval of the law hit the news and generated a significant debate in the political

[29] Again, we only consider boards of directors, which are formed on average by 10 members, without significant change in the board size after the reform (see Table 1). Thus, on average the quota requires at least 2 women per board, and they are 3 when the quota is exceeded. We will also control for board size.

arena.

The key focus of an event study is measuring the sample securities' mean and mean cumulative abnormal returns around the time of an event (Kothari and Warner, 2007). As the law announcement affected all firms at the same time, we estimate the abnormal returns around the event date using a procedure that allows accounting for the cross-correlation of stock returns (Eckbo et al., 2021).

First, we use a portfolio approach and check how different portfolios perform on the event dates relative to the holding period preceding the event. The holding period lasts 253 days, and we require each stock to be observed at least 100 days in order to be included in the estimation. For each stock j, we convert daily stock prices in dollar amount and calculate the stock price return on day t as the difference in the (log of the) closing stock price between t and $t-1$. We then obtain the excess return for each stock j, r_{jt}, by subtracting the daily return on 3-month US T-bills.

We estimate the following equation:

$$r_{jt} = \alpha + \beta R_{w,t} + \gamma_{AR} \cdot I(t = k) + \epsilon_{jt}, \tag{2.5}$$

where r_{jt} is the return for stock j on day t, $R_{w,t}$ is the excess return of the S&P Global index on day t relative to 3-month US T-bills, and $I(t = k)$ is an indicator variable for every day in the event window. We use the three days around the

event as our estimation window: $k \in \{-1, 0, 1\}$, where 0 represents the event day. Therefore, γ_{AR} identifies the average daily excess return on the portfolio over the event window relative to the holding period preceding the event, holding constant global stock market factors as captured by the global index.

We form portfolios of Italian firms and compare their performance around the event date. We focus on the full sample of Italian stocks, and then split stocks into two portfolios based on the distance of the company's share of female board members from the law-mandated threshold. The results are presented in Table B.0.5 in the Appendix. On June 28, the average daily excess return on the Italian portfolio is negative (but small), and not significantly different from zero. When splitting the portfolio, the return on the portfolio of less gender-diverse firms (i.e. above the median distance from the required threshold) is very similar to the return on the portfolio of more gender-diverse firms. None of the portfolio's returns is significantly different from zero, nor are the portfolios significantly different from one another. On March 15, the return on the Italian portfolio is positive. Again, the excess return of more and less gender-diverse portfolios is very similar, with no significant difference between the two groups. Therefore, using a portfolio approach we do not find convincing evidence of investors' reaction to the law on either date.

We then run a cross-sectional OLS regression of cumulative abnormal returns of

Italian stocks around the event dates, again exploiting the cross-sectional gender composition of the board. Differently from the portfolio analysis, we now allow average daily abnormal returns around the event date to be firm-specific. Therefore, we estimate Equation 2.5 separately for each stock j; as our event window includes the three days around June 28, 2011 (March 15, 2011), the cumulative abnormal returns over the event are computed as $CAR_j = 3 \cdot \hat{\gamma}_{AR,j}$. We estimate the following cross-sectional specification:

$$CAR_j = \alpha + \delta \text{ Distance from threshold}_j + \phi X_j + \epsilon_j \tag{2.6}$$

where Distance from threshold$_j$ represents the difference between the share of female directors at the announcement date and the threshold required by the law, i.e. 20%. X_j is a vector of firm-specific control variables including board size, the log of assets, and industrial sectors. Table 2.7 presents the results. The estimates are consistent with the portfolio analysis: the coefficient on the distance from the law-mandated threshold is close to zero and insignificant in both regressions. Therefore, we are unable to conclude that the quota law triggered any significant reaction from investors, neither on the day of its approval on June 28, 2011, nor during the policy debate on March 15, 2011.

2.7.2 The board elections

We now turn to stock market reactions to board elections. We have not found a significant stock market effect on the day of the approval of the law (or on March 15, during the policy discussion). Our approach is agnostic: even if investors did not react at the approval of the law, it may still be the case that the change of the boards induced by the appointment of women translate into an effect on stock market returns at the date of board election. In fact, it is only at the board election date that uncertainty on the new composition of the board is resolved. As shown by Tables 2.3 and 2.4, not only quotas led to the appointment of women directors, but also triggered a change in the composition of the entire board, where older members were replaced by younger and more educated directors.[30]

We run an event study over each board election date in the period 2011-2014, and calculate the abnormal returns for different event windows. We only focus on the election of the board of directors. For each stock j, we convert daily stock prices in dollar amount and calculate the stock price return on day t as the difference in the (log of the) closing stock price between t and $t-1$. We then obtain the excess return for each stock j, r_{jt}, by subtracting the daily return on 3-month US T-bills.

[30] Board members' characteristics may have an impact on stock market returns *per se*. We find some evidence that in absence of quotas the election of board members with higher education and lower age is associated with better returns. Results are available upon request and shown in a previous version of the manuscript.

For each stock j, we estimate abnormal returns using a single factor market model:

$$r_{jt} = \alpha_j + \beta_j R_{w,t} + \epsilon_{j,t}, \tag{2.7}$$

where $R_{w,t}$ is the excess return of the S&P Global index on day t relative to 3-month US T-bills. The firm-specific parameters α_j and β_j are estimated over the period $(-252;-11)$, where day 0 represents the election day. The abnormal return for stock j at time t is then obtained as the estimated residual from the previous regression:

$$AR_{j,t} = r_{j,t} - (\hat{\alpha}_j + \hat{\beta}_j R_{w,t}) \tag{2.8}$$

while the cumulative abnormal return $CAR_{j,(-T,T')}$ for stock j is the sum of the abnormal returns over the corresponding event window, from day $-T$ to day T'. We focus on the three, four, and five days around the board election date, and calculate cumulative abnormal returns for the windows $(-1,+1)$, $(-1,+2)$, and $(-2,+2)$.

In order to understand investors' reaction at the board election date for firms facing different requirements imposed by the law, we exploit the cross-sectional variation coming from the distance in the share of women directors from the required threshold at the board election date and the staggered introduction of quotas. We interact the distance in the share of women directors from the required threshold

with three groups of board elections. The first group includes board elections before June 2011, when the quota law had not yet been approved; the second group includes phase-in elections, namely elections that took place when the quota law was not binding. The third group includes elections that took place after the quota law became binding, that is elections in 2013 and 2014.[31] Our estimation sample includes 183 firms, for which we are able to collect the exact board election date, stock price data, firm performance, and board characteristics before and after the election, including the gender composition of the board, education, and age of board members. Out of these firms, 42 had elections before quotas, 49 in the phase-in period, and 93 after quotas.

We estimate the following equation using the sample of cumulative abnormal returns over the period 2011-2014:

$$CAR_{jt} = \alpha + \beta \text{ Dist. from threshold}_{jt} + \gamma_1 \text{ Phase-in}_{jt} + \gamma_2 \text{ After quotas}_{jt}$$
$$+\delta_1 \text{ Dist. from threshold} \times \text{Phase-in}_{jt} + \delta_2 \text{ Dist. from threshold} \times \text{After quotas}_{jt} + \phi \boldsymbol{Z}_{jt} + \epsilon_{jt} \tag{2.9}$$

Dist. from threshold$_{jt}$ is the usual variable that represents the difference between

[31] To address the possible anticipation effect, we also use the date of announcement of the lists of candidates for board membership. Results, available upon request, are very similar, although we miss some information due to more limited data availability.

the target threshold of 20% and the share of women on the board before the election. Phase-in$_{jt}$ is an indicator variable for elections in the phase-in period, and After quotas$_{jt}$ is an indicator for elections after quotas became binding, namely elections in 2013 and 2014. Z_{jt} is a vector of firm-level control variables, and ϵ_{jt} represents the error term. Our preferred specification includes firm size, ROA, and the number of elected board members. We also check whether our estimates are robust to the introduction of additional characteristics of elected board members, such as age and education. As shown in the previous sections, however, these characteristics are endogenous to the reform and likely correlate with the board-specific distance from the threshold at the election date.[32] We cluster the standard errors in the month-year of board election to account for the cross-correlation of the error term for stocks of companies that changed boards in the same time period.

Table 2.8 shows the results. For elections taking place after quotas, the coefficient on the interaction term with the distance from threshold variable is positive: a higher distance from the target threshold, i.e. a larger number of women appointed in order to comply with the quota, results in higher cumulative returns over the election period. The coefficient on the interaction term is also positive and large for elections in the phase-in period. This is due to the fact that many companies adapted to the new rules even when not required to do so by law. As

[32]Results are not shown.

shown by Table 2.8, in the estimating sample of firms the share of women elected in the phase-in period is 14% on average. Relative to the pre-reform period, the share of women on boards almost doubled in the phase-in, an increase that is much larger than what the trend would have predicted (see Figure 2.3).[33] For pre-reform elections, the omitted group in the regression, the coefficient on the distance from the target threshold is small in column (1), and flips in sign in columns (2) and (3). The point estimates in column (1) imply that increasing the distance from the target threshold by 1 percentage point increases cumulative abnormal returns by 0.0006 after quotas ($\frac{0.0043+0.0526}{100}$). Relative to an average CAR of 0.003, this effect is large (16%). It's worth noting, however, that our estimates are noisy, possibly due to the small number of observations in our sample. Changing the event window in columns (2) and (3) changes the implied CARs and the quantitative interpretation of our estimates, but the main qualitative findings in column (1) are confirmed. While we are cautious in quantitatively interpreting the magnitude of our coefficients, we feel more confident in excluding any negative effects of quotas at the election of board members. If anything, investors positively reacted to the appointment of women on boards, both in the phase-in period and when women directors were mandated by law.

[33]The difference across the two groups is significant at the 5% level.

2.8 Discussion and conclusions

We have analyzed the effects of the introduction of a gender quota law on boards of listed Italian companies according to several dimensions: the change to board characteristics, the effects on firm performance, and the stock market reaction to the announcement of the law and to board elections. We show that quotas are associated with a larger share of women directors, well above the required threshold, with higher average education of board members and a lower share of older members. These results suggest that the gender quota law significantly changed the selection process of board members. Changes may be costly, at least in the short run. However, we are able to reject the existence of a negative impact of gender quotas on economic performance, while we observe a lower variability of stock market prices. As we analyze the short-term effects of the reform, we are not able to provide a final answer on firm performance. However, we do find a non-negative reaction of investors at the board elections.[34]

Are board gender quotas improving female representation beyond the boards, i.e. generating a cascade effect from top to bottom? Bertrand et al. (2019) doubt that this happened in Norway. Future research will assess whether similar doubts apply also to the Italian case.

[34] Another unintended consequence of the Norwegian law is the delisting of companies (Bøhren and Staubo, 2014). We do not find evidence of this effect for the Italian case.

Our results for the Italian case are particularly interesting because the pre-existing causal evidence on the effects of board gender quotas is limited to Norway. We study a different context and a more gradual law reform, thus expanding the set of cases that have been studied. Italy offers a different, clear scenario which contributes substantially to our broad knowledge of the economic effects of board gender quotas. Italy and Norway are very different when looking at gender statistics: Panel A of Table 2.9 shows gender attitudes measured by the share of people who agree with specific statements posed by the European Value Survey and Panel B shows female labor force participation rates in the two countries. In these different contexts, the characteristics of board of directors in the status quo before the reform were also very different between the two countries (see Table 2.9 Panel C and D). Norway presents a higher level of education and a lower average age of board members in the period preceding the reform compared to Italy. In both cases, the level of education increases after the introduction of the quota (Ahern and Dittmar, 2012), while the change in age appears to be stronger in the Italian case.[35] In other words, in Norway, the status quo before the reform was less critical than in Italy, where it was less favorable to qualified people. More generally, as Norway is nurtured by

[35] (Ahern and Dittmar, 2012) highlight the role of previous experience as CEO, while we focus on the role of education, because the share of Italian women with CEO experience is close to zero, and thus the effect would be mechanical. We thus consider education a more interesting characteristic, as explained in Part I.

gender equality principles and Italy is far from that, consistently with our findings, the overall re-organization process induced by the reform is likely to be stronger in Italy. Note also that, although there exist specific requirements for boards of less than 10 members, the Norwegian reform imposed a jump up to 40% in female representation (existing firms have to comply by January 2008, while new firms by January 2006), while the gradualism of the target imposed by the Italian law (20% and then 33%) may have helped firms to select the more appropriate candidates and to fully exploit the beneficial effects of the re-organization process. As already noticed, if we compare directly the characteristics of new female members and exiting male members to understand the effects of the changes on the overall composition of the board, we find that the substitution between men and women increases the qualifications of board members. This result is in line with what has been found in the context of politics by Baltrunaite et al. (2014) and Besley et al. (2017).

The result that quotas may be associated with an increase in quality is a general one. Theoretically, quotas may have ambiguous effects on efficiency: on one side they are a constraint which may reduce quality if the status quo was already efficient, and on the other side they may improve quality if they force to change an inefficient status quo. Thus, as acknowledged by recent research (Bagues et al.,

2017), quotas may be desirable or not depending on the specific context. Our results show that boards of directors in Italy may be in this second situation. Our argument is supported by some interesting anecdotal evidence. During the discussion around the introduction of the Italian law, two facts rapidly became clear: first, the law had the potential to threaten the so-called "old-boys club", which dominated boards of directors prior to the introduction of the quota law, not necessarily because of their competence. Second, competent women were abundant: several lists with thousands of CVs of board-ready women were collected by women's associations, institutions, and business schools (Profeta et al., 2014). Our results suggest that gender quotas may be a policy tool to exploit abundant unused female talent.

As other countries and states, such as France, Germany, and recently California, have introduced board gender quotas, future studies will assess whether the results obtained for the Italian case are confirmed in other contexts. Further analyses are also needed to investigate whether performance will change in the long-run and whether the new selection process initiated by the introduction of gender quotas will survive when quotas, which are temporary, will be no longer in force.

2.9 Figures

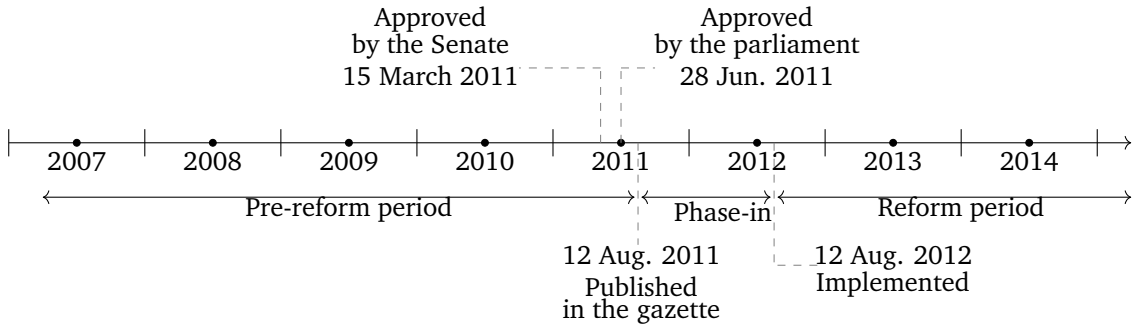

Figure 2.1: Timeline of implementation of the gender quota law

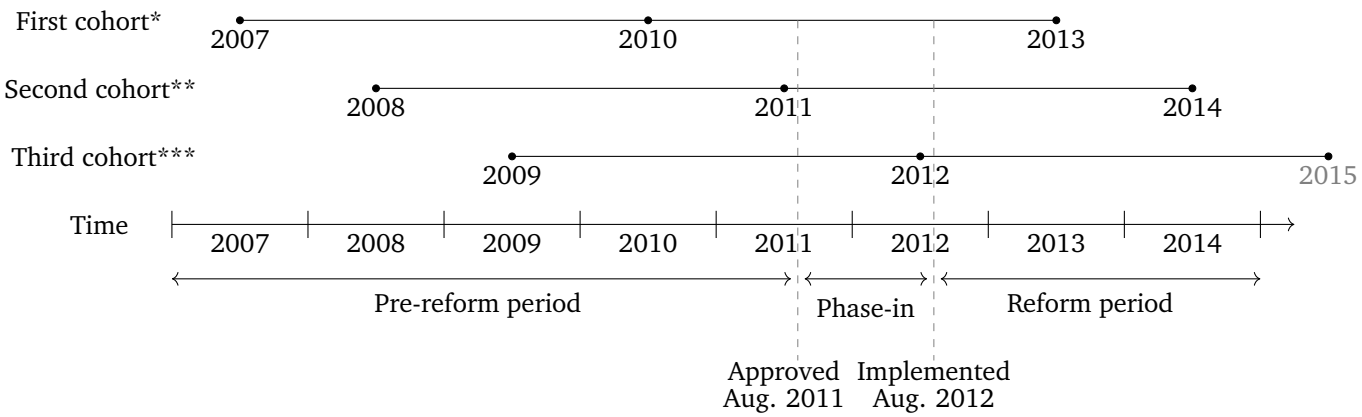

Figure 2.2: Timeline of board elections

Notes: The figure represents the timeline of the implementation of the law (on the time axis) and the timing of board elections.
*Board elections in years 2007, 2010, and 2013, from April to June. **Board elections in years 2008, 2011, and 2014, from April to June. ***Board elections in years 2009 and 2012, from April to June.

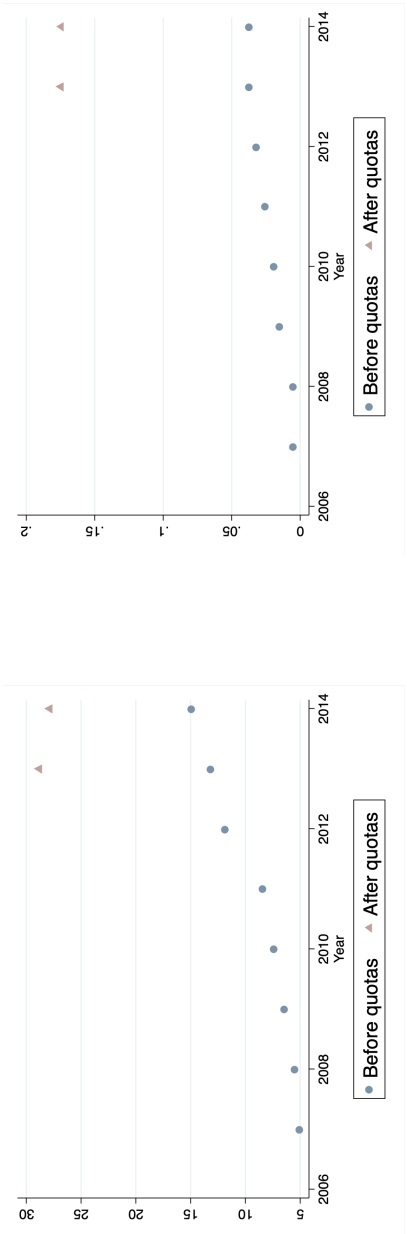

(a) Percentage of women

(b) More than 20% of women

Figure 2.3: Women's empowerment

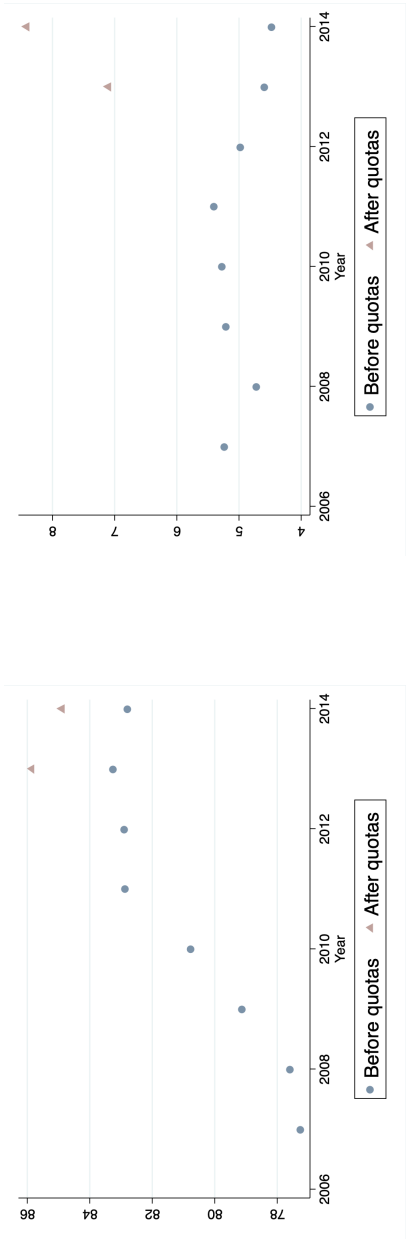

(a) % college degree

(b) % graduate degree

Figure 2.4: Education

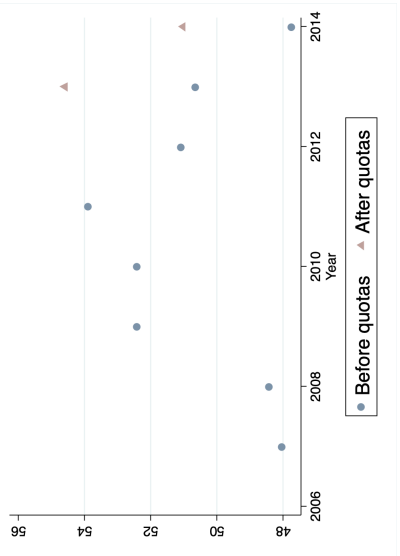

Figure 2.5: % younger than 55

2.10 Tables

Table 2.1: Summary statistics: board characteristics

		2007	2008	2009	2010	2011	2012	2013	2014
Panel A. Board of Directors									
Number of members		9.90	10.02	9.95	9.97	9.94	9.91	9.93	9.86
Percentage of women		5.90	6.15	6.84	7.65	9.42	12.46	17.83	20.97
More than 20% women		0.00	0.01	0.02	0.03	0.04	0.04	0.12	0.17
Female CEO		0.03	0.03	0.06	0.07	0.06	0.06	0.05	0.09
Fem. president		0.01	0.03	0.05	0.05	0.05	0.04	0.04	0.03
% college degree		71.38	72.33	74.43	75.62	77.90	78.87	80.27	79.87
	F	66.87	66.57	66.16	70.54	76.95	81.44	81.68	82.19
	M	71.50	72.69	75.01	76.16	78.16	78.71	79.85	79.17
% graduate degree		8.02	7.57	8.15	7.62	7.54	7.02	7.58	9.82
	F	10.64	9.69	10.70	9.57	10.84	9.38	12.34	16.49
	M	7.80	7.36	7.94	7.46	7.24	6.68	6.62	7.68
% studied abroad		4.52	4.67	4.43	4.26	4.44	5.33	5.56	6.89
	F	3.01	2.91	4.21	4.40	5.51	7.49	7.41	10.13
	M	4.46	4.56	4.28	4.07	4.28	4.81	4.82	5.34
Field diversity		0.47	0.47	0.47	0.47	0.47	0.48	0.47	0.46
% economics degree		34.39	35.57	37.05	38.59	40.41	42.24	42.22	41.85
	F	36.35	35.08	30.77	35.02	39.15	42.56	39.71	41.54
	M	34.15	35.47	37.33	38.76	40.56	42.46	42.92	41.67
% law degree		12.32	12.06	12.46	12.87	12.68	12.68	13.96	14.51
	F	10.34	7.66	10.30	10.52	11.99	14.47	17.66	19.11
	M	12.44	12.38	12.76	13.12	12.75	12.43	13.12	13.38
% younger than 55		46.59	45.53	48.57	47.75	49.76	48.08	49.03	47.25
	F	61.79	60.49	63.81	66.23	70.70	73.50	74.69	72.48
	M	45.70	44.72	47.56	46.45	47.58	45.17	43.70	40.81
% family ties		13.94	12.87	7.79	8.04	7.96	11.93	11.18	11.51
	F	37.55	34.21	19.02	17.65	15.34	19.01	14.66	13.29
	M	12.47	11.37	6.81	7.06	7.25	10.88	10.69	11.17
Number of positions		1.46	1.44	1.39	1.33	1.30	1.29	1.31	1.30
	F	1.50	1.51	1.46	1.32	1.18	1.19	1.24	1.30
	M	1.46	1.44	1.39	1.34	1.32	1.31	1.33	1.30

Continued on next page

Table 2.1: Summary statistics: board characteristics

		2007	2008	2009	2010	2011	2012	2013	2014
Retained		.	.	.	0.67	0.56	0.54	0.49	0.54
Number of boards		202	204	212	220	229	232	245	235
Panel B. Board of Auditors									
Number of members		3.03	3.06	3.11	3.12	3.10	3.09	3.11	3.12
Percentage of women		4.20	4.86	6.01	7.04	7.24	11.11	18.41	23.70
More than 20% women		0.01	0	0.01	0.01	0.01	0.02	0.04	0.05
Fem. president		0.01	0.01	0.02	0.01	0.02	0.04	0.07	0.21
% college degree		82.97	82.60	84.20	86.37	88.30	87.29	88.19	88.77
	F	75	82.76	80.56	87.80	84.09	85.32	86.69	89.83
	M	83.13	82.76	84.52	86.46	88.39	87.05	88.16	87.92
% graduate degree		2.29	1.79	2.31	2.69	3.04	2.70	2.98	3.05
	F	0.00	0.00	0.00	0.00	0.00	0.00	1.68	4.04
	M	2.64	1.79	2.47	2.86	3.20	3.02	3.50	3.07
% studied abroad		0.69	0.34	0.00	0.17	0.16	0.16	0.00	0.33
	F	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.47
	M	0.78	0.43	0.00	0.17	0.16	0.16	0.00	0.00
Field diversity		0.85	0.85	0.86	0.88	0.88	0.89	0.89	0.88
% economics degree		70.35	70.21	73	76.27	78.12	78.37	79.22	79.48
	F	62.50	62.07	66.67	75.61	75	78.86	78.43	80.39
	M	70.40	70.74	73.33	76.48	78.12	77.98	78.93	78.28
% law degree		7.80	7.68	7.37	7.32	7.84	7.14	7.52	7.48
	F	12.50	20.69	13.89	12.20	9.09	6.47	7.42	7.97
	M	7.63	7.00	6.88	6.83	7.61	6.97	7.58	7.55
% younger than 55		49.51	51.44	56.44	57.47	58.42	54.35	54.99	52.31
	F	68.75	68.97	66.67	70.73	75.00	76.37	80.25	78.68
	M	48.17	50.47	56.07	56.72	56.81	51.69	49.00	43.73
% family ties		0.69	0.34	0.33	0.33	0.32	1.11	1.33	1.49
	F	0.00	0.00	0.00	0.00	0.00	1.49	1.68	1.47
	M	0.87	0.51	0.50	0.50	0.48	1.27	1.40	1.58
Number of positions		1.57	1.47	1.39	1.36	1.33	1.26	1.29	1.29
	F	1.08	1.14	1.07	1.06	1.06	1.14	1.23	1.29
	M	1.59	1.48	1.41	1.37	1.35	1.29	1.32	1.31
Retained		.	.	.	0.68	0.66	0.59	0.49	0.47

Continued on next page

Table 2.1: Summary statistics: board characteristics

	2007	2008	2009	2010	2011	2012	2013	2014
Number of boards	192	195	200	201	208	210	226	201

Notes: Averages of board characteristics of Italian listed companies over the period 2007-2014.

Table 2.2: Summary statistics: firm characteristics

	2011	2012	2013	2014	2015
Firm characteristics:					
Employees	2,184.46	2,169.54	2,532.78	1,620.92	7,083.59
Production	696,312.46	728,536.15	692,906.69	679,326.92	4,101,473.99
Profits	-9,517.09	63,515.36	1,956.88	15,829.75	37,436.78
Assets	7,488,368.07	7,722,821.31	7,206,096.60	6,399,380.74	18,624,508.62
ROA	-0.43	-2.14	-1.90	-0.59	1.96
Short-term debt	0.64	0.63	0.70	0.67	0.44
Tobin's Q	0.90	0.76	0.78	0.85	0.98
Industrial sectors:					
Consumer	0.26	0.26	0.27	0.27	0.27
Industrial	0.17	0.17	0.17	0.17	0.18
Financials	0.24	0.25	0.24	0.24	0.24
Other	0.32	0.32	0.32	0.32	0.31
Number of firms	213	212	209	209	204

Notes: Averages of firm characteristics of Italian listed companies over the period 2011-2015.

Table 2.3: Board of Directors – Effect on board characteristics

Assumption on time trend:			No trend			Linear trend			Non-parametric		
		Mean before quotas		R^2	N		R^2	N		R^2	N
Percentage of women		9.42	16.32***	0.53	582	11.41***	0.54	582	12.70***	0.46	1779
			(0.90)			(1.77)			(1.02)		
More than 20% women		0.04	2.42***		109	1.24		109	4.25***		311
			(0.48)			(1.18)			(1.47)		
Female CEO		0.06	1.26**		59	0.94		59	0.67		200
			(0.56)			(1.27)			(1.19)		
Fem. president		0.05	-1.03		33	-0.96		33	-2.52*		109
			(0.83)			(1.39)			(1.42)		
% university degree		77.90	4.63***	0.04	582	-5.20**	0.10	582	-1.81	0.11	1779
			(1.28)			(2.47)			(1.44)		
	F	76.95	3.08	0.02	396	-10.52	0.04	396	-7.48**	0.04	1090
			(3.39)			(7.00)			(3.68)		
	M	78.16	3.65***	0.03	581	-5.68**	0.08	581	-2.05	0.10	1778
			(1.26)			(2.44)			(1.46)		
% graduate degree		7.54	4.05***	0.08	582	5.75***	0.09	582	2.54**	0.05	1779
			(1.03)			(2.05)			(1.12)		
	F	10.84	9.02***	0.09	396	9.62	0.09	396	7.72**	0.08	1090
			(3.13)			(6.55)			(3.07)		
	M	7.24	1.14	0.03	581	2.85	0.03	581	0.76	0.03	1778
			(0.94)			(1.86)			(1.06)		
% studied abroad		4.44	2.51***	0.05	582	3.60**	0.05	582	1.73**	0.02	1779
			(0.80)			(1.59)			(0.86)		
	F	5.51	5.08**	0.04	396	7.89*	0.04	396	3.28*	0.03	1090
			(1.97)			(4.12)			(1.96)		
	M	4.28	0.86	0.01	581	1.34	0.01	581	0.19	0.00	1778
			(0.67)			(1.34)			(0.79)		
Field diversity		0.47	-0.02	0.04	581	-0.01	0.05	581	-0.01	0.03	1774
			(0.01)			(0.03)			(0.02)		
% economics degree		40.41	3.18**	0.02	582	-4.73*	0.05	582	-1.95	0.06	1779

Continued on next page

Table 2.3: Board of Directors – Effect on board characteristics

Assumption on time trend:			No trend			Linear trend			Non-parametric		
		Mean before quotas		R^2	N		R^2	N		R^2	N
			(1.46)			(2.86)			(1.69)		
	F	39.15	1.13	0.05	396	-5.64	0.05	396	-4.42	0.06	1090
			(3.62)			(7.55)			(3.98)		
	M	40.56	2.11	0.01	581	-5.00*	0.04	581	-1.74	0.06	1778
			(1.51)			(2.97)			(1.74)		
% law degree		12.68	3.04***	0.04	582	0.86	0.05	582	3.25***	0.03	1779
			(0.81)			(1.61)			(0.99)		
	F	11.99	5.82**	0.05	396	-1.82	0.06	396	0.24	0.05	1090
			(2.81)			(5.83)			(3.00)		
	M	12.75	1.30	0.01	581	-0.56	0.02	581	1.81*	0.01	1778
			(0.85)			(1.69)			(1.04)		
% younger than 55		49.76	3.31**	0.04	582	-0.19	0.05	582	4.07**	0.07	1779
			(1.47)			(2.92)			(1.89)		
	F	70.70	-2.64	0.01	396	-5.65	0.01	396	0.30	0.03	1090
			(4.14)			(8.66)			(4.74)		
	M	47.58	-2.72*	0.06	581	-4.37	0.06	581	0.01	0.07	1778
			(1.53)			(3.05)			(2.00)		
% family ties		7.96	-0.46	0.15	582	0.25	0.15	582	1.95	0.11	1779
			(0.99)			(1.98)			(1.32)		
	F	15.34	-7.59**	0.10	396	-0.95	0.11	396	-0.30	0.11	1090
			(3.21)			(6.69)			(3.89)		
	M	7.25	1.00	0.14	581	0.55	0.14	581	2.55**	0.10	1778
			(0.94)			(1.87)			(1.25)		
Number of positions		1.30	-0.07**	0.03	582	0.12**	0.08	582	0.03	0.09	1779
			(0.03)			(0.05)			(0.03)		
	F	1.18	0.01	0.06	396	0.51***	0.12	396	0.19**	0.12	1092
			(0.08)			(0.17)			(0.09)		
	M	1.32	-0.08***	0.03	580	0.07	0.06	580	0.02	0.07	1776
			(0.03)			(0.06)			(0.03)		

Notes: The table shows the coefficient on the reform indicator in a regression where the dependent variable is shown in the first column. In every regression, we control for the number of board members and the phase-in period. The three specifications for every regression correspond to different assumptions on the time trend. In the first two specifications ("No trend" and "Linear trend") observations are for election years over the period 2007-2014. In the third specification, we include observations over all years between 2007 and 2014, and add an election year fixed effect. The models with dependent variable *More than 20% women*, *Female CEO*, and *Female President* are estimated using a logit model, which explains the lower number of observations. The others are estimated using a linear model. Every regression controls for board fixed effects. Standard errors are clustered at the board level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 2.4: Characteristics of retained, exiting, and new members

Panel A. 2010-2013 cohort

	Ret.	Exit.	New				Ret.	Exit.	New						
	(1)	(2)	(3)	1–2	2–3	1–3	(4)	(5)	(6)	4–5	5–6	4–6	1–4	2–5	3–6
	Before quotas: 2010						*After quotas: 2013*						*Difference 2010-2013*		
													Ret.	*Exit.*	*New*
Panel A1. All directors															
% college	82.35	91.26	91.7	***		***	87.29	88.73	93.48		**	***	*		
% graduate	6.90	8.7	6.94				6.23	8.63	10.86			***			*
% < 55 yrs	36.43	45.14	49.56	**		***	37.2	43.06	54.38		***	***			
Panel A2. Women															
% college							84.90	100	93.75			**			
% graduate							3.63	0	15.91			**			
% < 55 yrs							67.27	55.55	76.36						
Panel A3. Men															
% college							87.63	88.32	93.31		**	**			
% graduate							6.58	8.92	7.69						
% < 55 yrs							33.17	42.64	40.57	**		**			

Panel A4. Comparison: new women and exiting men

	After quotas: 2013		
	New women	Exiting men	Difference
% college	93.75	88.32	5.42**
% graduate	15.9	8.9	7.00**
% < 55 yrs	76.36	42.64	33.72***

Continued on next page

Panel B. 2011-2014 cohort

	Ret. (1)	Exit. (2)	New (3)	1–2	2–3	1–3	Ret. (4)	Exit. (5)	New (6)	4–5	5–6	4–6	1–4	2–5	3–6
	Before quotas: 2011						*After quotas: 2014*						*Difference 2010-2014*		
													Ret.	*Exit.*	*New*
Panel B1. All directors															
% college	86.86	86.58	92.5		**	**	89.12	89.94	93.47						
% graduate	3.81	7.94	8.33	**		***	5.36	6.66	25.49		***	***			***
% < 55 yrs	43.99	39.18	55.75		***	***	39.95	54.06	62.18	***	**	***		***	
Panel B2. Women															
% college							87.75	92.68	92.66						
% graduate							11.11	7.14	32.11		***	***			
% < 55 yrs							67.92	76.19	75.23	*	**	***			
Panel B3. Men															
% college							89.32	89.61	94.11						
% graduate							4.49	6.61	20.42		***	***			
% < 55 yrs							35.17	51.42	51.87	***		***			

Panel B4. Comparison: new women and exiting men

	After quotas: 2014		
	New women	Exiting men	Difference
% college	92.66	89.61	3.5
% graduate	32.11	6.61	25.4***
% < 55 yrs	75.23	51.42	23.81***

Notes: Data are averages of board members' characteristics, distinguishing between retained, exiting, and new members for each board election. Panel A shows the average characteristics of members of firms in the cohort 2007-2010-2013; Panel B shows the average characteristics of members of firms in the cohort 2008-2011-2014. *% college* indicates the share of members with at least a college degree; *% graduate* indicates the proportion of board members with a graduate degree (Master's degree, PhD, MBA); *% < 55* indicate the percentage of members younger than 55. The significance of the difference between groups 1, 2 and 3, and 4, 5 and 6 is tested in the adjacent columns. The last three columns test the difference between each group across different years. The level of significance of the difference is reported as: $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

Table 2.5: The effect of women directors on economic performance and variability of stock market prices

	(1) Log(empl.)	(2) Log(prod.)	(3) Log(profits)	(4) Log(assets)	(5) ROA	(6) Tobin's Q	(7) Short-term debt	(8) Sd(stock prices)
Panel A. OLS								
Percentage of women	-0.004	-0.007	0.002	0.004*	-0.060	0.000	-0.000	-0.002
	(0.005)	(0.007)	(0.009)	(0.002)	(0.045)	(0.002)	(0.001)	(0.002)
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
Time FE	Y	Y	Y	Y	Y	Y	Y	Y
R-squared	0.418	0.343	0.082	0.184	0.034	0.082	0.251	0.168
Panel B. ITT								
Reform	0.146	-0.097	-0.262	0.074	-0.758	-0.021	0.014	-0.002***
	(0.205)	(0.280)	(0.312)	(0.085)	(1.871)	(0.076)	(0.046)	(0.000)
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
Time FE	Y	Y	Y	Y	Y	Y	Y	Y
R-squared	0.418	0.342	0.084	0.183	0.031	0.083	0.257	0.165
Panel C. IV								
Percentage of women	0.009	-0.006	-0.009	0.008	-0.039	-0.000	0.004	-0.016***
	(0.014)	(0.018)	(0.026)	(0.006)	(0.136)	(0.006)	(0.003)	(0.003)
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
Time FE	Y	Y	Y	Y	Y	Y	Y	Y
R-squared	0.412	0.343	0.078	0.179	0.033	0.082	0.228	0.159
Mean of dep. var.	5.676	11.183	9.583	13.152	-0.674	0.847	0.613	0.021
N	874	782	566	892	899	875	776	7,662

Notes: Results of the regressions of firm performance measures (yearly observations from 2011 to 2015, column 1 to 7) and monthly standard deviation of stock prices (monthly observations, column 8) on the lagged percentage of female directors (observations from 2010 to 2014). The table shows estimates for OLS, intention to treat (ITT), and instrumental variable (IV). Year fixed effects are included for columns 1 to 7 and month fixed effects for column 8. We include dummies for election years (months in column 8). The ITT regression in Panel B includes an indicator for the phase-in period. The first stage of the IV regression is reported in Table 2.6. Standard errors are clustered at the firm level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 2.6: IV regression: first stage

	(1) Percentage of women
Reform	10.775***
	(1.379)
Phase-in	2.759**
	(1.251)
Year of election FE	Y
Firm FE	Y
Year FE	Y
Observations	1,047
Mean of dep. var.	13.381
F	119.88

Notes: First stage regression. Yearly observations between 2010 and 2014. Standard errors are clustered at the firm level. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

Table 2.7: Effect of the announcement of the quota law on cumulative abnormal returns

	(1)	(2)
	June 28, 2011	March 15, 2011
Distance from thres.	-0.0004	0.0000
	(0.0003)	(0.0005)
Board size	0.0002	0.0026**
	(0.0007)	(0.0011)
Log(assets)	0.0021	-0.0028
	(0.0015)	(0.0022)
N	188	187
R-squared	0.1211	0.0970
Mean of dep. var.	-0.0016	0.0123

Notes: Results of the event study on June 28, 2011 and March 15, 2011. Regressions are cross-sectional OLS regression of cumulative abnormal returns of Italian firms. Cumulative abnormal returns are the sum of abnormal returns over the three days surrounding the reform announcement ($(-1;+1)$ event window). *Distance from threshold* is a continuous variable constructed as 20% − percentage of women on board at the date of announcement. Board size is the number of board members. Robust standard errors in parenthesis. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

Table 2.8: Effect of the quota law on cumulative abnormal returns around the board election date

	(1)	(2)	(3)
	CARs over the event window:		
	(−1;+1)	(−2;+1)	(−2;+2)
Distance from threshold	0.0043	-0.0703+	-0.0887
	(0.0307)	(0.0460)	(0.1058)
Phase-in	-0.0094***	-0.0132**	-0.0318*
	(0.0025)	(0.0057)	(0.0168)
After quotas	-0.0050	-0.0205**	-0.0321*
	(0.0039)	(0.0084)	(0.0168)
Distance from threshold × Phase-in	0.0964**	0.1376**	0.1919+
	(0.0436)	(0.0665)	(0.1312)
Distance from threshold × After quotas	0.0526+	0.1247*	0.1602+
	(0.0380)	(0.0608)	(0.1115)
Log(assets)	-0.0006	0.0013	0.0012
	(0.0017)	(0.0019)	(0.0019)
ROA	-0.0002	-0.0003	0.0001
	(0.0005)	(0.0004)	(0.0004)
Board size	-0.0013**	-0.0022**	-0.0022**
	(0.0006)	(0.0009)	(0.0008)
N	183	183	183
R-squared	0.0481	0.0546	0.0547
Mean CAR	0.0034	0.0024	-0.0009
Mean share of women elected:			
Before quotas	0.086		
Phase-in	0.141		
After quotas	0.245		

Notes: OLS regressions of cumulative abnormal returns (CARs) of Italian listed companies on board-level variables and firm characteristics. CARs are computed over different windows around the election day. *Distance from threshold* is constructed as 20% − percentage of women on board before the board election. *Phase-in* is a dummy variable for board elections in the period from August 2011 to August 2012; *After quotas* is a dummy variable for board elections in 2013 and 2014. The omitted category are board elections before quotas (in 2011, up to June). Standard errors are clustered at the election date level (month-year). $+p<0.20$ $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

Table 2.9: Status quo: Italy and Norway

		Italy			Norway		
Panel A. Gender attitudes, 2008-2009							
% Agrees with:							
(1) "When jobs are scarce men should have more right to a job than women"		21%			3%		
(2) "A pre-school child suffers if the mother works"		75%			19%		
Panel B. Female labor force participation		2005	2010	2015	2005	2010	2015
Female LFP		38.1	38.2	39.8	68.7	68.7	68.3
Panel C. Board-level variables		2009	2010	2011	2001	2002	2003
Number of members		10.27	10.01	9.93	5.54	5.53	5.39
Share of women (%)		6.54	7.42	9.12	5.42	7.47	10.97
Higher education (%)					25.38	26.15	28.14
Graduate degree (%)		7.91	7.75	7.65			
Age		54.67	54.89	54.74	50.47	51.25	51.47
Retained from previous year			49.50	47.08		78.22	80.23
Number of positions		1.41	1.33	1.27	1.94	2.03	2.13
Panel D. Individual variables							
Age	Women	49.87	49.82	49.10	46.46	47.88	47.55
	Men	55.44	55.82	55.83	50.51	51.18	52.34
Higher education (%)	Women				25	26.67	34.15
	Men				23.80	22.66	22.83
Graduate degree (%)	Women	12.69	9.88	11.22			
	Men	7.70	7.77	7.55			
Number of positions	Women	1.41	1.30	1.19	1.08	1.22	1.22
	Men	1.48	1.41	1.34	1.18	1.18	1.21

Notes: Sources: European Value Survey 2008-2009 (Panel A), OECD Statistics (Panel B), Ahern and Dittmar (2012) (Panel C and D for Norway), and own data (Panel C and D for Italy). In Panel C and D, data for Norway are from Ahern and Dittmar (2012). Data on education are not strictly comparable. In Ahern and Dittmar (2012), higher education refers to board members with a postbaccalaureate degree, including MA, MS, MD, JD, and PhD. In our analysis, graduate degree refers to members with a Master's program, PhD, and MBA.

Chapter 3

Revisiting the Childcare Gap Between High- and Low-Skill Mothers

"It took very little time on the ground in America before I found myself becoming unrecognizable. I bought an SUV. I signed my unathletic elder daughter up for soccer. [...] I bought a small library of pre-K skill books.

I went around in a state of quiet panic."[1]

3.1 Introduction

Mothers of the recent generations are more likely to participate in the labor market, at the expense of leisure and housework time. Mothers today also spend more time with their children relative to the past. The difference is more pronounced for high-skill, high-earning mothers. The overall increase in childcare hours and the widening of the gradient by education dates back to the 1990s and has since been

[1] Excerpted from "Perfect Madness: Motherhood in an Age of Anxiety," by Judith Warner.

puzzling to economists.

In this paper, I revisit some of the evidence on the evolution of childcare time and the explanations that have been given to such changes. Two facts related to the evolution of childcare time are worth explaining. First, the overall level of childcare time has increased relative to the past, especially for the high-educated. Second, the increase has been heterogeneous for children of different ages. Whereas for cohorts of children born in the 1980s the relative focus was at older ages, for children of the recent generations childcare time is much higher at young ages. Previous research has focused on the increase in levels, but not on the relative changes over the lifecycle.

The increase in the level of childcare hours and the widening of the education gradient have been related to the increase in returns from human capital, and the comparative advantage of college-educated parents in supplying educational childcare (Ramey and Ramey, 2010). I argue that if complementarities between time and money exist, spending inequality can rationalize the increase in the level of childcare time spent on older children by the high-educated. At older ages, monetary investments are particularly productive, and complementarities with parents' time may exist: high-quality education, extracurricular activities, and enrichment goods need coordination and planning. Therefore, spending inequality in the late

1980s and early 1990s can rationalize the increase in childcare time spent on older children by high-educated parents.

I then show that neither increasing returns from human capital nor spending inequality are sufficient to explain the large increases in childcare time at very young ages documented for the cohorts born after the 1990s. I argue that information diffusion on the importance of early investments and on the complementarity of investments across the lifecycle can play an important role towards rationalizing the choice to allocate childcare time at young ages.

The paper is divided into two parts. In the first part, I establish the key facts that motivate the research question and analysis in the paper. First, I use historical time use data from the American Heritage Time Use Survey (AHTUS) to revisit some of the recent findings on the evolution of childcare time. The literature analyzes trends in parental time by calendar year (Ramey and Ramey, 2010; Guryan et al., 2008). Instead, I study how the age profile of investments changes across cohorts of children, separately for children of mothers with different educational attainment.

Taking a cohort-age approach uncovers novel evidence on the timing and the size of the relative change in investments between high- and low-skill mothers. In levels, childcare time spent by mothers of the 1970s was similar between high- and low-skill mothers. Starting with children born in the 1980s, the level of investments

increased, especially for older children of high-skill mothers. For the cohorts born after the 1980s, the shape of the education gradient over the lifecycle becomes U-shaped, with the largest differences at younger and older ages. Starting with the cohorts born in the 1990s, the gradient at young ages has been widening across cohorts.

I then show evidence on the evolution of education-related spending in high- and low-skill households. For teenage children born in the 1980s and 1990s, the level of spending in high-skill households was much higher than for the previous cohorts. The trend continued for the subsequent cohorts, reaching an annual amount of $4,600 for the cohort born in the 2000s. For young children, the trend throughout the period has been flat.

Finally, I present evidence supporting the role of information diffusion on the technology of skill formation. To document information diffusion, I construct *n-grams* from books published in the United States over the period 1960-2010. The availability of information on the importance of early child development increased fast between the mid-80s and the 1990s, in line with the timing of the increase in childcare time at very young ages.

In the second part of the paper, I introduce a two-period lifecycle model that can help understand some of the changes documented in the first part. I show that

higher returns from investing in the child's human capital and higher household spending can increase the level of time investment in every period, but does not change the optimal ratio of time investments across periods. With perfect capital markets, the optimal ratio of childcare hours across periods mostly depends on the perceived parameters of the technology of skill formation. When complementarities over the lifecycle are high, parents will tend to equalize investments at every age, as it is the case for children of high-skill mothers in the 1980s. When complementarities over the lifecycle are high, and the productivity of investments is perceived to be higher in the early childhood stage, the optimal ratio of early to late investments will tend to increase, as it is the case for high-skill mothers in the 1990s and 2000s.

Finally, I test whether differences in the perceived parameters of the technology of skill formation exist between high- and low-skill mothers. I use the Child Development Supplement of the Panel Study of Income Dynamics (PSID) to provide evidence that high- and low-skill mothers have different beliefs on the importance of early relative to late investments, and that beliefs are predictive of actual behavior. The Child Development Supplement provides information on a cohort of children born between the 1990s up to the early 2000s. I document that the difference on the perceived importance of early relative to late investments across educational levels is large: mothers with a graduate degree are 77% more likely to believe in

the long-lasting effects of early life conditions relative to high school dropouts, and have double the probability to have ever breastfed. The results provides strong support for the hypothesis that mothers of the recent generations are likely to have different beliefs on the importance of early investments, which then translate into differences in outcomes.

Related Literature Time use data show that parents today spend more time in childcare activities relative to the past, especially the college-educated (Sayer et al., 2001; Ramey and Ramey, 2010). Ramey and Ramey (2010) are the first to document the widening of the education gradient in childcare hours in the late 1990s. Guryan et al. (2008) show that the relationship between parents' education and childcare time holds within and across countries, and speculate that the positive gradient is due to the investment nature of parental time, which differs from household production and leisure time. The literature proposed a number of explanations for the overall increase in childcare hours and the widening education gradient. These include, for example, competition for college admission (Ramey and Ramey, 2010), assortative mating in the marriage market (Chiappori et al., 2017; Lundberg and Pollak, 2014), and cultural differences between parents of different socioeconomic status (Lareau, 2011). However, the determinants of the

increase in time spent with children are still little understood (Bertrand, 2018). I contribute by showing novel evidence on the evolution of childcare hours, and argue that information diffusion on the technology of skill formation can contribute towards explaining some of the increase in childcare time, especially at younger ages.

This paper is also related to the literature on the technology of human capital formation. Heckman's work on skill development challenged the traditional Beckerian view of childhood as a single investment period. Heckman builds upon evidence in psychology, education, and neuroscience to propose a multi-period model of skill formation in which human capital is formed at different stages, and there exist complementarities in investments across periods. Cunha and Heckman (2007) argue that gaps in cognitive abilities open up at very early stages in the child's development, and that the return from investing in cognitive abilities is highest at younger ages. Early investments, however, need to be followed by later investments: dynamic complementarities between development stages produce a multiplier effect, and skills and abilities beget skills and abilities.[2]

The literature has estimated the parameters of the technology of skill formation. The consensus is that parental time is more productive at young ages (Cunha et al., 2010; Del Boca et al., 2014; Agostinelli and Wiswall, 2016). What matters for

[2]See Cunha et al. (2006) for an extensive discussion.

parental investments, however, is not the "true" production function of the child's human capital, but rather what parents believe the production function to be. This paper is related to empirical work documenting the role of informational frictions between parents of different socioeconomic status. By eliciting parents' preferences through hypotetical scenarios, Boneva and Rauh (2018) show that parents' beliefs are predictive of actual investment behavior, and that parental beliefs about the skill development process differ by socieconomic status.

3.2 Motivating facts

I start by showing composition-adjusted means of weekly childcare hours, separately for high- and low-skill mothers and for children of different ages. Weekly childcare hours include time spent in non-basic care, such as playing, reading and talking to the child, or helping with homework. Figure 3.1 shows that childcare hours increased both for high- and low-educated mothers, but the increase has been more dramatic for high-educated mothers. Childcare hours increased fast for adolescent children (ages 11-17) in the late 1980s, and for younger children (ages 0-5) in the late 1990s.

Time investments over the child's lifecycle In Figure 3.2 I rearrange the

composition-adjusted means in Figure 3.1 to show the age profile of investments for children of different cohorts. Each line in Figure 3.2 represents a different cohort of children, and each dot within a line represents the child's age. Panel (a) of Figure 3.1 shows mothers at the bottom of the education distribution, and Panel (b) the top of the education distribution. For each cohort, high-educated mothers supply more childcare hours relative to low-educated mothers. Changes in childcare hours over time can be summarized by considering three groups of cohorts. For the first group, corresponding to children born in the 1970s, the age profile of investments looks similar between high- and low-educated mothers. The similarities arise both in levels and when looking at age profile over the lifecycle. The second group of children corresponds to the cohorts born between the 1980s and the early 1990s. Mothers of the cohorts born after the 1980s spend more time taking care of older children relative to the past. Such pattern is much more pronounced for high-skill mothers. For high-skill mothers, the age profile of investments looks flatter, and childcare hours equalize over the child's lifecycle. The third group of children corresponds to the cohorts born between the late 1990s and early 2000s. High-skill mothers increased substantially maternal care at young ages, whereas childcare hours increased at a low rate for low-skill mothers. With the most recent cohorts the divide in childcare hours has reached the largest size at every age in the child's

lifecycle.

Figure 3.3 shows the implied differential in childcare hours, separately for different cohorts of children and over the lifecycle of the child. The estimated differentials confirm the main patterns shown by the composition-adjusted means in Figure 3.2. Controlling for employment status, the differences between high- and low-skill mothers amplify. For children born in the 1970s, the differential is insignificant over the lifecycle. For children born in the 1980s, large and significant differences show up at older ages. For the cohorts born after the 1990s, the age profile of the education gradient becomes U-shaped, with the largest differences at younger and older ages. Starting with the cohorts born in the 1980s, the gradient has been widening across cohorts, and the intercept of the age profile has shifted upwards for the recent cohorts relative to the early cohorts.

Increase in educational spending Parental spending in educational goods and activities is largest at old ages. Figure 3.4 shows the evolution of average annual household spending separately for low- and high-skill mothers, and for children of different cohorts and ages. Average annual spending is in thousands of 2010 dollars. The measure includes education-related expenditures, such as tuition, books and supplies, tutoring and other lessons, and expenditures on enrichment

goods, such as books and magazines, musical instruments, and sports equipment.

For older children born in the 1980s and 1990s, the level of spending in high-skill households was much higher than for the previous cohorts. The trend has continued for the subsequent cohorts, reaching an annual amount of $4,600 for the cohort born in the 2000. For young children, the trend throughout the period has been flatter. Older children born in the 1980s and 1990s in low-skill households also experienced higher levels of spending relative to the older cohorts, but the amounts have been far from those spent in high-skill households. Similarly to their high-skill counterparts, for young children of low skill mothers the trend has been flat.

There are multiple reasons behind the increase in spending at older ages. First, the rise in income inequality starting from the late 1970s led to the divergence in spending possibilities between high- and low-educated households. If monetary investments are especially productive at older ages, an increase in parental resources translates into an increase in the level of spending on older children. The increase in demand for college education since the early 1990s may also have increased the productivity of monetary investments at older ages. Bound, Hershbein, and Long (2009) document that from 1992 to 2004 the number of college applicants grew by 44%, due to both increasing cohort size and rising fraction of high school grad-

uates applying for college. At the same time, the supply of college slots on the side of 4-year institutions did not keep pace with the increase in demand: undergraduate enrollment in top 20 private universities and top 20 liberal arts colleges increased by just 0.7% from 1986 to 2003 (Bound, Hershbein, and Long, 2009). The high increase in demand and the relative lack of supply is one of the reason for the large increase in university tuition costs (see Figure 3.5). High school students who intend to apply for college face pressure to signal their ability to college admission committees through performance in standardized tests, and involvement into extracurricular activities such as sports, music, volunteering, and participation to school clubs. Since the early 1990s competition for college admission became harsher, putting greater pressure on teenagers and their parents.

Complementarity between time and money For the cohorts of children born in the 1980s and 1990s, the level of spending increased at every age in high-skill households, especially for older children. If complementarities between time and money exist, the increase in educational and recreational spending in high-skill households can rationalize some of the increase in childcare hours found in time use data for older children of the 1980-1990 cohorts. Figure 3.6 shows evidence on the existence of some degree of complementarity between mothers' time and children's activities in recent years. The light line in Figure 3.6 shows the difference in

childcare hours spent on older children between high- and low-educated mothers. The darker series in Figure 3.6 shows the difference in average time spent on homework and extracurricular activities by children aged 15-17, again by the mother's education. The series for children in Figure 3.6 is obtained using more recent data from the American Time Use Survey (ATUS), which can be merged with data from the Current Population Survey (CPS) to obtain the respondent's household characteristics.[3] The differentials estimated for children are conditional on family income, suggesting that the gradient in hours devoted to homework and extracurricular activities does not merely reflect differences in income.

Information on early child development, 1960-2010 I document changes in the availability of information on early child development using *n-grams* from books published in English in the United States over the period 1960-2010. *N-grams* are often used to calculate the occurrence probability of groups of words or phrases in publications. Figure 3.7 shows the top 10 word substitutions in the middle of the words "early" and "development". The expression "early childhood development" has been the top substitution starting from the mid 1970s, but the frequency of its occurrence increased fast between the late 1980s and the early 1990s.

[3]Control variables include the child's age, sex, race, the mother's age, employment status, marital status, and family income.

Figure 3.8 shows the dependency relation between the words "skills" and "early", namely the frequency with which the word "early" modifies "skills". Examples include "early literacy skills", "early skills", "early math and literacy skills". The frequency of the dependency between "skills" and "early" increased at a constant rate from the 1970s until the 1990s. From the mid 1990s, the dependency relation started becoming more frequent, and became 8 times more present in the early 2000s relative to its level in the 1990s.

3.3 A model of the lifecycle and human capital development

I provide a framework for understanding the changes in childcare hours, and the role of increasing returns from human capital, complementarity between time and money, and information diffusion on the technology of skill formation. The mother's problem lasts for two periods, corresponding to two different stages of the child's lifecycle, early and late. Mothers get utility from consumption, and the level of well-being of their child at the end of the second period, when child development is complete. Mothers are endowed with one unit of time, which they allocate between labor hours and childcare hours. In every period, mothers choose the optimal level

of consumption and investment in their child. Mothers can borrow from the second period to finance consumption in the first period, but cannot leave negative savings to their child. When a child is born, a mother solves the problem:

$$\max_{c_1,c_2,X_1,X_2} U(c_1,c_2,X_1,X_2) = u(c_1) + \beta u(c_2) + \beta^2 V(Q)$$

subject to the intertemporal budget constraint:

$$c_1 + \frac{c_2}{1+r} + wH_1X_1 + \frac{wH_2X_2}{1+r} = \left(\tau H_1 + \frac{H_2}{1+r}\right)w \tag{3.1}$$

where c_j represents consumption and X_j represents maternal investments in each period. Because the opportunity cost of X_j is the market wage w, X_j represents childcare time. H_j represents the skill premium for the mother in each period of her life. τ is a parameter that captures the fact that when the child is young a certain fraction of the mother's time must be devoted to basic care activities. The value of the child's well-being is given by $V(Q) = \xi^e \log Q$, where Q represents the child's human capital and ξ^e the return from the child's skill. Therefore, $V(Q)$ can be thought of as the present discounted value of the child's stream of future income. The child's human capital is produced combining investments in the first and in the second period according to the production function:

$$Q = \left[\gamma X_1^{\phi} + (1-\gamma)X_2^{\phi}\right]^{\frac{1}{\phi}} \tag{3.2}$$

The CES production function allows to capture investments' self productivity, and complementarity across investments in different periods (Cunha and Heckman, 2007). Self productivity implies that investments in one period can augment investments at later periods. The parameter capturing self productivity is γ in Equation 3.2, where $0 < \gamma < 1$. The degree of complementarity across periods is represented by the parameter $\phi \in (-\infty, 1)$, capturing the extent to which early investments can be substituted by late investments, and viceversa. The first order conditions for the problem are:

$$u'(c_1) = \lambda \tag{3.3}$$

$$\beta u'(c_2) = \frac{\lambda}{1+r} \tag{3.4}$$

$$\beta^2 \xi^e Q^{-\phi} \gamma (X_1)^{\phi-1} = \lambda H_1 w \tag{3.5}$$

$$\beta^2 \xi^e Q^{-\phi} (1-\gamma)(X_2)^{\phi-1} = \frac{\lambda H_2 w}{1+r} \tag{3.6}$$

where λ corresponds to the multiplier associated with the intertemporal budget constraint. From the first order condition for consumption, it follows that setting $\beta = \frac{1}{1+r}$ implies that $c_1 = c_2$ at the optimum. From the first order condition for X_1 (or X_2), it is straightforward to see that increasing the return from the child's skill increases the mother's level of investments. Holding everything else constant, increasing ξ^e increases the marginal utility of consumption in every period. Because

consumption falls in every period, the budget constraint implies that X_1 (and X_2) must increase. An increase in the return from the child's skill, ξ_e, increases the level of investment in every period, but does not change the optimal ratio across periods. From the first order conditions of the problem, the optimal ratio of early to late investments is:

$$\frac{X_1}{X_2} = \left[\frac{\gamma}{(1-\gamma)(1+r)}\right]^{\frac{1}{1-\phi}} \tag{3.7}$$

This expression is independent from the return from the child's skill, ξ^e. With perfect capital markets, the optimal investment ratio across periods only depends on the parameters of the production function and the real rate, r.

Time and money The production function in Equation 3.2 can be augmented to allow for complementarity or substitutability between time and money at different stages of the life cycle. In particular, I set the levels of investments in every period X_j, for $j = 1,2$, equal to a composite of parental time and monetary resources:

$$X_j = \left[\omega_j t_j^{\zeta_j} + (1-\omega_j) m_j^{\zeta_j}\right]^{\frac{1}{\zeta_j}} \tag{3.8}$$

where t_1 and m_1 denote time and money in period 1, and t_2 and m_2 denote time and money in period 2. The choice variables now become the optimal amounts of

consumption, childcare hours, and monetary resources in every period. From the first order conditions, the optimal ratio between the optimal time and money in period j is:

$$\frac{t_j}{m_j} = \left(\frac{\omega_j}{w(1-\omega_j)}\right)^{\frac{1}{1-\zeta_j}} \tag{3.9}$$

Setting $\zeta_j < 0$ implies that there exists some degree of complementarity between time and money in every period: when monetary investments are high, time investments will tend to be high as well, and viceversa. From Equation 3.9, the higher the opportunity cost of maternal time w, the less the share of time resources over total investment in period j. For simplicity, I set $\zeta_1 = \zeta_2 = \zeta < 0$, so that complementarity between time and money is the same in the two periods. The optimal ratio of maternal time across periods becomes:

$$\frac{t_1}{t_2} = \left[\frac{\gamma}{(1-\gamma)(1+r)}\frac{\omega_1}{\omega_2}\right]^{\frac{1}{1-\zeta}} \left(\frac{X_1}{X_2}\right)^{\frac{\phi-\zeta}{1-\zeta}} \tag{3.10}$$

Figure 3.9 shows the optimal ratio of early to late childcare time when mothers' investments are perceived as substitutes (Panel a) or complements (Panel b) across periods, as a function of the productivity of early investments, γ. Panel (a) shows that when investments across periods are substitutes, mothers will allocate more time whenever it is perceived as more productive: the higher the productivity of

early investments γ, the higher the optimal ratio of early to late childcare time. When investments over the lifecycle are complements, instead, the optimal ratio of early to late investments goes to one, for any level of γ. When complementarities over the lifecycle are high, parents equalize investments at every age, and the productivity of early investments γ plays a less important role.

Figure 3.9 and the model provide a sufficient framework for understanding the role of complementarities between money and time and across the lifecycle towards rationalizing some of the changes documented in the first part of the paper. To make the exposition easier, in Table 3.1 I group cohorts of children into 10-year intervals, and only focus on investments at ages 0-5 ("Early") and ages 11-17 ("Late").

Table 3.1 shows the evolution of childcare hours for children of different cohorts, in the early and late stages of their lifecycles.

Substitutability across the lifecycle (1970 cohort) In the 1970s, the availability of information on the complementarity of investments across different stages of the lifecycle was low, and the wage premium for the college-educated was modest. The investment behavior of high- and low-skill mothers of children born in the 1970s was quite similar, with large differences in spending documented especially for older children. Table 3.1 shows that for children born in the 1970s, time invest-

ments are high at young ages, and low at older ages. The opposite is true for money: monetary investments are low at young ages, and high when the child is older. The pattern is consistent with the existence of some degree of substitutability between investments across the lifecycle. The existence of some degree of substitutability in investments across periods implies that mothers will intensify their investments whenever they are perceived as more productive. As shown by Figure 3.9, when $\phi > 0$ and the productivity of early investments is perceived to be high, mothers will spend more time on young children, and $\frac{t_1}{t_2} > 1$. The existence of substitutability between investments across the lifecycle and the fact that maternal time is perceived to be more productive at young ages will make mothers more willing to spend time with their children early on, and delay monetary investments to later stages. This investment behavior is consistent with a number of studies finding that time investments are more productive at younger ages, whereas monetary investments are usually more productive for older children (see for example Del Boca, Flinn, and Wiswall (2014).

Focus on older children (1980-1990 cohorts) In the period 1980–1990, the availability of information on the complementarity of investments across different stages of the lifecycle becomes more abundant. At the same time, monetary invest-

ments at older ages increase considerably, especially for high-skill parents. Complementarities between time and money in the production of late investments – X_2 in the model – push up childcare hours spent on older children (see Equation 3.9). At the same time, the greater availability of information on the complementarity of investments across the lifecycle tends to equalize time across the lifecycle. Greater availability of information on the complementarity of investments maps to a decrease in ϕ– the parameter governing the degree of complementarity between investments over the lifecycle. The Appendix shows how the optimal ratio of maternal time across the two stages of the lifecycle changes as the level of parents' investments and the degree complementarity between investments across the lifecycle increase at the same time.

Focus on young children (2000 cohort) In the late 1990s and early 2000s, the growth in income inequality keeps increasing, albeit at a lower rate. In the same time period, information diffusion on the complementarity across investments and the importance of early skill development accelerates. As shown by Figure 3.9, a further decrease in the parameter ϕ – governing degree of complementarity between investments over the lifecycle – would not suffice to rationalize the large increase in mothers' time spent on young children. The increase in mother's hours

at young ages can be rationalized through a change in the perceived values of γ and ω_1 (or ω_2), which reflect the productivity of investments at early ages. In the next section, I provide evidence corroborating the existence of different beliefs on the importance of early investments

3.4 Do high- and low-skill mothers have different beliefs on human capital formation?

I provide evidence on the existence of different beliefs on human capital formationof betweeen high- and low-skill mothers.[4] I use data of the Child Development Supplement (CDS) of the Panel Study of Income Dynamics (PSID), which contains detailed information on a sample of children and their parents. The available information includes, for example, test score data, children's time diaries, and the response to a large number of questions asked to parents. I focus on the perceived importance of early investments by analyzing the likelihood to agree to the following statement: *"The way a parent treats a child in the first four years has important life-long effects."* I also focus on differences in breastfeeding rates between mothers

[4] A similar question has been answered by Boneva and Rauh (2018) and Cunha et al. (2013), who focus on a sample of UK parents and disadvantaged African American mothers respectively. Attanasio et al. (2019) study whether parents in the UK perceive time and material investments to be complements or substitutes in the production of human capital.

of different educational attainment. Although breastfeeding represents a measure of actual behavior rather than a belief, it is suggestive of the existence of different early child-rearing attitudes between high- and low-skill mothers. Table 3.2 shows that high-skill mothers are more likely to believe in the importance of early investments, and that the gradient monotonically increases with the mother's level of education. All the regressions control for family income, so that the correlations should not be thought of as merely capturing the effect of wealth. The differences are very large: for example, mothers with a graduate degree are 77% more likely to believe in the long-lasting effects of early life conditions relative to high school dropouts, and have double the probability to have ever breastfed.

In Table 3.3 I check whether mothers' beliefs correlate with actual behavior, and regress total weekly childcare spent on very young children on the belief measure. I find that beliefs correlate positively with time spent with young children: for both measures, mothers who agree with the statement spend between 2.5 and 3 hours more with their young children relative to mothers who do not agree. When controlling for beliefs the coefficient on the mother's skill level turns insignificant. The point estimate on the indicator for mothers with a graduate degree is large, but it is imprecisely estimated. Far from capturing a causal effect, these results are nevertheless suggestive that beliefs are important predictors of actual behavior.

3.5 Conclusion

This paper revisits the evidence on the evolution of mothers' childcare time over the period 1970-2000. Whereas for the low-educated the increase in childcare time has been rather homogeneous for children of different ages, for the high-educated the change been heterogeneous at different ages. For high-educated parents of children born in the 1980s, the relative focus was at older ages. For children high-educated parents of the recent generations, childcare time is much higher at young ages. I then present evidence on contemporaneous trends that may offer candidate explanations for the change in childcare time: the increase in spending and the availability of information on the importance of early child development. Finally, I present a stylized model that helps understand which parameters may have changed in order to rationalize the observed trends.

My results show that informational frictions may result in large and persistent differences: parent's time is crucial for children's wellbeing, school performance, and intergenerational mobility (Del Boca, Flinn, and Wiswall, 2014; Daruich, 2019). My paper suggests that policies aimed at raising awareness on the importance of early life experiences among parents from less advantaged backgrounds can be an effective solution to close some of the gap in childcare time between high- and low-skill mothers.

3.6 Figures

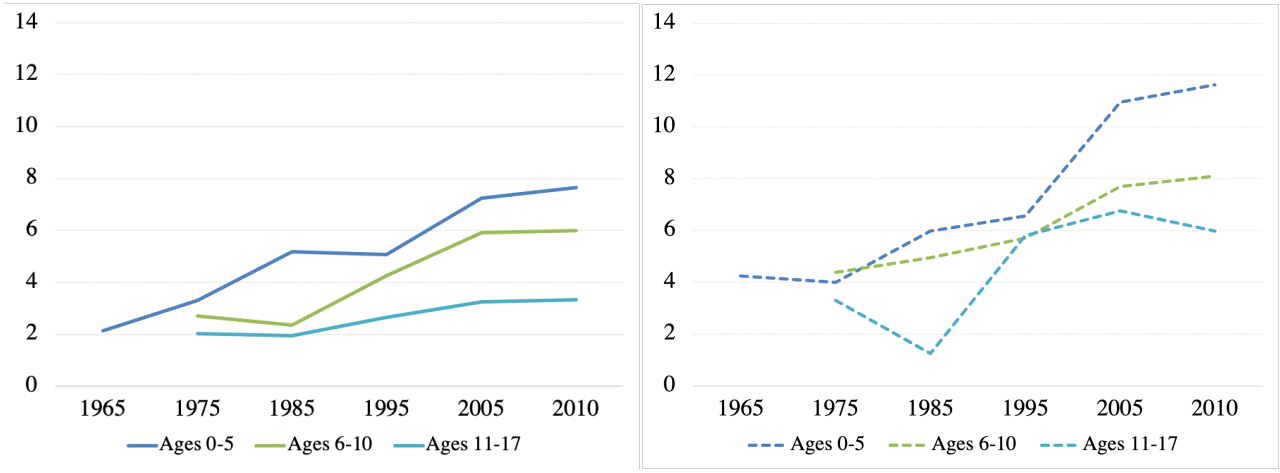

(a) Low-skill mothers (b) High-skill mothers

Figure 3.1: Childcare hours, by child's age

Notes: AHTUS, 1965-2010. Predicted average weekly childcare hours spent by mothers by year of birth of the last child in the household. Predicted values are computed by regressing weekly childcare hours separately for mothers of children of different ages (0-5, 6-10, 11-16) on dummies for the calendar year, the mother's education tercile, the mother's age category (18-24, 25-34, 35-44, 45-54, 55-64), and a full set of interactions between the calendar year and education, and the calendar year and the mother's age. Composition-adjusted means are obtained separately for children of different ages in different years using a fixed set of weights for the mother's age category.

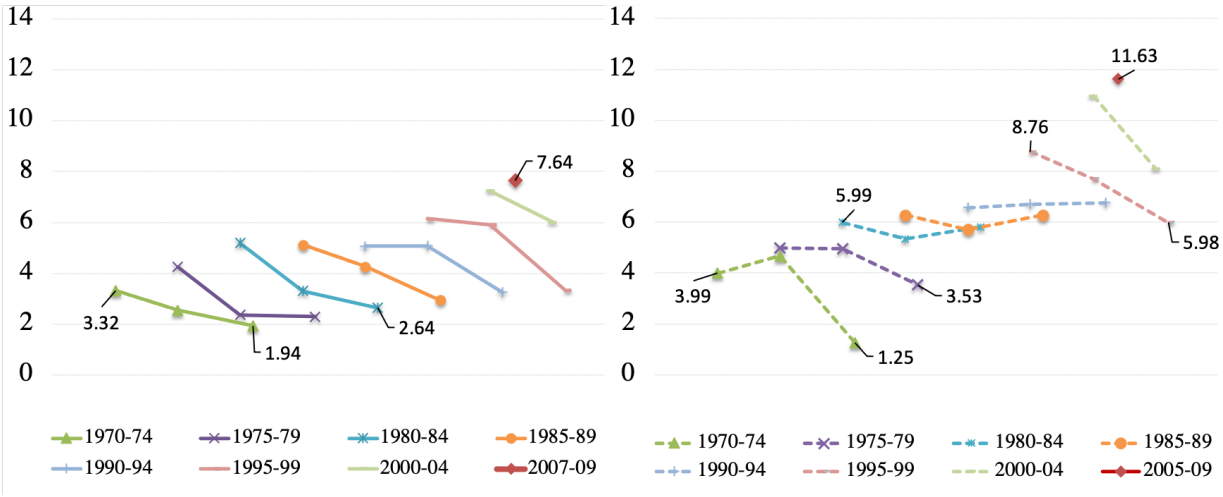

(a) Low-skill mothers (b) High-skill mothers

Figure 3.2: Childcare hours, by child's cohort

Notes: AHTUS, 1965-2010. Predicted average weekly childcare hours spent by mothers by year of birth of the last child in the household. Predicted values are computed by regressing weekly childcare hours separately for mothers of children of different ages (0-5, 6-10, 11-16) on dummies for the calendar year, the mother's education tercile, the mother's age category (18-24, 25-34, 35-44, 45-54, 55-64), and a full set of interactions between the calendar year and education, and the calendar year and the mother's age. Composition-adjusted means are obtained separately for children of different ages in different years using a fixed set of weights for the mother's age category.

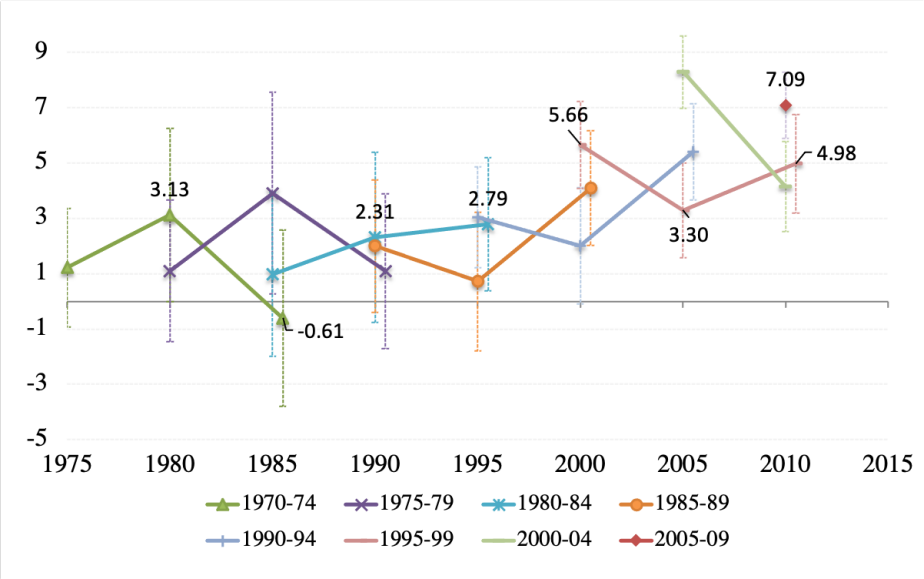

Figure 3.3: Gaps in childcare hours, high- vs low-skill mothers

Notes: AHTUS, 1965-2010. Differentials in weekly average childcare hours between mothers at the top tercile of the education distribution and mothers at the bottom tercile of the education distribution, by year of birth of the last child in the household and over the lifecycle of the child. Each dot represents a regression coefficient obtained by regressing weekly childcare hours separately for each cross section on dummies for the child's age category, dummies for the mother's tercile of the education distribution, a set of interactions between the child's age and the mother's education, and a set of controls (the mother's age category, race, employment status, marital status, number of children). The 90% confidence interval is plotted along with the coefficient. Since I do not observe every cohort at every age, missing age-year cells are obtained with a linear interpolation between adjacent age-year cells.

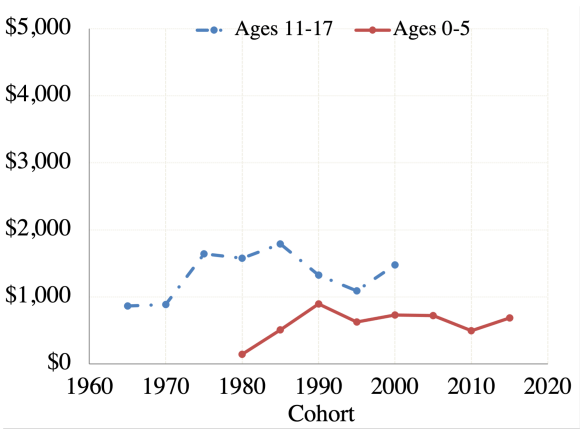

(a) Low-skill mothers

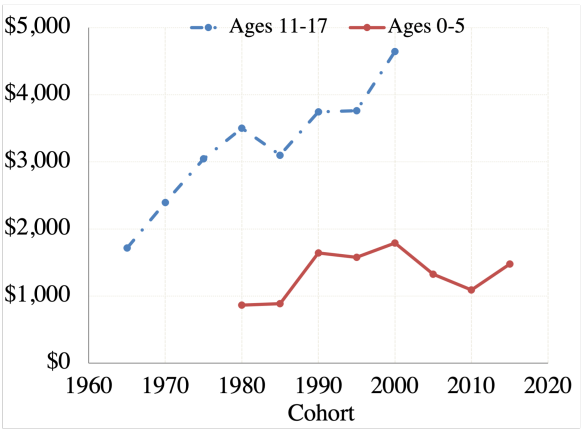

(b) High-skill mothers

Figure 3.4: Average annual expenditures ($2010)

Notes: CEX, 1980-2015. Predicted average annual expenditures spent by households in which the mother is at the bottom tercile of the education distribution (Figure a) and at the top tercile of the education distribution (Figure b), by age of the last child in the household. See the Appendix for a full description of the methodology used.

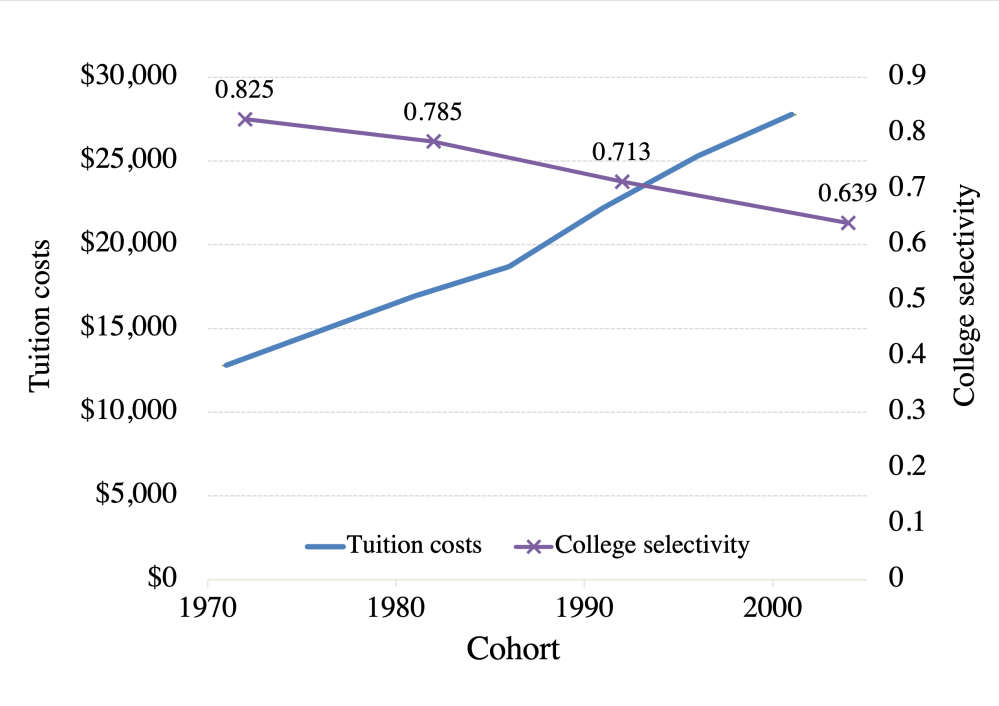

Figure 3.5: Tuition costs and college selectivity

Notes: Tuition costs in 4-year private and public institutions (National Center For Education Statistics, constant 2018-2019 dollars) and college selectivity (Bound, Hershbein, and Long, 2009). College selectivity in Bound, Hershbein, and Long (2009) is calculated as the counterfactual probability that a high school senior with the same average characteristics as in 1972 would be admitted to college in later years. The covariates used to calculate the counterfactual probability are test decile and regional indicators.

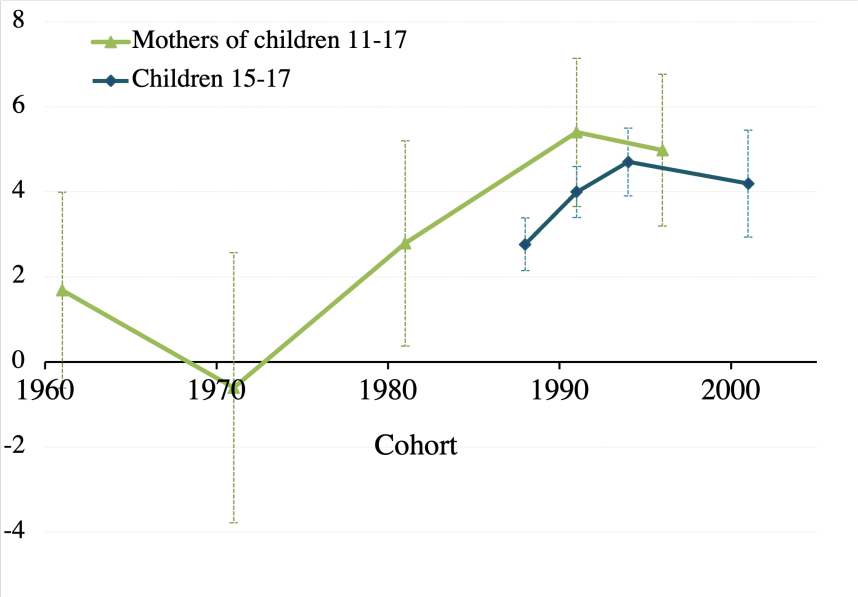

Figure 3.6: Older children: mothers' and children's time

Notes: Mothers of children 11-17: difference in childcare hours spent by mothers of children aged between 11 and 17, by skill level (high-low skill). The series is the same as the one shown in Figure 3.3 for children 11-17 (Source: AHTUS 1965-2011). Children 15-17: difference in homework and extracurricular hours spent by a sample of respondents aged 15-17, by the mother's skill level (high-low skill). Control variables include the child's age, sex, race, the mother's age, employment status, marital status, and family income (Source: ATUS 2003-2017, linked to CPS data to obtain household's characteristics).

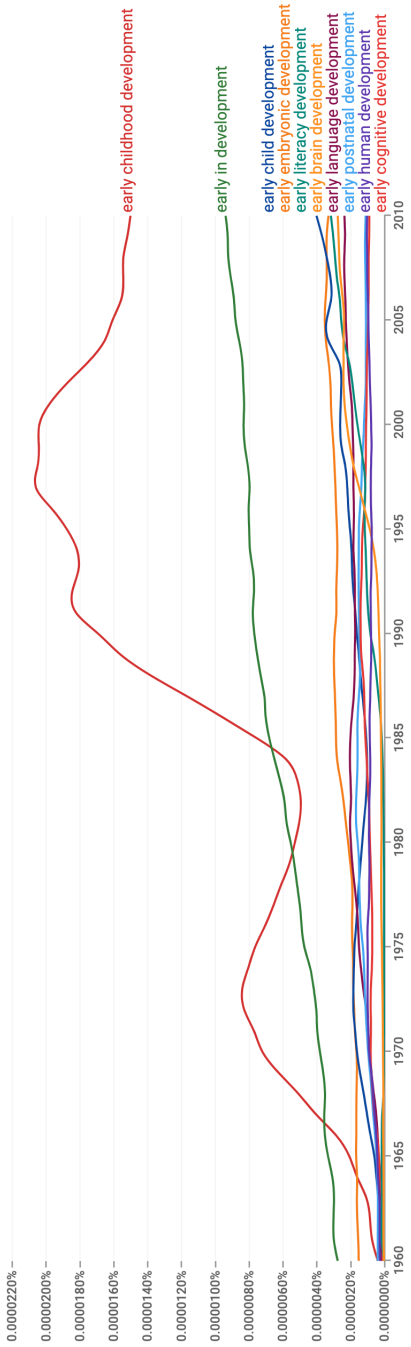

Figure 3.7: early – *word* – development

Notes: Source: Google's Books Ngram Viewer. Frequency of the top 10 substitutions for *word*. Books predominantly in the English language that were published in the United States, 1960-2010.

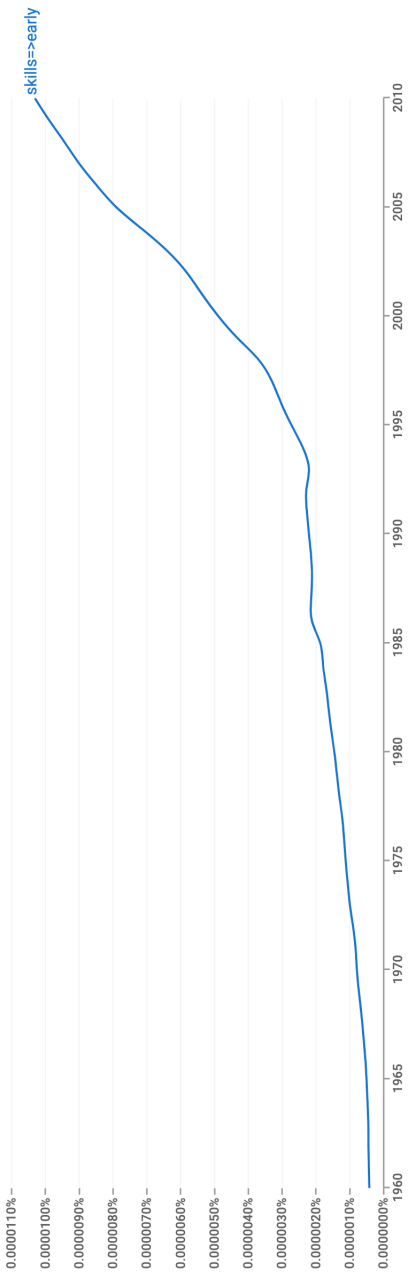

Figure 3.8: skills → early

Notes: Source: Google's Books Ngram Viewer. Dependency relation between the word *skills* and the word *early*. The series shows the frequency with which *early* modifies *skills*. These include, for example, *early literacy skills*, *early skills*, *early math and literacy skills*, and all the other instances in which the word *early* is applied to *skills*.
Books predominantly in the English language that were published in the United States, 1960-2010.

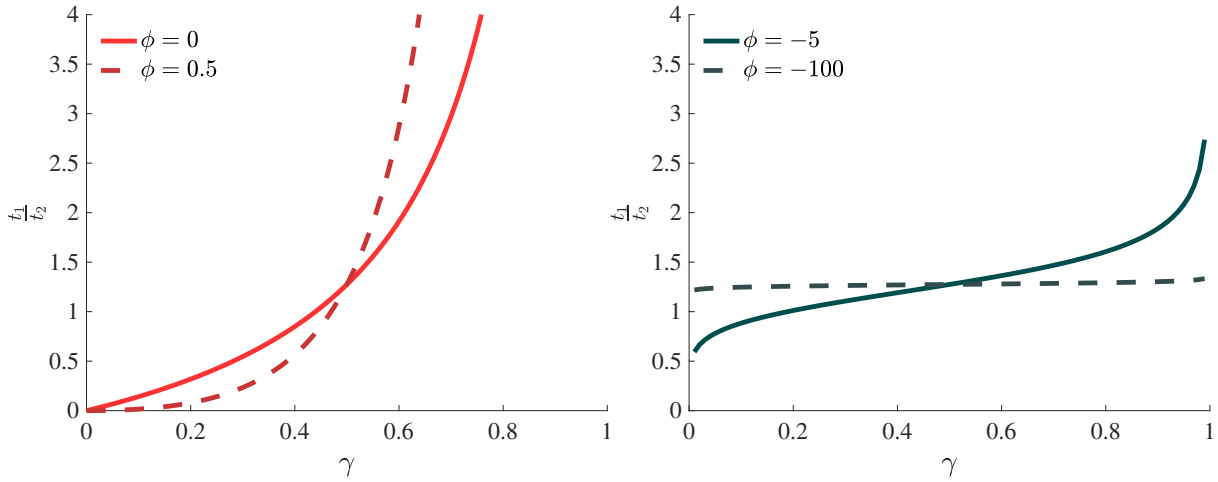

(a) Substitutability (b) Complementarity

Figure 3.9: Optimal ratio of early to late time investment, as a function of γ

Notes: $\zeta = -3$, $\omega_1 = 0.8$, $\omega_2 = 0.3$, $r = 0.01$.

3.7 Tables

Table 3.1: Ratio of early to late childcare time

	Low-skill mothers				*High-skill mothers*			
Child born:	1970	1980	1990	2000	1970	1980	1990	2000
Childcare hours:								
Early (ages 0-5)	3.78	5.15	5.60	7.44	4.49	6.13	7.66	11.29
Late (ages 11-17)	2.12	2.79	3.29	3.77	2.39	6.04	6.37	5.27
Ratio E/L	1.79	1.84	1.70	1.97	1.88	1.01	1.20	2.14

Notes: Source: AHTUS, 1965-2000. Childcare hours are the same as in Figure 3.2, aggregated in 10-year cohorts instead of 5-year cohorts. For children of ages 11-17 of the 2000 cohort, childcare hours are from the ATUS, 2003-2017.

Table 3.2: Importance of early Investments

	First 4 years have long-lasting impact:		
	Strongly Agree	Agree	Ever breastfed
	(1)	(2)	(3)
Graduate	0.109***	0.055***	0.346***
	(0.021)	(0.022)	(0.046)
Bachelor's	0.082***	0.027	0.235***
	(0.016)	(0.018)	(0.035)
Some College	0.019*	0.011	0.065***
	(0.011)	(0.011)	(0.023)
Individual-level controls	Y	Y	Y
Child's cohort FE	Y	Y	Y
Year FE	Y	Y	Y
Observations	6224	6224	2963
Mean of dep.var. in reference group	0.142	0.282	0.333

Notes: Source: Child Development Supplement of the PSID, 1997-2007. Sample: mothers of children aged under 19. Individual-level controls include the child's age, mother's education, marital status, age, number of children, employment status, and family income. Standard errors are clustered at the location-year level. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

Table 3.3: Importance of early Investments and total childcare time

Dependent Variable: Total Weekly Childcare Hours		
	(1)	(2)
First 4 years important, strongly agree	2.487***	
	(0.711)	
First 4 years important, agree		2.974***
		(0.711)
Graduate	1.863	2.081
	(1.710)	(1.700)
Bachelor's	-0.032	0.077
	(1.302)	(1.297)
Some College	0.234	0.252
	(0.864)	(0.862)
Individual-level controls	Y	Y
Child's cohort FE	Y	Y
Year FE	Y	Y
Observations	1,695	1,695
Mean of dep.var. in reference group	21.716	20.869

Notes: Source: Child Development Supplement of the PSID,1997-2007. Sample: mothers of children aged under 6. Individual-level controls include the child's age, mother's education, marital status, age, number of children, employment status, and family income. Standard errors are clustered at the location-year level. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

Appendix A

Media Focus, Executive Turnover, and Female Leadership

A.1 Model

A.1.1 Proofs of propositions

Proof of Proposition 1

Proof. Rewrite $P(S=1|q)$ and $P(S=0|q)$ using Bayes' rule:

$$P(S=1|q) = \frac{P(q|S=1)\cdot P(S=1)}{P(q)}$$

$$P(S=0|q) = \frac{P(q|S=1)\cdot P(S=0)}{P(q)}$$

Therefore:

$$\frac{P(S=1|q)}{P(S=0|q)} = \frac{P(q|S=1)}{P(q|S=0)} \cdot \frac{P(S=1)}{P(S=0)}$$

For fixed $\frac{P(S=1)}{P(S=0)}$, this implies that $\frac{P(q|S=1)}{P(q|S=0)}$ is decreasing in q, and therefore $\frac{P(q|S=0)}{P(q|S=1)}$ is increasing in q.

Denote $f_0(q)$ and $f_1(q)$ the density functions of $P(q|S=0)$ and $P(q|S=1)$. We have:

$$\frac{f_0(q_i)}{f_1(q_i)} \geq \frac{f_0(q_j)}{f_1(q_j)} \ \forall q_i \geq q_j$$

or equivalently:

$$f_0(q_i)f_1(q_j) \geq f_0(q_j)f_1(q_i) \ \forall q_i \geq q_j \tag{A.1}$$

Integrate both sides of the last expression from the minimum in the range of q to q_j, with respect to q_j:

$$\int_{\min q \in Q}^{q_j} f_0(q_i) \cdot f_1(q_j) dq_j \geq \int_{\min q \in Q}^{q_j} f_0(q_j) \cdot f_1(q_i) dq_j$$

which simplifies to:

$$\frac{f_0(q)}{f_1(q)} \geq \frac{F_0(q)}{F_1(q)} \tag{A.2}$$

Integrate both sides of equation A.1 from q_i to the maximum in the range of q, with respect to q_i:

$$\int_{q_i}^{\max q \in Q} f_0(q_i) \cdot f_1(q_j) dq_j \geq \int_{q_i}^{\max q \in Q} f_0(q_j) \cdot f_1(q_i) dq_j$$

which simplifies to:

$$\frac{1 - F_0(q)}{1 - F_1(q)} \geq \frac{f_0(q)}{f_1(q)} \tag{A.3}$$

Combine inequalities A.2 and A.3 and rearrange terms to obtain:

$$F_0(q) \leq F_1(q)$$

□

Proof of Proposition 2

Proof. W.l.o.g., assume that $E(q_t) = 0$. From Proposition (1), $P(q_t \leq q|S_t = 0) \leq P(q_t \leq q|S_t = 1)$, which implies that $E(q_t|S_t = 0) \geq E(q_t|S_t = 1)$. Therefore, since $E(q_t) = 0$:

$$E(q_t) = E(q_t|S_t = 0) \cdot P(S_t = 0) + E(q_t|S_t = 1) \cdot P(S_t = 1) = 0$$

Since $E(q_t|S_t = 0) \geq E(q_t|S_t = 1)$, it must be $E(q_t|S_t = 1) \leq 0$, which in turn implies that $E(q_t|S_t = 1) \leq E(q_t)$.

□

A.1.2 Model's predictions

Recall that an executive j is dismissed at tenure time t if $\mu_{\alpha,j}(t) < \mu^*_{\alpha,t}$, where $\mu_{\alpha,j}(t)$ is the posterior belief on the ability of executive j, and $\mu^*_{\alpha,t}$ is the endogenous firing threshold set by the board. Standard results on Bayesian updating with Gaussian distributions imply that:

$$\mu_\alpha(t) = \frac{\tau_0 \alpha_0 + \tau_q (t-1) h(t)}{\tau_\alpha + (t-1)\tau_q} \tag{A.4}$$

where α_0 is the mean prior ability of an executive, τ_0 is the prior precision, with $\tau_0 = \frac{1}{\sigma_0^2}$, and τ_q is the signal precision, with $\tau_q = \frac{1}{\sigma_q^2}$. $h(t)$ is the average value of the signal up to time t: $h(t) = \frac{\sum_{s=1}^{t-1} q_s}{t-1}$.[1] Equation A.4 shows that the posterior mean ability of an executive is an increasing function of the average realized signal up to time $t-1$: a signal above $h(t)$ increases the posterior mean at time $t+1$, whereas a signal below $h(t)$ decreases the posterior mean at time $t+1$. For simplicity, I drop the subscript α from $\mu_{\alpha,j}(t)$ and $\mu^*_{\alpha,t}$, and I set $\mu^*_t = \mu^* \forall t$. Partition the state space M_t of all possible values of $\mu(t)$ and define:

$$M_t^0 = \{\mu(t) \in M_t | \mu(t) > \mu^*\}$$

$$M_t^1 = \{\mu(t) \in M_t | \mu(t) \leq \mu^*\}$$

In every period, a signal $q_t \in Q$ is realized, where q_t is normally distributed and centered around α, the true underlying ability of an executive. Because the performance space Q is continuous, I need to define a negative news. Partition the space

[1] I ignore the complications arising from news selection, which modifies the average value of the signal up to time t to take into account publication decisions: $\hat{h}(t) = \frac{\sum_{s=1}^{t-1}(q_s|S_s=1)}{\sum_{s=1}^{t-1} I(S_s=1)}$. The posterior mean ability of an executive at time t becomes:

$$\hat{\mu}_\alpha(t) = \frac{\tau_0\alpha_0 + \tau_q \sum_{s=1}^{t-1} I(S_s=1)\hat{h}(t)}{\tau_\alpha + \sum_{s=1}^{t-1} I(S_s=1)\tau_q}$$

Because the intuition is analogous to the standard case, I provide the proof for the standard case only.

Q and define the subset Q^L of negative news such that:

$$\log\left\{\frac{P(q_t \in Q^L | \mu(t) \in M_t^0)}{P(q_t \in Q^L | \mu(t) \in M_t^1)}\right\} < 0 \tag{A.5}$$

Inequality A.5 states that the values in M_t^0 become less likely than the values in M_t^1 when $q_t \in Q^L$ is realized.

Prediction 1: *Conditional on tenure t, the probability of turnover increases as $h(t)$ decreases.*

Proof. Consider two histories $h_1(t)$ and $h_2(t)$ such that $h_1(t) < h_2(t)$. Conditional on the prior, the two posterior beliefs $\mu_1(t)$ and $\mu_2(t)$ are known at time t, and $\mu_1(t) < \mu_2(t)$. Define the log likelihood ratio $\lambda_{j,t}$ as the relative probability that an executive is retained:

$$\lambda_{j,t} = \log\frac{P(\mu_j(t) \in M_t^0)}{P(\mu_j(t) \in M_t^1)}$$

The log likelihood ratios $\lambda_{1,t}$ and $\lambda_{2,t}$ will be:

$$\lambda_{1,t} = \log\frac{P(\mu_1(t) \in M_t^0)}{P(\mu_1(t) \in M_t^1)} < \lambda_{2,t} = \log\frac{P(\mu_2(t) \in M_t^0)}{P(\mu_2(t) \in M_t^1)} \tag{A.6}$$

where inequality A.6 follows from the fact that $\mu_1(t) < \mu_2(t)$, so that $P(\mu_1(t) < \mu^*) >$

$P(\mu_2(t) < \mu^*)$. □

Prediction 2: *Conditional on tenure t, the sensitivity of turnover to the arrival of a negative news increases as $h(t)$ decreases.*

Proof. $q_t \in Q^L$ realizes at time t. Set $\mu^* < \mu_1(t) < \mu_2(t)$, so that the two executives have not yet been dismissed at time t. Using the definition of conditional probability, we have:

$$P(\mu(t) \in M_t^0 | q \in Q^L) = \frac{P(q_t \in Q^L | \mu(t) \in M_t^0) P(\mu(t) \in M_t^0)}{P(q_t \in Q^L)}$$

$$P(\mu(t) \in M_t^1 | q_t \in Q^L) = \frac{P(q_t \in Q^L | \mu(t) \in M_t^1) P(\mu(t) \in M_t^1)}{P(q_t \in Q^L)}$$

And therefore:

$$\frac{P(\mu(t) \in M_t^0 | q \in Q^L)}{P(\mu(t) \in M_t^1 | q \in Q^L)} = \frac{P(q_t \in Q^L | \mu(t) \in M_t^0)}{P(q_t \in Q^L | \mu(t) \in M_1)} \frac{P(\mu(t) \in M_t^1)}{P(\mu(t) \in M_t^1)}$$

Therefore the log likelihood ratio at time $t+1$ can be expressed as:

$$\lambda_{t+1} = \lambda_t + b_t$$

where $b_t = \log \frac{P(q_t \in Q^L | \mu(t) \in M_t^0)}{P(q_t \in Q^L | \mu(t) \in M_t^1)} < 0$ is independent of the prior belief. Using A.6, we have:

$$\lambda_{1,t+1} = \lambda_{1,t} + b_t < \lambda_{2,t+1} = \lambda_{2,t} + b_t$$

Define λ^* as the log-likelihood threshold such that an executive is fired, where $\lambda^* < 0$. Since at t the executive has not yet been fired, it must be:

$$\lambda^* < \lambda_{1,t} < \lambda_{2,t}$$

Set $\lambda_{1,t} = \lambda^* + E_1$ and $\lambda_{2,t} = \lambda^* + E_1$, with $E_1, E_2 > 0$ and $E_1 < E_2$. Then at time $t+1$, executive 1 is fired if $\lambda_{1,t+1} < \lambda^*$, which implies $E_1 < -b_t$. Executive 2 is fired if $E_2 < -b_t$. Because $E_1 < E_2$, executive 1 is more likely to be fired in period $t+1$ relative to executive 2. □

A.1.3 News selection parametrization

The distributional assumptions of the learning model and the structure imposed by the news selection rule give enough conditions to set the parameters of the distributions and produce simulations.

Definition 1 states that the publication rule is such that $\frac{P(S=1|q)}{P(S=0|q)}$ is decreasing in q.

Using Bayes' rule:

$$P(S=1|q) = \frac{P(q|S=1) \cdot P(S=1)}{P(q)}$$

$$P(S=0|q) = \frac{P(q|S=1) \cdot P(S=0)}{P(q)}$$

Setting $\omega = P(S=1)$, the two expressions imply that the odds ratio can be rewritten as:

$$\frac{P(S=1|q)}{P(S=0|q)} = \frac{P(q|S=1)}{P(q|S=0)} \cdot \frac{P(S=1)}{P(S=0)} = \frac{P(q|S=1)}{P(q|S=0)} \cdot \frac{\omega}{(1-\omega)}$$

The unconditional probability $P(q)$ is a mixture of two distributions:

$$P(q) = P(S=1) \cdot P(q|S=1) + P(S=0) \cdot P(q|S=0) = \omega \cdot P(q|S=1) + (1-\omega) \cdot P(q|S=0)$$

Under the assumption that $P(q|S=1)$ and $P(q|S=0)$ are normal distributions, then $P(q)$ is also a normal distribution. Assume that:

$$P(q|S=0) \sim N(\mu_0, \sigma_0^2)$$

$$P(q|S=1) \sim N(\mu_1, \sigma_1^2)$$

Set $\sigma_0^2 = \gamma\sigma_1^2$. Then we have:

$$\frac{P(S=1|q)}{P(S=0|q)} = \sqrt{\gamma} e^{\frac{1}{2\sigma_1}\left[\left(\frac{q-\mu_0}{\sqrt{\gamma}}\right)^2 - (q-\mu_1)^2\right]}$$

The right hand side is decreasing in q if the exponent is decreasing in q. Therefore, the following condition must be met:

$$q\left(\frac{1}{\sqrt{\gamma}}-1\right)<\mu_0-\mu_1$$

Setting $\gamma=1$, the condition is met for every q if $\mu_0-\mu_1>0$.[2] Note, moreover, that we must choose values ω, μ_0, and μ_1 such that:

$$E(q)=\mu_1\cdot\omega+\mu_0\cdot(1-\omega)$$

A.1.4 Model solution

State space At each tenure time t, the state space is represented by realized performance signals $q_1,...,q_{t-1}$ and publication decisions $S_t,...,S_{t-1}$. For private learning, the average of the signals $q_1,...,q_{t-1}$ is a sufficient statistic for past performance realizations. The statement is not true for public learning. In fact, at every point in time public beliefs are updated using the average *published* signals $q_1,...,q_{t-1}$, which depend on the realization of the sequence of random variables $S_t,...,S_{t-1}$. Keeping track of the full history of published $q_1,...,q_{t-1}$ would imply that, for a discretized

[2] Some values of μ_0,μ_1 and σ_0 may introduce kurtosis in $P(q)$. In order to avoid bimodality in $P(q)$ one must set $\mu_0-\mu_1<2\sigma_0$.

performance state of K_q points and a discretized public performance state of $K_{q|S}$ points, at each point in time the state space has dimension $K_q \times K_{q|S}^{t-1}$. To avoid such a high-dimensional state space, I simplify the problem as follows. First, I need to keep track of the history of publication decisions $S_t,...,S_{t-1}$, as the variance of posterior public beliefs depends on how many times S_t has turned on. To summarize past publications, at each point in time I calculate the average number of publications up to time $t-1$: $\bar{S}_t = \sum_{j=1}^{t-1} \frac{S_j}{t-1}$, and then discretize the interval [0,1] into K_S equally spaced points. I discretize the continuous state space of firm performance using a grid of K_q equally spaced points. Recall that the bias introduced by news selection makes the performance state look "worse": in Figure 1.7, the distribution of published events is shifted to the left relative to the true distribution. Therefore, I map the true performance space K_q to the published performance space K_q^S by re-centering K_q according to the bias introduced by news selection. The simplification I introduce implies that at each point in time the state space has dimension $K_q \times K_S$.

Turnover probability I start from time T, when learning is complete. Recall that at time T the asymptotic choice-specific value functions are:

$$V^K(\mathbf{x_T}) = E_T(\kappa_1 q_T|\mathbf{x_T}) + E_T(\kappa_2 \hat{q}_T|\mathbf{x_T}) + \delta V_{T+1}(\mathbf{x_{T+1}})|\mathbf{x_T}) + \epsilon^S = \bar{V}^K + \epsilon^K$$

$$V_T^D(\mathbf{x_T}) = -c + V_0(\mathbf{x_0}) + \epsilon^D = \bar{V}^D + \epsilon^D$$

and the optimization problem is $V(\mathbf{x}) = \max_{d\in\{0,1\}}\left(V^K(\mathbf{x}), V^Q(\mathbf{x})\right)$. The taste shocks are distributed with a Type 1 Extreme Value distribution with scale parameter τ, which has cumulative distribution function $\Lambda(x) = \frac{\exp(x)}{1+\exp(x)}$. At time T, the probability of keeping the CEO given the state variables is:

$$P(keep_T|\mathbf{x_T}) = Pr(V_T^K > V_T^D|q_1,...,q_{T-1},y_1,...,y_{T-1},S_1,...,S_{T-1}) =$$

$$P(E_T(\kappa_1 q_T|q_1,...,q_{T-1}) + E_T(\kappa_2\hat{q}_T|y_1,...,y_{T-1},S_1,...,S_{T-1}) + \delta E_T V_{T+1}(\mathbf{x_{T+1}})|\mathbf{x_T}) + \epsilon^K >$$

$$-c + V_0(\mathbf{x_0})) + \epsilon^D) =$$

$$\Lambda\left(\frac{E_T(\kappa_1 q_T|q_1,...,q_{T-1}) + E_T(\kappa_2\hat{q}_T|y_1,...,y_{T-1},S_1,...,S_{T-1}) + \delta E_T V_{T+1}(\mathbf{x_{T+1}})|\mathbf{x_T}) + c - V_0(\mathbf{x_0}))}{\tau}\right) \tag{A.7}$$

The expectations $E_T(q_T|q_1,...,q_{T-1})$ and $E_T(\hat{q}_T|q_1,...,q_{T-1},S_1,...,S_{T-1})$ can be calculated using the standard results in Bayesian inference with Gaussian distributions.

For a general period t, the probability of keeping the CEO is:

$$P(keep_t|\mathbf{x_t}) =$$

$$= \Lambda\left(\frac{E_t(\kappa_1 q_t|q_1,...,q_{t-1}) + E_T(\kappa_2\hat{q}_t|y_1,...,y_{t-1},S_1,...,S_{t-1}) + \delta E_t V_{t+1}(\mathbf{x_{t+1}}|\mathbf{x_t}) + c - V_0(\mathbf{x_0})}{\tau}\right) \tag{A.8}$$

Calculating $E_t V_{t+1}(\mathbf{x_{t+1}})|\mathbf{x_t})$ requires integrating expectations of future perfor-

mance realizations, publications, and taste shocks:

$$E_t V_{t+1}(\mathbf{x_{t+1}})|\mathbf{x_t}) = E_{S_t} E_{q_t|S_t} E_{\epsilon|q_t,S_t} V_{t+1}(\mathbf{x_{t+1}})|\mathbf{x_t}) = E_{S_t} E_{q_t|S_t} E_{\epsilon|q_t,S_t} \left(max\{\bar{V}^S_{t+1} + \epsilon^S_{t+1}, \bar{V}^D_{t+1} + \epsilon^D_{t+1}\} \right) \tag{A.9}$$

Fix the state space of past publications, summarized by $\bar{S}_t = \sum_{j=1}^{t-1} \frac{S_j}{t-1}$ as described above. Then for every grid point $k_S \in K_S$ in the publication state space:

$$\begin{aligned} \int_q \tau \log\left(\exp\left(\frac{\bar{V}^S_{t+1}}{\tau}\right) + \exp\left(\frac{\bar{V}^D_{t+1}}{\tau}\right)\right) f(q_t|q_1,...,q_{t-1}) dq_t = \\ \sum_{k_q \in K_q} \log\left(\exp\left(\frac{\bar{V}^S_{t+1}}{\tau}\right) + \exp\left(\frac{\bar{V}^D_{t+1}}{\tau}\right)\right) P(q_t^{k_q}|q_1,...,q_{t-1}) \end{aligned} \tag{A.10}$$

Assuming a taste shock with Type 1 Extreme Value distribution allows having a closed form for the expectation in Equation A.9. Note that going from equation A.9 to A.10 for a given publication state requires that $E_{q_t|S_t}(\cdot) = E_{q_t}(\cdot)$, which follows from the assumption that the Board of Directors does not learn CEO quality through publications, and therefore the Board's expectation of firm performance is independent of past publications. However, the expected value of the current CEO, $E_t V_{t+1}$ in Equation A.9, depends on the publication state k_s, because public beliefs affect the value of the current firm-CEO match.

Transition probabilities The expression $P(q_t^{k_q}|q_1,...,q_{t-1})$ in Equation A.10 represents the Board's perceived probability of the CEO realizing performance k_q at

time t, given past performance $q_1, ..., q_{t-1}$.

$$P(q_t^{k_q}|q_1,...,q_{t-1}) = \Phi\left(\frac{q_t^{k_q} + 0.5 \times kstep - E(q_t|q_1,...,q_{t-1})}{\sqrt{\Omega_{t-1}}}\right) - \Phi\left(\frac{q_t^{k_q} - 0.5 \times kstep - E(q_t|q_1,...,q_{t-1})}{\sqrt{\Omega_{t-1}}}\right) \tag{A.11}$$

where $\Omega_{t-1} = \left(\tau_\alpha + (t-1)\tau_{\tilde{q}}\right)^{-1}$, $kstep$ is the distance between grid points, and $q_t^{k_q}$ is the value of firm performance at grid point k_q. Since past performance realizations are summarized by the average realized performance up to $t-1$, and I have discretized the performance state, I use the transition probability matrix of average performance moving from grid point $k_{j'}$ at $t-1$ to point k_j at time t:

$$P(\bar{q}_t^{k_{j'}}|\bar{q}_{t-1}^{k_j}) = \Phi\left(\frac{t\cdot(\bar{q}_t^{k_j} + 0.5 \times kstep) - (t-1)\bar{q}_{t-1} - E(q_t|q_1,...,q_{t-1})}{\sqrt{\Omega_{t-1}}}\right) - \Phi\left(\frac{t\cdot(\bar{q}_t^{k_j} - 0.5 \times kstep) - (t-1)\bar{q}_{t-1} - E(q_t|q_1,...,q_{t-1})}{\sqrt{\Omega_{t-1}}}\right) \tag{A.12}$$

Model solution I use value function iteration to solve the dynamic programming problem numerically. The algorithm is similar to Rust (1987).

I guess a value for V_0, that is the value from hiring a CEO:

1. I start from time T and solve for the asymptotic value functions V^K and V^D using value function iteration. I set $T = 130$.

2. I use backwards recursion to solve for the choice-specific value functions V_t^K and V_t^D at every $t = 1, ..., T$.

3. I obtain V_0.

I iterate steps 1–3 and stop at the i-th iteration whenever $|V_0^i - V_0^{i-1}| < 10^{-15}$.

A.2 Additional Figures

Feb 2, 2005
Dow Jones News Service

CITY OF INDUSTRY, Calif. (Dow Jones) -- Hot Topic Inc.'s (HOTT) same-store sales fell 2.5% in January, as a shift in the timing of holidays led to weaker comparisons in the last two weeks of the month. [...]
The teen apparel and music retailer said it expects to meet analysts' mean estimate for earnings of 38 cents a share in the fourth quarter. [...]

Entity id	Relevance	Source	Date	Time	Story group	Sentiment
HOTT	100	DJNS	2feb2006	10:59 AM	revenues	47

(a) Negative news

Oct 8, 2003
Market News Publishing

CITY OF INDUSTRY, Calif. -- [...] California-based Hot Topic said that comparable store sales for the period increased 9.0% from fiscal September 2002. The company added that net sales for the month of September 2003 increased 31% to $48.9 million from net sales of $37.4 million posted in the year ago period. [...]

Entity id	Relevance	Source	Date	Time	Story group	Sentiment
HOTT	100	DJNS	8oct2006	9:15 AM	earnings	83

(b) Positive news

Figure A.2.1: News examples

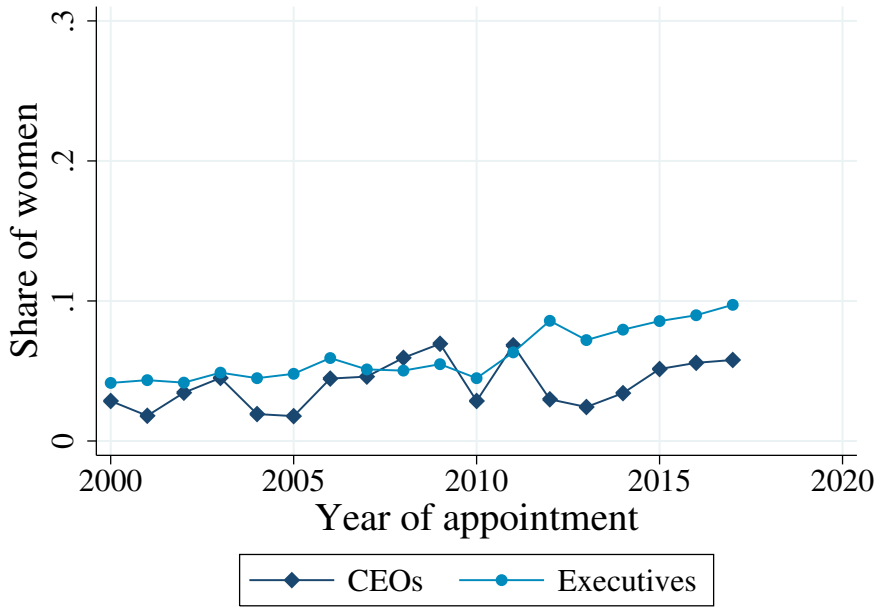

Figure A.2.2: Share of female executives, by year of appointment

Notes: Executives include Chairs, CEOs, Presidents, CFOs, COOs, and other Chief Officers.

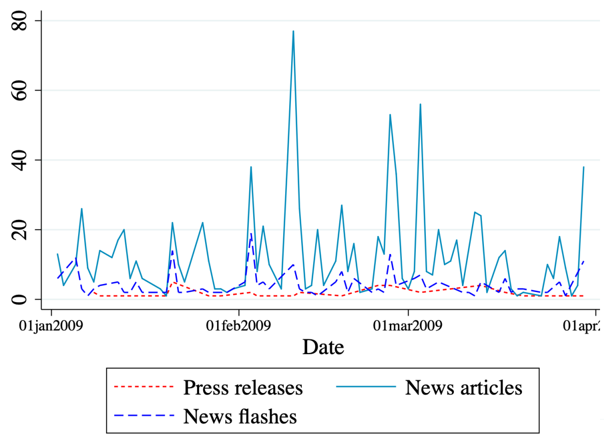

(a) Press releases and news

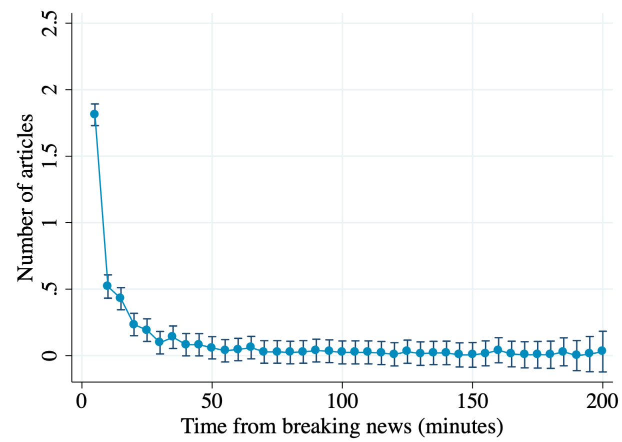

(b) News articles after breaking news

Figure A.2.3: News data

Notes: General Motors, period January-March 2009. (a) Number of press releases, full articles, and news flashes released over the period January-March 2009. (b) Average number of full articles and news flashes released after the first news release.

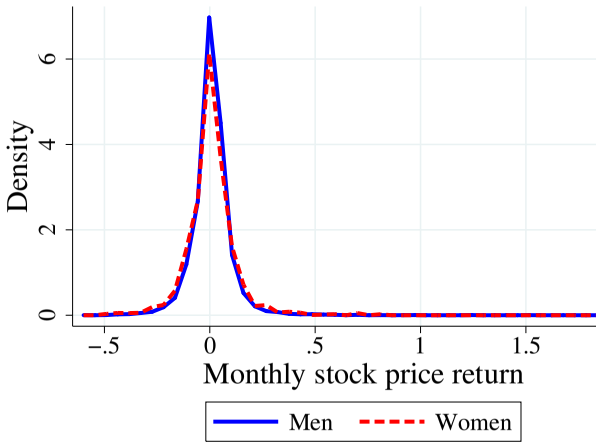

(a) Monthly stock price return

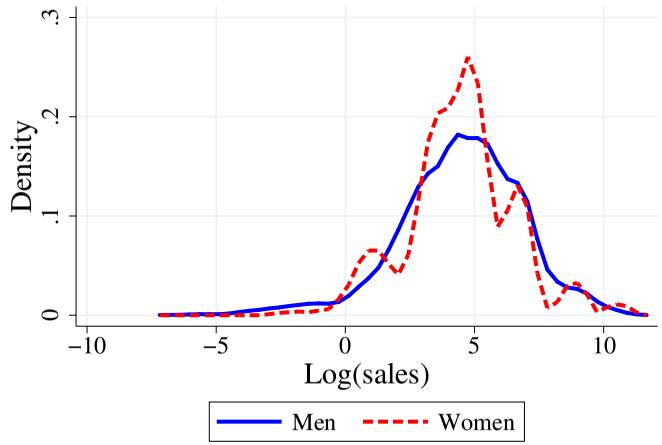

(b) Log(sales)

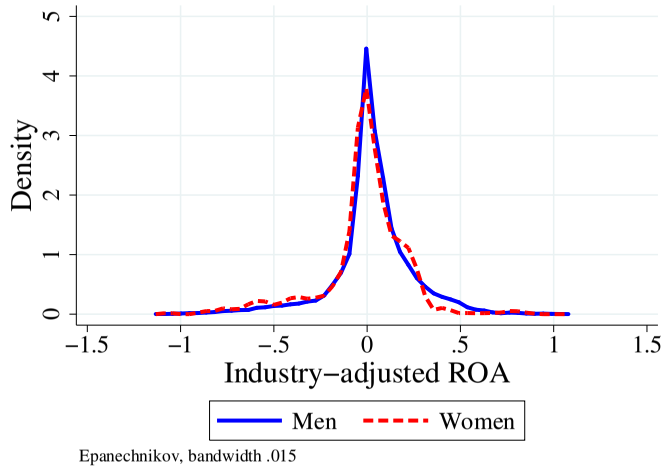

(c) Industry-adjusted ROA

Figure A.2.4: Firm performance

Notes: (a) Monthly stock price returns (Source: CRSP) (b) Log(sales) (Source: Compustat) (c) Industry-adjusted Return on Assets (Source: Compustat). Quarterly observations between 2000 and 2017. Monthly stock price returns in Figure (a) are averaged over the corresponding quarter. Industry-adjusted Return on Assets in Figure (c) is calculated as quarterly ROA minus average industry ROA over the corresponding quarter.

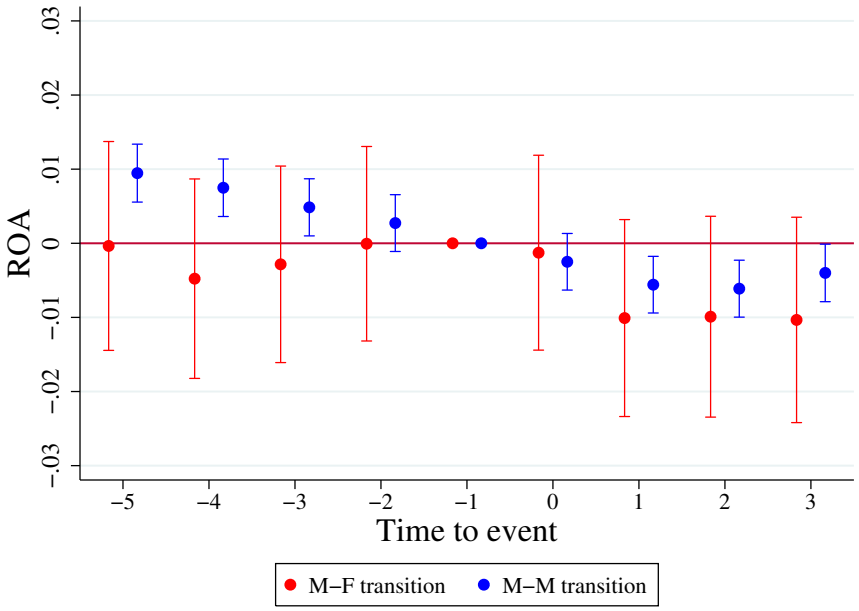

Figure A.2.5: Firm performance around appointment

Notes: Return on Assets around CEO appointment. Coefficients and 95% confidence intervals from a regression of quarterly ROA on year fixed effects, firm fixed effects, and dummies for the leads and lags from the transition event. Standard errors are clustered at the firm level. The sample includes 200 male-to-female transition events and 3,293 male-to-male transition events.

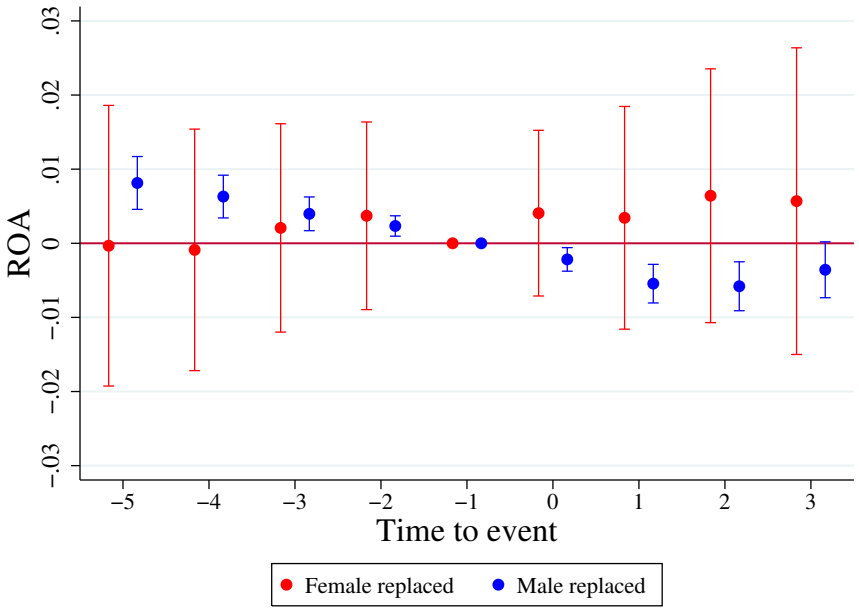

Figure A.2.6: Firm performance around replacement

Notes: Return on Assets around CEO replacement. Coefficients and 95% confidence intervals from a regression of quarterly ROA on year fixed effects, firm fixed effects, and dummies for the leads and lags from the replacement event. Standard errors are clustered at the firm level. The sample includes 120 female replacement events and 3,493 male replacement events.

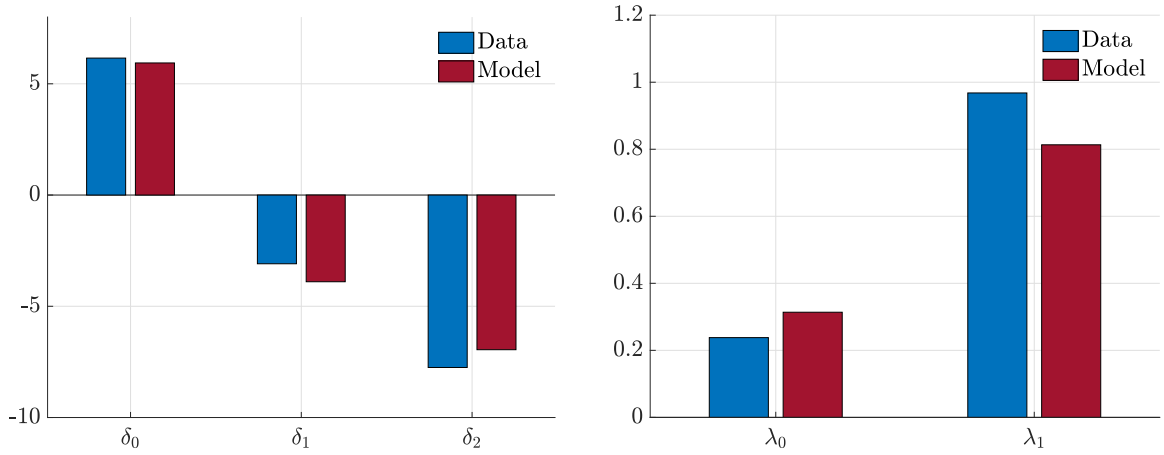

(a) Profitability by publication state

(b) Profitability $AR(1)$

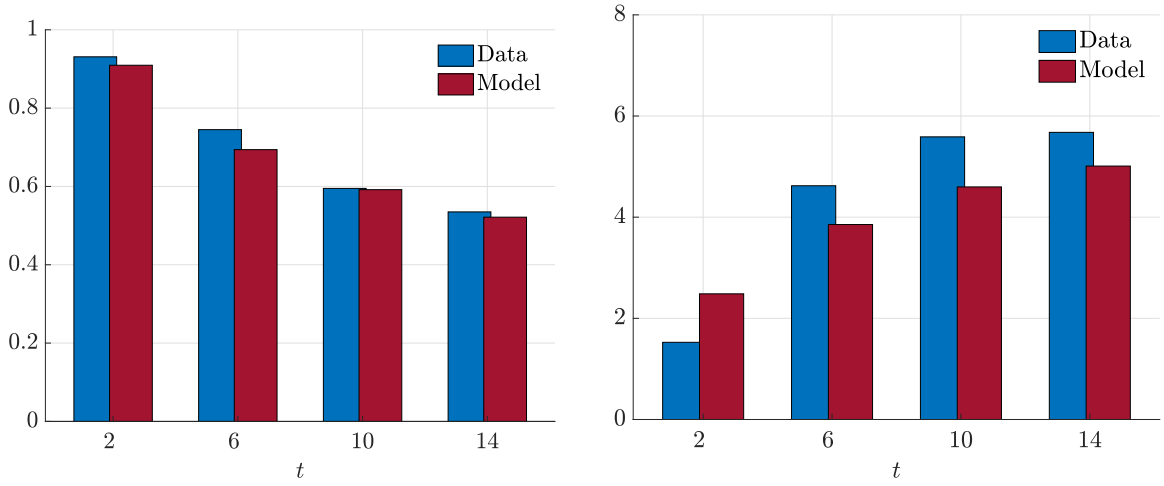

(c) Survival function

(d) Profitability

Figure A.2.7: Model fit: Target moments

Notes: (a) Profitability by publication state. The coefficients δ_0, δ_1, and δ_2 are obtained from the regression $q_{it} = \delta_0 + \delta_1 pub2_{it-1} + \delta_2 pub3_{it-1} + \epsilon_{it}$ where q_{it} is industry-adjusted ROA for firm i in quarter t, and $pub2_{it}$ and $pub3_{it}$ are two dummies for whether the history of negative publications in firm i and quarter t belong to the second or third tercile. (b) Profitability $AR(1)$. The coefficients λ_0 and λ_1 are obtained from the $AR(1)$ regression $q_{it} = \lambda_0 + \lambda_1 q_{it-1} + \epsilon_{it}$. (c) Survival function. Survival function at different tenure times. (d) Profitability. Average firm profitability at different tenure times.

A.3 Additional Tables

Table A.3.1: Comparison between unmatched and matched positions and firms

	Unmatched			Matched				
	Mean	Standard deviation	N	Mean	Standard deviation	N	Difference	p-value
Panel A. Individual characteristics								
Female	0.02	0.141	100	0.043	0.202	3,026	-0.023	0.267
Age	51.156	8.977	77	52.6	8.174	2,970	-1.444	0.127
Appointed after 2010	0.41	0.494	100	0.527	0.499	3,026	-0.117	0.021
Born in the US	0.778	0.424	27	0.919	0.273	1,329	-0.141	0.009
Bachelor's degree	0.287	0.455	87	0.288	0.453	2,756	-0.001	0.982
Master's/MBA/Professional degree	0.425	0.497	87	0.449	0.497	2,756	-0.024	0.664
Doctorate degree	0.149	0.359	87	0.152	0.359	2,756	-0.003	0.954
Number of qualifications	2.033	1.326	30	1.919	1.09	2,070	0.114	0.571
Appointment duration (days)	676.068	693.603	74	695.772	764.012	2,749	-19.704	0.826
Tenure in company (years)	3.217	3.962	30	6.569	8.338	2,070	-3.352	0.028
Network size	920.798	1338.01	99	1175.933	1429.706	2,970	-255.135	0.080
Total number of boards	1.867	1.57	30	1.936	1.649	2,029	-0.069	0.818
Number of positions			100			3,026		
Panel B. Board characteristics								
Gender ratio	0.97	0.08	30	0.91	0.11	2,070	0.06	0.001
Number of directors	6.90	1.52	30	8.45	2.50	2,070	-1.55	0.001
Panel C. Firm characteristics								
Assets	2,243.82	8,572.30	38	7,997.51	72,424.21	2,449	-5753.69	0.624
Employees	3.22	10.04	37	8.43	28.77	2,376	-5.21	0.271
Sales	1,227.37	4,300.95	38	2,597.52	10,289.15	2,432	-1370.15	0.413
Gross profits	501.79	2,154.67	38	845.87	3,383.40	2,432	-344.09	0.532
Market value	3,311.84	17,723.33	36	3,660.20	16,344.71	2,216	-348.37	0.899
Primary sector	0.08	0.28	96	0.14	0.35	2,870	-0.06	0.104
Consumer sector	0.17	0.38	96	0.15	0.36	2,870	0.02	0.664
Service sector	0.75	0.44	96	0.71	0.46	2,870	0.04	0.369
Number of firms			75			2,039		

Notes: Source: Panel A and B: BoardEx, 2000-2017, Panel C: Compustat, 2000-2017. Individual and board characteristics are measured in the year of the appointment (except *Appointment duration*), whereas firm characteristics are measured the year before the appointment.

Table A.3.2: CEOs: Differences in news coverage

Dependent variable: News coverage (z-scores)					
	(1)	(2)	(3)	(4)	(5)
Female	0.032	−0.005	0.314***	0.327***	0.387***
	(0.058)	(0.044)	(0.109)	(0.111)	(0.117)
Network size		0.000***	−0.000	−0.000	−0.000*
		(0.000)	(0.000)	(0.000)	(0.000)
Born in the US		0.035*	−0.029	−0.044	−0.048
		(0.020)	(0.101)	(0.098)	(0.099)
Number of qualifications		0.003	-0.015	−0.022	−0.048
		(0.010)	(0.049)	(0.047)	(0.048)
Age		0.019	0.026	0.033	0.010
		(0.014)	(0.029)	(0.029)	(0.029)
Age sq.		−0.000	−0.000	−0.000	−0.000
		(0.000)	(0.000)	(0.000)	(0.000)
Tenure		−0.007***	−0.004	−0.004	−0.004
		(0.002)	(0.004)	(0.004)	(0.004)
Tenure sq.		0.000***	0.000	0.000	0.000
		(0.000)	(0.000)	(0.000)	(0.000)
Appointment news=1				0.213***	0.212***
				(0.028)	(0.029)
Resignation news=1				0.442***	0.448***
				(0.066)	(0.067)
Sentiment score				0.000	0.000
				(0.001)	(0.001)
Number of listed boards					0.066*
					(0.038)
Tenure in company					−0.010**
					(0.004)
Year FE	Y	Y	Y	Y	Y
Year of appointment FE	N	Y	Y	Y	Y
Firm FE	N	N	Y	Y	Y
N	18703	18703	18703	18703	18300

Notes: Observations are news events released between 2000 and 2017. Every news event specifically mentions the CEO as the primary individual involved in the news event. The dependent variable is represented by the total number of articles for a news event, standardized into z-scores. Standard errors are clustered at the position level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.3.3: Other Chief Officers: Differences in news coverage

Dependent variable: News coverage (z-scores)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	CFOs/COOs				Other Chief Officers (CAOs, CMOs, CTOs)			
Female	0.031	−0.003	0.355**	0.263**	0.037	−0.000	0.238+	0.231+
	(0.052)	(0.054)	(0.166)	(0.135)	(0.086)	(0.090)	(0.179)	(0.174)
Network size		0.000***	−0.000	−0.000		0.000**	−0.000+	−0.000+
		(0.000)	(0.000)	(0.000)		(0.000)	(0.000)	(0.000)
Born in the US		−0.001	−0.305***	−0.191**		0.039	−0.404	−0.417
		(0.035)	(0.107)	(0.082)		(0.062)	(0.351)	(0.355)
Number of qualifications		0.016	0.091**	0.075**		−0.040	−0.092	−0.076
		(0.014)	(0.040)	(0.035)		(0.033)	(0.107)	(0.106)
Age		0.038	0.163**	0.182***		0.015	0.125	0.152+
		(0.024)	(0.080)	(0.065)		(0.029)	(0.104)	(0.102)
Age sq.		−0.000	−0.002**	−0.002***		−0.000	−0.001	−0.001+
		(0.000)	(0.001)	(0.001)		(0.000)	(0.001)	(0.001)
Tenure		−0.008**	0.002	0.001		−0.001	0.026	0.024
		(0.003)	(0.004)	(0.003)		(0.007)	(0.032)	(0.029)
Tenure sq.		0.000***	0.000**	0.000		0.000	0.000+	0.000
		(0.000)	(0.000)	(0.000)		(0.000)	(0.000)	(0.000)
COO=1		−0.011	0.089	0.102				
		(0.040)	(0.093)	(0.077)				
Appointment news=1				0.786***				0.445+
				(0.061)				(0.321)
Resignation news=1				1.576***				0.680+
				(0.265)				(0.425)
Sentiment score				−0.000				0.006
				(0.002)				(0.027)
Year FE	Y	Y	Y	Y	Y	Y	Y	Y
Year of appointment FE	N	Y	Y	Y	N	Y	Y	Y
Firm FE	N	N	Y	Y	N	N	Y	Y
Observations	11295	11295	11295	11295	1271	1271	1271	1271

Notes: Observations are news events released between 2000 and 2017. Every news event specifically mentions an executive as the primary individual involved in the news event. The dependent variable is represented by the total number of articles for a news event, standardized into z-scores. Standard errors are clustered at the position level. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$. $+p < 0.20$ $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

Table A.3.4: Differences in firm performance

	(1)	(2)	(3)	(4)
		A. Stock price returns		
	OLS	Q(0.25)	Q(0.5)	Q(0.75)
Female	-0.008	-0.005	-0.001	0.001
	(0.017)	(0.004)	(0.003)	(0.006)
CEO char.	Y	Y	Y	Y
Firm size	Y	Y	Y	Y
Firm FE	Y	N	N	N
Year FE	Y	Y	Y	Y
N	15,742	15,742	15,742	15,742
		B. Log(sales)		
	OLS	Q(0.25)	Q(0.5)	Q(0.75)
Female	-0.020	0.256***	0.309***	0.401***
	(0.053)	(0.088)	(0.055)	(0.056)
CEO char.	Y	Y	Y	Y
Firm size	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Firm FE	Y	N	N	N
N	18,133	18,133	18,133	18,133
		C. Industry-adjusted ROA		
	OLS	Q(0.25)	Q(0.5)	Q(0.75)
Female	0.031	0.002	-0.012	-0.020**
	(0.028)	(0.004)	(0.008)	(0.009)
CEO char.	Y	Y	Y	Y
Firm size	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Firm FE	Y	N	N	N
N	13,281	13,281	13,281	13,281

Notes: Quarterly observations between 2000 and 2017. The dependent variable is represented by quarterly stock price returns (Panel A), the logarithm of quarterly sales (Panel B), and industry-adjusted Return on Assets (Panel C). Quarterly stock price returns are calculated as monthly returns averaged over the corresponding quarter. Industry-adjusted ROA is calculated as quarterly ROA minus quarterly average industry ROA. OLS regression in column 1 and quantile regressions in columns 2–4. CEO characteristics include a quadratic in age and a quadratic in tenure. Firm size is represented by the the logarithm of assets. Standard errors are clustered at the position level in column 1 and bootstrapped in columns 2–4. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.3.5: Turnover and news: Cox proportional hazard

	(1)	(2)	(3)	(4)	(5)	(6)
	End of:		*CEO moves to:*	*End of:*		*CEO moves to:*
	CEO app.	All app.	Private or smaller firm, missing move	CEO app.	All app.	Private or smaller firm, missing move
	A. High-coverage firms			*B. All firms*		
Negative articles	0.0165	0.0561**	0.0444*	0.0209	0.0581***	0.0549**
	(0.0141)	(0.0221)	(0.0242)	(0.0138)	(0.0207)	(0.0232)
Positive articles	-0.0189	-0.0290	-0.0293	-0.0315*	-0.0632**	-0.0637*
	(0.0174)	(0.0316)	(0.0338)	(0.0174)	(0.0310)	(0.0339)
Female	0.2218	0.1766	-0.1737	0.2214	0.1815	-0.0750
	(0.1839)	(0.2863)	(0.3515)	(0.1415)	(0.2127)	(0.2480)
ROA	-0.1260	-0.4163**	-0.4351**	-0.0299	-0.1590	-0.1713
	(0.1384)	(0.1833)	(0.1962)	(0.0856)	(0.1058)	(0.1083)
CEO controls	Y	Y	Y	Y	Y	Y
Sector FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
N	9,673	9,673	9,673	15,944	15,944	15,944

Notes: Quarterly observations between 2000 and 2017. High coverage firms (Panel A) include firms for which the median number of articles in a quarter is above the median across all firms. CEO controls include network size, the number of qualifications, a quadratic in age. All regressions include controls for the total number of articles released. Standard errors are clustered at the firm level. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

A.4 Additional Results

A.4.1 Addressing potential biases in news coverage results

In order to understand news selection decisions, it would be ideal to check how the *unconditional* distribution of events maps to the distribution of *reported* news events. The crucial issue is that observing the unconditional distribution of all events in a firm is not possible, and one would need to make assumptions on the shape of such distribution. The main advantage of measuring coverage decisions using the sample of reported events is to avoid imposing structure on the underlying unconditional distribution. Moreover, looking at the extensive margin of publication decisions – rather than the intensive margin – would carry the risk of confounding media decisions with firms' decisions. In fact, much of the information reported by financial outlets is day-to-day information released by the company itself. In order to better understand whether news media decisions significantly differ at the extensive margin between male- and female-headed firms, I check whether the number of news events released in a given unit of time differs for male- versus female-headed firms. The underlying assumption is that the unconditional distribution of all events is the same across male- and female-headed firms. Such assumption is supported by empirical evidence: the data do not reveal any differences in

performance between male- and female-led firms. I then aggregate news events at the quarterly level and estimate the equation:

$$\text{Number of news events}_{it} = \alpha + \text{CEO char}_{it}\,\delta + \text{Perf}_{it}\,\eta + \phi_{f(it)} + \tau_t + \nu_{it} \tag{A.13}$$

where Number of news events$_{it}$ is the number of events linked to CEO i in quarter t, CEO char$_{it}$ are CEO characteristics, including a quadratic in age and in tenure, Perf$_{it}$ are performance controls, and $\phi_{f(it)}$ and τ_t are firm and time fixed effects. I estimate both OLS and quantile regressions. Because I cannot control for sector or firm fixed effects when running quantile regressions, I add controls for firm size, represented by the logarithm of assets. I focus on performance-related news events only, and exclude events related to acquisition and mergers, legal and labor issues, and products and services. I also exclude all performance events related to bankruptcy.[3] In Table A.4.1 I present coefficient estimates for the OLS regression in Equation A.13 and quantile regressions, separately for positive and negative events. On average, there is no significant difference in the number of news events released for male- and female-led firms, neither when looking at positive or negative events. A small, positive difference shows up in the quantile regressions for the sample of negative events: at the 75th percentile of the distribution, the difference for female-

[3] I run the same analysis on the full sample of events, and find very similar results.

headed firms is 0.14 news events.

Table A.4.1: Number of negative and positive news events in a quarter

	(1)	(2)	(3)	(4)
		A. Negative events		
	OLS	Q(0.25)	Q(0.5)	Q(0.75)
Female	0.377	0.021**	0.120***	0.142***
	(0.235)	(0.009)	(0.034)	(0.044)
CEO char.	Y	Y	Y	Y
Firm performance	Y	Y	Y	Y
Firm size	N	Y	Y	Y
Firm FE	Y	N	N	N
Year FE	Y	Y	Y	Y
N	18,133	18,133	18,133	18,133
		B. Positive events		
	OLS	Q(0.25)	Q(0.5)	Q(0.75)
Female	-0.001	-0.010	-0.008	0.049
	(0.181)	(0.023)	(0.038)	(0.053)
CEO char.	Y	Y	Y	Y
Firm performance	Y	Y	Y	Y
Firm size	N	Y	Y	Y
Firm FE	Y	N	N	N
Year FE	Y	Y	Y	Y
N	18,133	18,133	18,133	18,133

Notes: Quarterly observations between 2000 and 2017. The dependent variable is represented by the number of news events in a quarter. Negative events are represented by news events at the bottom 10% of the sentiment distribution (Panel A), whereas positive events belong to the top 90% of the sentiment distribution. OLS regression in column 1 and quantile regressions in columns 2–4. The estimating specification is Equation A.13 in the text. CEO characteristics include a quadratic in age and a quadratic in tenure. Firm performance is represented by the logarithm of sales and firm size by the the logarithm of assets. Standard errors are clustered at the firm level in column 1 and bootstrapped in columns 2–4. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

A.4.2 Assessing differences in prior beliefs

In this section, I check whether there exist any differences in prior beliefs on the ability of male versus female leaders. Given the underrepresentation of women in executive positions, the board may have more dispersed prior beliefs on the managerial ability of women.[4] Understanding the role of uncertainty is particularly important in my setting. First, if firm-level uncertainty increases following the appointment of a female CEO, the value of information to the general public may increase, thus explaining why female top executives are more monitored than their male counterparts. Second, differences in prior beliefs alone can rationalize the higher incidence of turnover measured for female appointments. Bayesian updating implies that the relative weight of new information depends on the precisions of prior information and the signal. For the same level of signal precision, the weight of new information is larger when prior information is less precise. Therefore, a less precise prior would increase the weight of new information and lead to faster board's learning. Holding everything else constant, faster board's learning would increase the turnover hazard over the first years of tenure for female CEOs relative to their male counterparts. A similar argument would apply to prior public beliefs. The higher the dispersion in prior beliefs, the more the weight of new in-

[4]Because the dispersion of firm performance is not different across male- and female-led firms, more dispersed prior beliefs on the ability of female CEOs would imply a departure from rational expectations.

formation provided by news media, which would then increase the incidence of turnover. In order to understand the dispersion in prior beliefs at the start of a female appointment, I check the evolution of (i) firm-level uncertainty and (ii) analysts' beliefs around the appointment of a new CEO, comparing male-to-female transitions to male-to-male transitions. Because data on expectations and volatility are sparse, I can only match 12% of my sample of individual CEOs to data on expectations and volatility. To increase the number of observations, I extend my sample of 3,026 individual CEOs to include CEOs that are also the company's President. As a measure of firm-level uncertainty, I use data on the volatility of firm equity options, calculated by OptionMetrics.[5] I form two portfolios of firms, corresponding to male-to-male and male-to-female transitions, and check the evolution of average monthly volatility of firm equity options around CEO appointment, separately for the two portfolios. The results are plotted in Figure A.4.1. In the 6 months before the appointment, the two portfolios closely follow each other. Firm-level uncertainty increases slightly in the month of CEO transition, but only for male-to-male appointments. In Appendix Figure A.4.1, Panel (b), I zoom-in closer and focus on 10-day volatility calculated in each of the 25 days around CEO appointment. CEO appointment increases firm-level uncertainty in both groups of firms, and the two

[5]Option volatility is commonly used in the corporate finance and macroeconomics literature to measure firm-level uncertainty. Two prominent examples include Baker, Bloom, and Davis (2016), and Kelly, Pástor, and Veronesi (2016).

portfolios very closely follow each other.

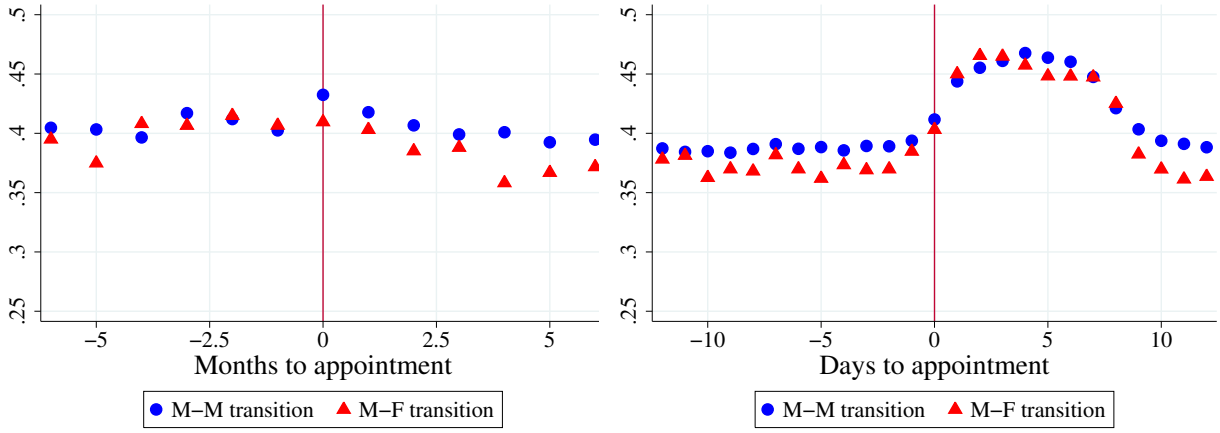

(a) Monthly volatility (b) Daily volatility

Figure A.4.1: Volatility of firm equity options around appointment

Notes: (a) Average volatility of firm equity options, measured on the last trading day of the month and calculated in the preceding 30-day horizon. The sample includes 117 male-to-female transitions and 1,817 male-to-male transitions. (b) Average volatility of firm equity options, measured daily and calculated over the preceding 10-day horizon. The red vertical bar corresponds to the day of CEO appointment. The sample includes 89 male-to-female transitions and 1,396 male-to-male transitions.

In order to have a more direct measure of dispersion in beliefs, I use IBES data on analysts' expectations.[6] I match to firms analysts' monthly forecasts of earnings per share (EPS) at a one-year horizon, and form two portfolios of firms, corresponding to male-to-male and male-to-female transitions. In order to proxy for uncertainty in analysts' beliefs, I focus on two measures. First, I calculate the forecast error, de-

[6]Such data are becoming increasingly common in recent work in corporate finance. Examples include Ben-David, Graham, and Harvey (2013), Greenwood and Shleifer (2014), Gennaioli, Ma, and Shleifer (2016), Bouchaud, Krueger, Landier, and Thesmar (2019) and Bordalo, Gennaioli, Porta, and Shleifer (2019).

fined as the difference between realized EPS and the average forecast. As a second measure, I use the standard deviation of analysts' forecasts. The results are plotted in Figure A.4.2. Again, I do not detect any significant increase in uncertainty following the appointment of a female CEO. In fact, the average forecast error – overoptimistic before appointment in both portfolios – converges to zero more quickly following the appointment of a female CEO. For male-to-male appointments, the transition is smoother, and I do not detect any deviation from the trend around the month of CEO appointment. In general, Figure A.4.2 suggests that analysts do not revise their forecasts dramatically following the appointment of a new CEO, and that expectations are highly-path dependent, at least in the short term.[7] Similarly, Figure A.4.2 shows no evidence of higher disagreement among analysts when evaluating female-led firms: the trend is flat both before and after the appointment, with no significant change in the intercept around the time of CEO appointment.

[7] I find similar results when looking at forecasts of long-term earnings growth. On the persistency of forecast errors, see for example Ma et al. (2020).

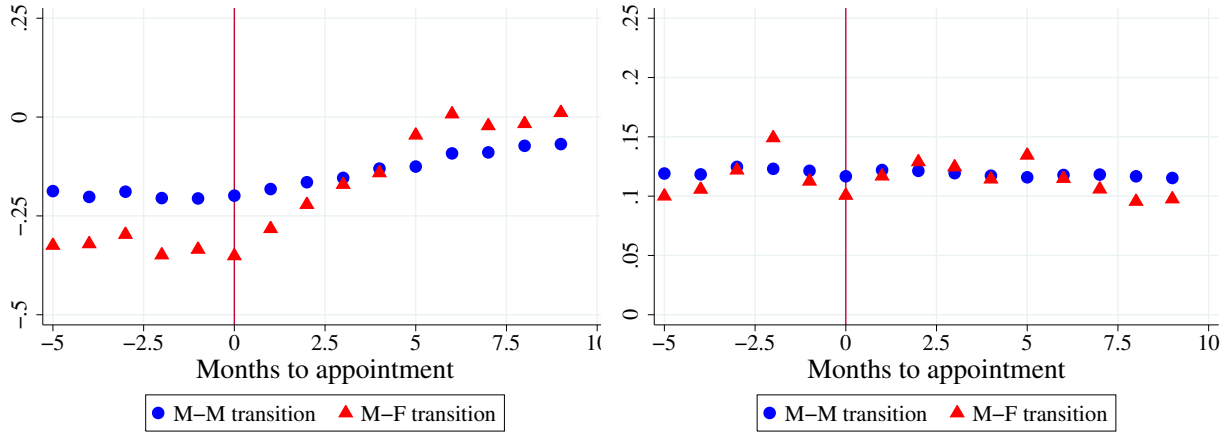

(a) Average forecast error

(b) Standard deviation

Figure A.4.2: Analysts' expectations around appointment

Notes: (a) Average forecast error, calculated as the difference between actual EPS and the average forecasted EPS. The forecast period corresponds to one year. The sample includes 53 male-to-female transitions and 1,047 male-to-male transitions. (b) Average standard deviation of analysts' EPS expectations. The forecast period corresponds to one year. The sample includes 53 male-to-female transitions and 1,047 male-to-male transitions.

A.4.3 Assessing differences in CEO power

As a final test, I check whether female CEOs are less powerful than their male counterparts, or more likely to be appointed following powerful CEOs. If women are systematically appointed following particularly influential or long-tenured leaders, investors' uncertainty regarding the new leadership may arise, even if not due to gender per se. This hypothesis is similar in spirit with the previous one, and is in line with the so called "glass cliff" hypothesis, according to which women and

other minorities are more likely to be appointed in particularly difficult or precarious positions. In Table A.4.2, I focus on my main sample of CEOs and check the characteristics of the current CEO and his or her predecessor, separately by gender. In the first panel, I compare male and female CEOs across firms and show that, on average, female CEOs are not less powerful than their male counterparts. The only significant difference arises when looking at the share of independent board members, as female-led firms tend to have slightly more independent boards. In the second panel, I check how male and female CEOs compare when considering their predecessors: again, I do not find evidence that women are more likely to be appointed following particularly powerful leaders. The results suggest that female CEOs are not less powerful than their male counterparts, and that uncertainty regarding female leadership is unlikely to account for the observed patterns.

Table A.4.2: CEO power: current CEOs and their predecessors

Current CEO is:	*Female*		*Male*			
	Mean	N	Mean	N	Diff.	p-val.
(a) Current CEO:						
First appointment	0.581	129	0.624	2897	−0.04	0.324
Tenure in company (years)	8.272	101	7.916	2226	0.36	0.673
Founder	0.101	129	0.075	2892	0.03	0.288
Share of indep. board members	0.891	101	0.852	2225	0.04	0.014
Appointment duration (days)	662	118	707	2660	−45	0.531
(b) Predecessor CEO:						
Female	0.514	109	0.015	2389	0.50	0.000
Tenure in company (years)	10.444	95	10.612	2029	−0.17	0.874
Founder	0.138	109	0.128	2387	0.01	0.764
Chair	0.349	109	0.31	2389	0.04	0.392
Share of indep. board members	0.853	95	0.832	2028	0.02	0.234
Appointment duration (days)	1415	109	1463	2388	−48	0.777

Notes: Average characteristics of the current CEO (in Panel a) and average characteristics of the predecessor CEO (Panel B), by gender of the current CEO. *Tenure in the company* refers to the number of years as employee in the appointing company at the time of CEO appointment.

Appendix B

Do Board Gender Quotas Matter? Selection, Performance and Stock Market Effects

Table B.0.1: Summary statistics, individual-level data

		2007	2008	2009	2010	2011	2012	2013	2014
Panel A. Board of Directors									
College degree	All	0.83	0.84	0.85	0.84	0.85	0.86	0.88	0.88
	F	0.68	0.69	0.71	0.76	0.83	0.88	0.90	0.90
	M	0.84	0.85	0.86	0.85	0.85	0.86	0.88	0.87
Graduate degree	All	0.08	0.07	0.08	0.08	0.08	0.07	0.08	0.1
	F	0.09	0.08	0.10	0.09	0.11	0.10	0.14	0.18
	M	0.08	0.07	0.08	0.08	0.07	0.07	0.07	0.08
Studied abroad	All	0.08	0.09	0.08	0.08	0.08	0.09	0.07	0.08
	F	0.06	0.07	0.07	0.08	0.09	0.12	0.10	0.12
	M	0.08	0.09	0.08	0.08	0.08	0.08	0.07	0.07
Economics degree	All	0.37	0.38	0.37	0.38	0.4	0.42	0.42	0.42
	F	0.39	0.36	0.31	0.34	0.40	0.43	0.41	0.43
	M	0.37	0.38	0.38	0.39	0.40	0.42	0.42	0.42
Law degree	All	0.15	0.14	0.14	0.14	0.14	0.13	0.15	0.15
	F	0.10	0.08	0.10	0.10	0.11	0.14	0.19	0.20
	M	0.15	0.15	0.15	0.14	0.14	0.13	0.14	0.14
Younger than 55	All	0.44	0.44	0.46	0.46	0.48	0.45	0.47	0.46
	F	0.65	0.62	0.59	0.64	0.71	0.73	0.72	0.71
	M	0.43	0.43	0.45	0.44	0.46	0.41	0.41	0.40
Family tie	All	0.13	0.12	0.08	0.08	0.08	0.11	0.11	0.11
	F	0.40	0.35	0.20	0.18	0.14	0.17	0.12	0.12
	M	0.11	0.11	0.07	0.07	0.07	0.10	0.10	0.11
Number of positions	All	1.38	1.41	1.35	1.31	1.29	1.27	1.26	1.25
	F	1.40	1.41	1.29	1.21	1.17	1.20	1.16	1.20
	M	1.38	1.40	1.35	1.32	1.30	1.29	1.29	1.27
Retained	All	.	.	.	0.49	0.38	0.33	0.26	0.28
	F	.	.	.	0.63	0.59	0.58	0.57	0.60
	M	.	.	.	0.62	0.57	0.55	0.51	0.54
N	All	2048	2092	2158	2194	2276	2299	2426	2291

Continued on next page

Table B.0.1: Summary statistics, individual-level data

		2007	2008	2009	2010	2011	2012	2013	2014
	F	116	119	137	160	207	284	432	474
	M	1932	1973	2021	2034	2069	2015	1994	1817
Panel B. Board of Auditors									
College degree	All	0.92	0.92	0.91	0.91	0.92	0.93	0.93	0.95
	F	0.88	0.90	0.86	0.88	0.85	0.90	0.92	0.94
	M	0.92	0.92	0.92	0.92	0.92	0.93	0.93	0.95
Graduate degree	All	0.02	0.01	0.02	0.03	0.03	0.03	0.03	0.03
	F	0.00	0.00	0.00	0.00	0.00	0.00	0.02	0.04
	M	0.02	0.02	0.02	0.03	0.03	0.03	0.03	0.03
Studied abroad	All	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	F	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01
	M	0.01	0.01	0.00	0.00	0.00	0.00	0.00	0.00
Economics degree	All	0.74	0.73	0.73	0.76	0.78	0.78	0.79	0.8
	F	0.75	0.69	0.73	0.77	0.76	0.78	0.78	0.80
	M	0.74	0.74	0.73	0.76	0.78	0.78	0.79	0.80
Law degree	All	0.07	0.07	0.07	0.07	0.08	0.07	0.08	0.07
	F	0.12	0.21	0.14	0.12	0.09	0.07	0.08	0.08
	M	0.07	0.06	0.06	0.07	0.08	0.07	0.07	0.07
Younger than 55	All	0.48	0.51	0.55	0.56	0.58	0.52	0.54	0.54
	F	0.73	0.70	0.66	0.70	0.74	0.74	0.77	0.76
	M	0.47	0.50	0.54	0.55	0.56	0.49	0.49	0.47
Family tie	All	0.00	0.00	0.00	0.00	0.00	0.01	0.01	0.01
	F	0.00	0.00	0.00	0.00	0.00	0.01	0.02	0.01
	M	0.00	0.00	0.00	0.00	0.00	0.01	0.01	0.01
Number of positions	All	1.43	1.38	1.34	1.32	1.31	1.24	1.23	1.21
	F	1.08	1.07	1.08	1.09	1.11	1.15	1.21	1.22
	M	1.45	1.40	1.35	1.34	1.33	1.25	1.23	1.20
Retained	All	.	.	.	0.63	0.62	0.57	0.48	0.46
	F	.	.	.	0.50	0.50	0.37	0.22	0.20

Continued on next page

Table B.0.1: Summary statistics, individual-level data

		2007	2008	2009	2010	2011	2012	2013	2014
	M	.	.	.	0.64	0.63	0.60	0.53	0.54
N	All	602	617	642	627	645	649	703	623
	F	26	30	38	43	46	72	128	148
	M	576	587	604	584	599	577	575	475

Notes: Averages of individual characteristics of Italian board members of listed companies, 2007-2014.

Table B.0.2: Board of Auditors – Effect on board characteristics

Assumption on time trend:			No trend			Linear trend			Non-parametric		
		Mean before quotas		R^2	N		R^2	N		R^2	N
Percentage of women		7.24	23.93***	0.51	519	21.01***	0.51	519	23.73***	0.43	1633
			(1.50)			(2.87)			(1.53)		
More than 20% women		0.01	2.21**		25	-18.73		25	19.13		84
			(1.10)			(20879.27)			(9466.95)		
Fem. President		0.02	4.10***		111	4.11*		111	4.08**		315
			(1.02)			(2.36)			(2.03)		
% college degree		88.30	5.56**	0.06	519	-5.26	0.09	519	0.05	0.05	1633
			(2.17)			(4.11)			(2.31)		
	F	84.09	2.44	0.02	200	-6.71	0.06	200	1.70	0.06	496
			(3.32)			(6.88)			(3.38)		
	M	88.39	4.16*	0.04	519	-6.57	0.07	519	-0.28	0.04	1633
			(2.38)			(4.53)			(2.45)		
% graduate degree		3.04	2.18***	0.03	519	0.33	0.03	519	1.25	0.01	1633
			(0.82)			(1.58)			(0.97)		
	F	0.00	6.71*	0.08	200	3.36	0.08	200	6.67***	0.08	496
			(3.37)			(7.13)			(2.40)		
	M	3.20	1.59	0.01	519	-0.76	0.02	519	1.58	0.01	1633
			(0.99)			(1.90)			(1.14)		
% studied abroad		0.16	0.35	0.00	519	0.34	0.00	519	0.54	0.02	1633
			(0.32)			(0.62)			(0.42)		
	F	0.00	0.00	.	200	0.00	.	200	0.00	.	496
			(.)			(.)			(.)		
	M	0.16	-0.26	0.01	519	-0.50	0.01	519	0.25	0.04	1633
			(0.18)			(0.34)			(0.42)		
Filed diversity		0.88	0.02	0.01	510	-0.04	0.02	510	0.00	0.02	1606
			(0.02)			(0.04)			(0.02)		
% economics degree		78.12	7.97***	0.08	519	-5.88	0.12	519	1.76	0.07	1633
			(2.35)			(4.41)			(2.51)		
	F	75.00	-2.44	0.02	200	-10.74	0.06	200	-4.60	0.03	496

Continued on next page

Table B.0.2: Board of Auditors – Effect on board characteristics

Assumption on time trend:			No trend			Linear trend			Non-parametric		
		Mean before quotas		R^2	N		R^2	N		R^2	N
			(3.32)			(6.91)			(3.44)		
	M	78.12	5.40**	0.04	519	-8.73*	0.08	519	0.34	0.05	1633
			(2.66)			(5.03)			(2.76)		
% law degree		7.84	-0.55	0.00	519	-1.26	0.01	519	-0.87	0.01	1633
			(1.19)			(2.30)			(1.35)		
	F	9.09	0.00	0.08	200	-0.00	0.08	200	-0.10	0.18	496
			(1.79)			(3.79)			(1.59)		
	M	7.61	0.88	0.00	519	0.11	0.00	519	0.57	0.01	1633
			(1.47)			(2.84)			(1.56)		
% younger than 55		58.42	5.95**	0.05	519	-2.27	0.06	519	6.77**	0.08	1633
			(2.64)			(5.05)			(3.14)		
	F	75.00	3.66	0.01	200	-18.79	0.06	200	-2.54	0.08	496
			(7.92)			(16.38)			(7.36)		
	M	56.81	-5.17*	0.05	519	-10.92*	0.05	519	-2.98	0.06	1633
			(3.01)			(5.79)			(3.49)		
% family ties		0.32	0.39	0.04	519	0.45	0.04	519	0.95	0.03	1633
			(0.77)			(1.48)			(0.81)		
	F	0.00	0.00	0.08	200	0.00	0.08	200	0.65	0.07	496
			(2.68)			(5.69)			(2.28)		
	M	0.48	0.29	0.02	519	0.33	0.02	519	0.74	0.02	1633
			(0.88)			(1.70)			(0.90)		
Number of positions		1.33	-0.05	0.05	519	0.17**	0.09	519	0.15***	0.10	1633
			(0.04)			(0.08)			(0.05)		
	F	1.06	0.18***	0.19	200	0.17	0.19	200	-0.07	0.10	496
			(0.05)			(0.11)			(0.07)		
	M	1.35	-0.03	0.03	518	0.19**	0.06	518	0.20***	0.09	1632
			(0.05)			(0.09)			(0.06)		

Notes: The table shows the coefficient on the reform indicator in a regression where the dependent variable is shown in the first column. In every regression, we control for the number of board members and the phase-in period. The three specifications for every regression correspond to different assumptions on the time trend. In the first two specifications ("No trend" and "Linear trend") observations are for election years over the period 2007-2014. In the third specification, we include observations over all years between 2007 and 2014, and add an election year fixed effect. The models with dependent variable *More than 20% women*, *Female CEO*, and *Female President* are estimated using a logit model, which explains the lower number of observations. The others are estimated using a linear model. Every regression controls for board fixed effects. Standard errors are clustered at the board level. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

Table B.0.3: The effect of women directors on economic performance and variability of stock market prices

	(1) Log(empl.)	(2) Log(prod.)	(3) Log(profits)	(4) Log(assets)	(5) ROA	(6) Tobin's Q	(7) Short-term debt	(8) Sd(stock prices)
Panel B – bis. ITT								
Reform	-0.124	-0.363	-0.100	0.004	0.000	-0.088	0.027	-0.002***
	(0.233)	(0.318)	(0.371)	(0.095)	(2.109)	(0.086)	(0.053)	(0.001)
Reform × Distance	0.023**	0.022*	-0.013	0.006	-0.065	0.006*	-0.001	0.004
	(0.009)	(0.012)	(0.016)	(0.004)	(0.083)	(0.003)	(0.002)	(0.004)
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
Time FE	Y	Y	Y	Y	Y	Y	Y	Y
R-squared	0.423	0.345	0.086	0.186	0.032	0.087	0.257	0.175
Panel C – bis. IV								
Percentage of women	0.021**	0.013	-0.014	0.008**	-0.068	0.005	0.001	-0.008***
	(0.010)	(0.012)	(0.017)	(0.004)	(0.087)	(0.003)	(0.002)	(0.003)
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
Time FE	Y	Y	Y	Y	Y	Y	Y	Y
R-squared	0.395	0.333	0.073	0.178	0.034	0.075	0.248	0.171
Mean of dep. var.	5.676	11.183	9.583	13.152	-0.674	0.847	0.613	0.021
N	874	782	566	892	899	875	776	7,446

Notes: Results of the regression of firm performance measures (yearly observations from 2011 to 2015, column 1 to 7) and monthly standard deviation of stock prices (monthly observations, column 8) on the lagged percentage of female directors (observations from 2010 to 2014). Panel B-bis shows the ITT specification, and Panel C-bis the IV regression. Year fixed effects are included for columns 1 to 7 and month fixed effects for column 8. We include dummies for the election years (months in column 8). The ITT specification in Panel B-bis also includes an indicator variable for the phase-in period. The first stage of the IV regression is reported in Table A4. Standard errors are clustered at the firm level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.0.4: IV regression in Table B.0.3: first stage

	(1) Percentage of women
Reform	3.509***
	(0.973)
Reform × Distance from threshold	0.699***
	(0.054)
Phase-in	2.929***
	(1.136)
Year of election FE	Y
Firm FE	Y
Year FE	Y
Observations	1,047
Mean of dep. var.	13.381
F	178.46

Notes: First stage regression for the IV regression in Table B.0.3:. Yearly observations between 2010 and 2014. Standard errors are clustered at the firm level. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

Table B.0.5: Daily average excess returns around the event date

A. June 28, 2011

	(1)	(2)	(3)	
	All firms	Above median distance	Below median distance	(2)-(3)
	-0.0011	-0.0016	-0.0002	-0.0014
	(0.0011)	(0.0014)	(0.0020)	(0.0024)
N. firms	222	103	100	203

B. March 15, 2011

	(1)	(2)	(3)	
	All firms	Above median distance	Below median distance	(2)-(3)
	0.0047***	0.0048***	0.0038***	0.0010
	(0.0010)	(0.0014)	(0.0014)	(0.0020)
N. firms	219	102	100	202

Notes: The table shows the estimated average abnormal returns over the three days surrounding the event for different portfolios of stocks. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$

Appendix C

Revisiting the Childcare Gap Between High- and Low-Skill Mothers

C.1 Data

Historical Time Use Data Time use data are from six cross-sections of the American Heritage Time Use Survey (AHTUS), covering the period 1965-2010.[1] The main sample includes mothers aged 18-64 who have completed education and whose eldest child in the household is aged less than 18.

Time Use Data I use data from the American Time Use Survey (ATUS) 2006-2017 to study the relative productivity between time of high- and low-skill mothers. The sample includes mothers of children aged less than 18. The ATUS sample is drawn from the CPS sample two months after completion of the eight CPS interview. Time diaries data are available only for one respondent in each household.

Spending Data To document changes in parental spending I use eight cross sections from the Consumer Expenditure Survey (CEX), 1980-2015. [2] The sample includes mothers aged 18-64 whose oldest child in the household is aged below 18 years and who have completed education. Total spending is adjusted for inflation to 2010 dollars and is defined as the sum of the following categories: education expenditures (tuition, books and supplies, tutoring activities and other lessons), equipment expenditures (musical instruments and sports equipment), and expenditures on books and magazines. The exact breakdown of the expenses by year is

[1] These include the years 1965, 1975, 1985, 1995, 2005, and 2010.
[2] These include years 1980, 1985, 1990, 1995, 2000, 2005, 2010, 2015.

shown below.

Immigration Data I use data from the American Community Survey (ACS) 2006-2017 and Census data for 1990 to construct yearly immigration measures and the Bartik instrument. I merge immigration data with the ATUS at the county level. Whenever the county is not identifiable either in the ATUS or in the Census, I merge the data at the metropolitan area level. Low-skill immigrants are defined as foreign-born individuals who have not completed high-school. When calculating the intensity of low-skill immigration over total population, I only consider individuals aged 25-54. To construct ancestry shares in 1990 – the “share" part of the instrument – I use information on ancestry whenever there exists a one-to-one mapping with country of origin. For the “shift" part of the instrument, I use information on country of origin and include individuals that migrated to the US between the previous decade and the current year.

Beliefs Data Data on mothers' beliefs are from the Child Development Supplement of the Panel Study of Income Dynamics. I focus on mothers of children aged under 19.

C.2 Methodology

C.2.1 Time

I distinguish between three distinct periods in the lifecycle of the child; these correspond to ages 0-5 (infants and kindergarten years), ages 6-10 (primary school), and ages 11-16 (middle school and high school). I analyze how parental investments evolved for different generations of mothers and children, with the idea that whenever a child is born, a new mother is also born. Therefore, I treat mothers of children born in different periods separately, and group cohorts of mothers/children in 5-year intervals.

The AHTUS provides information on the exact age of the youngest child in the household. Therefore, pseudo-cohorts can only be build based on the age of the youngest child in the household.[3] Consistently with previous literature[4], I define as parental time the time spent in active interaction with the child (e.g. playing, doing sports with the child, reading and talking to the child, helping with homework), the time spent taking (active) basic care of the child (e.g. feeding, medical care), and

[3] Price (2008) shows that parents spend on average less time with children of higher birth orders (i.e. those who are born later). Therefore, the analysis may be underestimating total parental time in levels. However, to the extent that higher and lower educated parents reduce time investments on children of higher birth orders proportionally, the results on parental gaps should not differ much when looking at children of lower birth orders. As a robustness check, I run the analysis for one-child families only. The results are similar, especially for the latest cross-sections. For the early cross-sections the estimates are very noisy due to very small sample sizes.

[4] See Ramey and Ramey (2010), and Guryan et al. (2008).

chauffering time. A complete list of all the activities included is shown in Table **??** in the Appendix.

I create composition-adjusted means to account for changes in parental hours that are due to the changing composition of the sample rather than mothers' behavior. In particular, I keep the age composition of mothers constant over time; very similar results are obtained if in addition I keep constant the share of mothers with children younger than 5 and children older than 5.

I divide mothers into three groups according to the age of their youngest child. I run the following regression separately for mothers whose youngest child belongs to age group g (g = 0-5, 6-10, 11-16):

$$\text{Total Childcare}_{it} = \alpha + \beta_1 \text{tert3}_{it} + \beta_2 \text{tert2}_{it} + \sum_{j \neq 1980} \gamma_j \cdot I(t=j) + \sum_{j \neq 1980} \delta_{1t}[tert3_{it} \times I(t=j)] + \sum_{j \neq 1980} \delta_{2j}[tert2_{it} \times I(t=j)] + \sum_{l} \eta_l \cdot I(a=l) + \sum_{j \neq 1980} \sum_{l} \eta_{jl}[I(a=l) \times I(t=j)] + \epsilon_{it} \quad \text{(C.1)}$$

where tert3_{it} and tert2_{it} represent respectively indicator variables for whether mother i belongs to the the third and the second tercile of the education distribution; $I(t=j)$ represents a year dummy, and $I(a=l)$ represents an indicator variable for whether mother i's age belongs to age group a (a=18-24, 25-34, 35-44,

45-54, 55-64). For every child's age group g, average childcare hours are predicted in every year for mothers of each education tercile and each age group a; the year×education×age cells are then aggregated using a fixed set of weights for the mother's age.

I use a richer specification to document the evolution of *gaps* in parental hours for children of mothers with different educational attainment. Instead of running separate regressions for mothers of children belonging to different age groups, I allow every coefficient to be cross-section specific.[5] Therefore, in each cross section I run the following regression:

$$\begin{aligned}\text{Total Childcare}_i = \alpha + \beta\,\text{tert2}_i + \gamma_1\,\text{age 6-10}_i + \gamma_2\,\text{age 11-16}_i + \\ \eta_1\,\text{tert2}_i \times \text{age 6-10}_i + \eta_2\,\text{tert2}_i \times \text{age 11-16}_i \\ + \delta_1\,\text{tert3}_i + \delta_2\,\text{tert3}_i \times \text{age 6-10}_i + \delta_3\,\text{tert3}_i \times \text{age 11-17}_i + x_i'\eta + \epsilon_i \quad \text{(C.2)}\end{aligned}$$

where i is the index for the mother, age 6-10$_i$ is a dummy variable for whether the mother's youngest child is aged 6-10, age 11-16$_i$ is a dummy for ages 11-16, and tert2$_i$ and tert3$_i$ represent respectively dummy variables for the second and third

[5] The main advantage of this methodology is that I can avoid dropping from the regression control variables that are present in the AHTUS in certain years but absent in others. For example, the race variable is absent in 1985, marital status, the number of children older than 5, and the number of children older than 5 are absent in 1995.

tercile of the education distribution. x_i is a vector of controls, which includes the mother's age category, marital status, race, the number of children aged less than 5, number of children older than 5, the employment status (full time, part time, not employed), and the sample year (if the cross-section is obtained by pooling together more than one sample year). In this regression, the coefficients of interest are represented by δ_1, δ_2, and δ_3. These represent the gaps in parental hours between mothers at the top and at the bottom of the education distribution in each cross section for children aged 0-5 (δ_1), children aged 6-10 ($\delta_1 + \delta_2$), and children aged 11-16 ($\delta_1 + \delta_3$).

C.2.2 Spending

I obtain composition-adjusted means to show the evolution of parental spending over time. The CEX provides information on the level of education of both parents, as well as the age of each member in the household. For consistency with the AHTUS, I construct pseudo-cohorts of children who are the youngest in their household. Similarly to the AHTUS, I divide mothers into three groups according to the age of their youngest child in the household. I run the following regression separately for households in which the youngest child is in age group g (g = 0-5,

6-10, 11-16):

$$\text{Total Spending}_{ht} = \alpha + \beta_1 \text{tert3}_{ht} + \beta_2 \text{tert2}_{ht} + \sum_{j \neq 1980} \gamma_j \cdot I(t=j) +$$
$$\sum_{j \neq 1980} \delta_{1t}[tert3_{ht} \times I(t=j)] + \sum_{j \neq 1980} \delta_{2j}[tert2_{ht} \times I(t=j)] +$$
$$\sum_{l} \zeta_l \cdot I(a=l) + \sum_{j \neq 1980} \sum_{l} \zeta_{jl}[I(a=l) \times I(t=j)] +$$
$$\eta_1 \log(\text{income})_{ht} + \sum_{j \neq 1980} \eta_{2j}[\log(\text{income})_{ht} \times I(t=j)] + \epsilon_{ht} \quad \text{(C.3)}$$

where tert3_{ht} and tert3_{ht} represent respectively dummy variables for whether the mother in household h at time t belongs to the top tercile or the middle tercile of the education distribution in year t; $I(t=j)$ represent a dummy for year $t=j$, $I(a=l)$ represent an indicator variable for whether the mother's age belongs to group a (a=18-24, 25-34, 35-44, 45-54, 55-64), and $\log(\text{income})_{ht}$ represents the log of family income in household h at time t. Average spending is predicted in every year for each education tercile and each age group, holding income constant at the median in the cross-section. The year×education×age cells are then aggregated using a fixed set of weights for the mother's age.

I use a richer specification in order to document $gaps$ in spending between households according to the mother's level of education. In particular, I add covariates and pairwise interactions between the covariates and the calendar year to the spec-

ification in equation C.3, where the covariates include the mother's race, employment status, marital status, and number of children.

C.3 Additional Figures

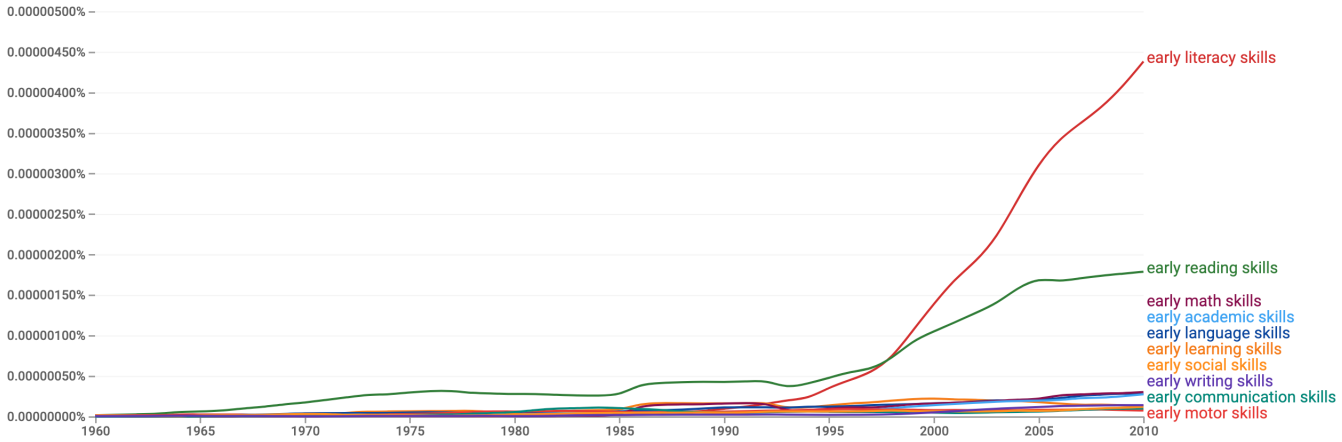

Figure C.3.1: early – *word* – skills

Notes: Source: Google's Books Ngram Viewer. Frequency of the top 10 substitutions for *word*. Books predominantly in the English language that were published in the United States, 1960-2010.

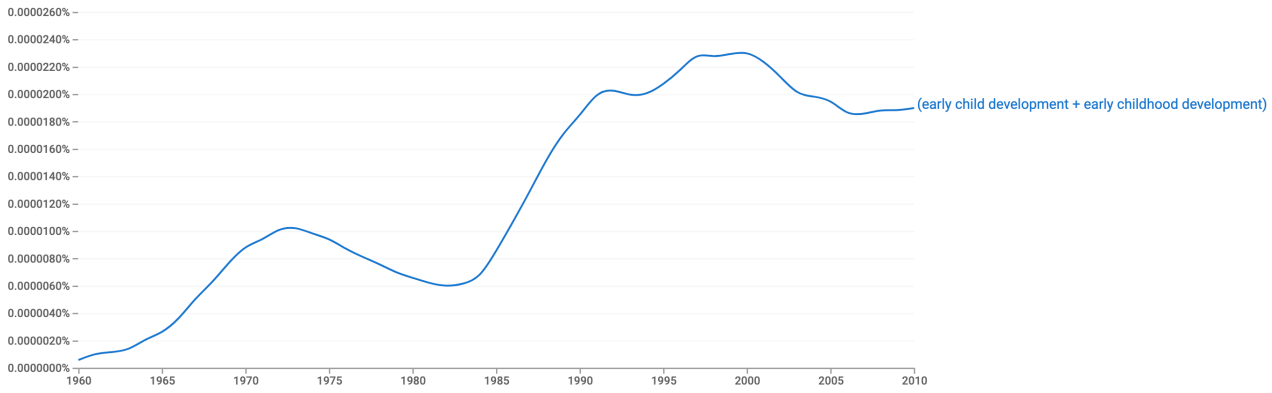

Figure C.3.2: "early child development" – "early childhood development"

Notes: Source: Google's Books Ngram Viewer. Frequency of *early childhood development* and *early child development*.
Books predominantly in the English language that were published in the United States, 1960-2010.

Bibliography

Renee Adams and Daniel Ferreira. Diversity and incentives in teams: Evidence from corporate boards. *Unpublished working paper*, 2004.

Renée B Adams and Daniel Ferreira. Women in the boardroom and their impact on governance and performance. *Journal of financial economics*, 94(2):291–309, 2009.

Renée B Adams and Patricia Funk. Beyond the glass ceiling: Does gender matter? *Management Science*, 58(2):219–235, 2012.

Renée B Adams and Vanitha Ragunathan. Lehman sisters. *Available at SSRN 3046451*, 2017.

Renée B Adams, Stephen Gray, and John Nowland. Does gender matter in the boardroom? evidence from the market reaction to mandatory new director an-

nouncements. *Evidence from the Market Reaction to Mandatory New Director Announcements (November 2, 2011)*, 2011.

Francesco Agostinelli and Matthew Wiswall. Estimating the technology of children's skill formation'. NBER Working Paper No. 22442, 2016.

Kenneth R Ahern and Amy K Dittmar. The changing of the boards: The impact on firm valuation of mandated female board representation. *The Quarterly Journal of Economics*, 127(1):137–197, 2012.

Stefania Albanesi, Claudia Olivetti, and Maria José Prados. Gender and dynamic agency: Theory and evidence on the compensation of top executives. In Solomon W. Polachek, Konstantinos Tatsiramos, and Klaus F. Zimmermann, editors, *Gender in the Labor Market*, volume 42 of *Research in Labor Economics*, pages 1–59. Emerald Publishing Ltd, January 2015.

Mario Daniele Amore, Orsola Garofalo, and Alessandro Minichilli. Gender interactions within the family firm. *Management Science*, 60(5):1083–1097, 2014.

Orazio Attanasio, Teodora Boneva, and Christopher Rauh. Parental beliefs about returns to different types of investments in school children. NBER Working Paper No. 25513, January 2019.

Manuel Bagues, Mauro Sylos-Labini, and Natalia Zinovyeva. Does the gender composition of scientific committees matter? *American Economic Review*, 107(4): 1207–38, 2017.

Manuel F Bagues and Berta Esteve-Volart. Can gender parity break the glass ceiling? Evidence from a repeated randomized experiment. *The Review of Economic Studies*, 77(4):1301–1328, 2010.

Scott R. Baker, Nicholas Bloom, and Steven J. Davis. Measuring economic policy uncertainty. *The Quarterly Journal of Economics*, 131(4):1593–1636, 2016.

Audinga Baltrunaite, Piera Bello, Alessandra Casarico, and Paola Profeta. Gender quotas and the quality of politicians. *Journal of Public Economics*, 118:62–74, 2014.

Itzhak Ben-David, John R. Graham, and Campbell R. Harvey. Managerial miscalibration. *The Quarterly Journal of Economics*, 128(4):1547–1584, 2013.

Peter Bergman, Danielle Li, and Lindsey Raymond. Hiring as exploration. *Available at SSRN 3630630*, 2020.

Marianne Bertrand. New perspectives on gender. In *Handbook of labor economics*, volume 4, pages 1543–1590. Elsevier, 2011.

Marianne Bertrand. The glass ceiling. Working Paper No. 2018-38, University of Chicago, June 2018.

Marianne Bertrand and Kevin F. Hallock. The gender gap in top corporate jobs. *ILR Review*, 55(1):3–21, 2001.

Marianne Bertrand and Antoinette Schoar. Managing with style: The effect of managers on firm policies. *The Quarterly Journal of Economics*, 118(4):1169–1208, 2003.

Marianne Bertrand, Sandra E Black, Sissel Jensen, and Adriana Lleras-Muney. Breaking the glass ceiling? The effect of board quotas on female labour market outcomes in Norway. *The Review of Economic Studies*, 86(1):191–239, 2019.

Timothy Besley and Andrea Prat. Handcuffs for the grabbing hand? Media capture and government accountability. *American Economic Review*, 96(3):720–736, 2006.

Timothy Besley, Olle Folke, Torsten Persson, and Johanna Rickne. Gender quotas and the crisis of the mediocre man: Theory and evidence from Sweden. *American economic review*, 107(8):2204–42, 2017.

Magda Bianco, Angela Ciavarella, and Rossella Signoretti. Women on corporate

boards in Italy: The role of family connections. *Corporate Governance: An International Review*, 23(2):129–144, 2015.

Nicholas Bloom. Fluctuations in uncertainty. *Journal of Economic Perspectives*, 28 (2):153–76, 2014.

Øyvind Bøhren and Siv Staubo. Does mandatory gender balance work? Changing organizational form to avoid board upheaval. *Journal of Corporate Finance*, 28: 152–168, 2014.

Teodora Boneva and Christopher Rauh. Parental beliefs about returns to educational investments— the later the better? *Journal of the European Economic Association*, 16(6):1669–1711, 2018.

Pedro Bordalo, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer. Diagnostic expectations and stock returns. *The Journal of Finance*, 74(6):2839–2874, 2019.

Jean-Philippe Bouchaud, Philipp Krueger, Augustin Landier, and David Thesmar. Sticky expectations and the profitability anomaly. *The Journal of Finance*, 74(2): 639–674, 2019.

John Bound, Brad Hershbein, and Bridget Terry Long. Playing the admissions game: Student reactions to increasing college competition. *Journal of Economic Perspectives*, 23(4):119–146, 2009.

Julia Cagé, Nicolas Hervé, and Marie-Luce Viaud. The production of information in an online world: Is copy right? *Available at SSRN 2672050*, 2019.

Ryan Chahrour, Kristoffer, and Stefan Pitschner. Sectoral media focus and aggregate fluctuations. Boston College Working Papers in Economics, Boston College Department of Economics, October 2019.

Larelle Chapple and Jacquelyn E Humphrey. Does board gender diversity have a financial impact? Evidence using stock portfolio performance. *Journal of business ethics*, 122(4):709–723, 2014.

Pierre-André Chiappori, Bernard Salanié, and Yoram Weiss. Partner choice, investment in children, and the marital college premium. *American Economic Review*, 107(8):2109–2167, 2017.

J Ignacio Conde-Ruiz, Juan-José Ganuza, and Paola Profeta. Statistical discrimination and the efficiency of quotas. *Documento de Trabajo*, page 04, 2017.

Flavio Cunha and James Heckman. The technology of skill formation. *American Economic Review*, 97(2):31–47, 2007.

Flavio Cunha, James J. Heckman, Lance Lochner, and Dimitry V. Masterov. Interpreting the evidence on life cycle skill formation. In Eric A. Hanushek and Finis

Welch, editors, *Handbook of the Economics of Education*, volume 1, pages 698–805. 2006 Elsevier B.V., 2006.

Flavio Cunha, James J. Heckman, and Susanne M. Schennach. Estimating the technology of cognitive and noncognitive skill formation. *Econometrica*, 78(3):883–931, May 2010.

Diego Daruich. The macroeconomic consequences of early childhood development policies. March 2019.

Ruth Mateos De Cabo, Siri Terjesen, Lorenzo Escot, and Ricardo Gimeno. Do "soft law" board gender quotas work? Evidence from a natural experiment. *European Management Journal*, 37(5):611–624, 2019.

Daniela Del Boca, Christopher Flinn, and Matthew Wiswall. Household choices and child development. *Review of Economic Studies*, 81:137–185, 2014.

Aaron A Dhir. *Challenging boardroom homogeneity: Corporate law, governance, and diversity*. Cambridge University Press, 2015.

Frank Dobbin and Jiwook Jung. Corporate board gender diversity and stock performance: The competence gap or institutional investor bias. *NCL Rev.*, 89:809, 2010.

B Espen Eckbo, Knut Nygaard, and Karin S Thorburn. Valuation effects of Norway's board gender-quota law revisited. *Management Science*, 2021.

Bernard Enjolras and Karl Henrik Sivesind. *Firms, boards and gender quotas: Comparative perspectives*. Emerald Group Publishing, 2012.

Kathleen A Farrell and David A. Whidbee. Monitoring by the financial press and forced CEO turnover. *Journal of Banking & Finance*, 26(12):2249–2276, 2002.

Raquel Fernández. Cultural change as learning: The evolution of female labor force participation over a century. *American Economic Review*, 103(1):472–500, 2013.

Raquel Fernández, Alessandra Fogli, and Claudia Olivetti. Mothers and sons: Preference formation and female labor force dynamics. *The Quarterly Journal of Economics*, 119(4):1249–1299, 2004.

Daniel Ferreira. Board diversity: Should we trust research to inform policy? *Corporate Governance: An International Review*, 23(2):108–111, 2015.

Eliezer Fich and Lawrence J. White. CEO compensation and turnover: The effects of mutually interlocked boards. *Wake Forest L. Rev.*, 38:935, 2003.

Luca Flabbi, Mario Macis, Andrea Moro, and Fabiano Schivardi. Do female exec-

utives make a difference? The impact of female leadership on gender gaps and firm performance. *The Economic Journal*, 129(622):2390–2423, 2019.

Stefano Gagliarducci and M Daniele Paserman. The effect of female leadership on establishment and employee outcomes: Evidence from linked employer-employee data. In *Gender convergence in the labor market*. Emerald Group Publishing Limited, 2015.

Vincenzo Galasso and Tommaso Nannicini. Competing on good politicians. *American political science review*, 105(1):79–99, 2011.

Kevin Gaughan and Edward Bishop Smith. Better in the shadows? Media coverage and market reactions to female CEO appointments. In *Academy of Management Proceedings*, volume 2016, page 13997. Academy of Management Briarcliff Manor, NY 10510, 2016.

George-Levi Gayle, Limor Golan, and Robert A. Miller. Gender differences in executive compensation and job mobility. *Journal of Labor Economics*, 30(4):829–872, 2012.

Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. Expectations and investment. *NBER Macroeconomics Annual*, 30(1):379–431, 2016.

Matthew Gentzkow and Jesse M Shapiro. Media bias and reputation. *Journal of Political Economy*, 114(2):280–316, 2006.

Mariassunta Giannetti and Mengxin Zhao. Board ancestral diversity and firm-performance volatility. *Journal of Financial and Quantitative Analysis*, 54(3): 1117–1155, 2019.

Robert Gibbons and Kevin J Murphy. Relative performance evaluation for chief executive officers. *ILR Review*, 43(3):30–S, 1990.

Vidhan K Goyal and Chul W Park. Board leadership structure and CEO turnover. *Journal of Corporate Finance*, 8(1):49–66, 2002.

Lorenz Graf-Vlachy, Abbie Griffith Oliver, Richard Banfield, Andreas König, and Jonathan Bundy. Media coverage of firms: Background, integration, and directions for future research. *Journal of Management*, 46(1):36–69, 2020.

Robin Greenwood and Andrei Shleifer. Expectations of returns and expected returns. *The Review of Financial Studies*, 27(3):714–746, 2014.

Vishal K. Gupta, Sandra C. Mortal, Sabatino Silveri, Minxing Sun, and Daniel B. Turban. You're fired! Gender disparities in CEO dismissal. *Journal of Management*, 46(4):560–582, 2020.

Jonathan Guryan, Erik Hurst, and Melissa Schettini Kearney. Parental education and parental time with children. *Journal of Economic Perspectives*, 22(3):23–46, 2008.

Tony Harcup and Deirdre O'Neill. What is news? *Journalism Studies*, 18(12): 1470–1488, 2017.

David E. Harrington. Economic news on television: The determinants of coverage. *The Public Opinion Quarterly*, 53(1):17–40, 1989.

Mitchell Hoffman and Stephen V Burks. Worker overconfidence: Field evidence and implications for employee turnover and firm profits. *Quantitative Economics*, 11 (1):315–348, 2020.

Harry J Holzer and David Neumark. Affirmative action: What do we know? *Journal of policy analysis and management*, 25(2):463–490, 2006.

Sander Hoogendoorn, Hessel Oosterbeek, and Mirjam Van Praag. The impact of gender diversity on the performance of business teams: Evidence from a field experiment. *Management Science*, 59(7):1514–1528, 2013.

Chang-Tai Hsieh, Erik Hurst, Charles I. Jones, and Peter J. Klenow. The allocation of talent and US economic growth. *Econometrica*, 87(5):1439–1474, 2019.

Morten Huse and Cathrine Seierstad. Getting women on to corporate boards: Consequences of the Norwegian gender balance law. *The European Financial Review*, 12:37–39, 2013.

Dirk Jenter and Fadi Kanaan. CEO turnover and relative performance evaluation. *The Journal of Finance*, 70(5):2155–2184, 2015.

Dirk Jenter and Katharina Lewellen. Performance-induced CEO turnover. *Available at SSRN 1570635*, 2019.

Nancy Ammon Jianakoplos and Alexandra Bernasek. Are women more risk averse? *Economic inquiry*, 36(4):620–630, 1998.

Boyan Jovanovic. Job matching and the theory of turnover. *Journal of Political Economy*, 87(5, Part 1):972–990, 1979.

Steven N. Kaplan and Bernadette A. Minton. How has CEO turnover changed? *International Review of Finance*, 12(1):57–87, 2012.

Wolfgang Keller, Teresa Molina, and William W Olney. The gender gap among top business executives. Technical report, National Bureau of Economic Research, 2020.

Bryan Kelly, L'uboš Pástor, and Pietro Veronesi. The price of political uncertainty:

Theory and evidence from the option market. *The Journal of Finance*, 71(5): 2417–2480, 2016.

Sagar P Kothari and Jerold B Warner. Econometrics of event studies. In *Handbook of empirical corporate finance*, pages 3–36. Elsevier, 2007.

Annette Lareau. *Unequal Childhoods*. University of California Press, second edition, September 2011.

Volker Laux. Board independence and CEO turnover. *Journal of Accounting Research*, 46(1):137–171, 2008.

Peggy M. Lee and Erika Hayes James. She'-e-os: gender effects and investor reactions to the announcements of top executive appointments. *Strategic Management Journal*, 28(3):227–241, 2007.

Shelly Lundberg and Robert A. Pollak. Cohabitation and the uneven retreat from marriage in the united states, 1950–2010. In Leah Platt Boustan, Carola Frydman, and Robert A. Margo, editors, *Human Capital in History: The American Record*, pages 241–272. University of Chicago Press, 2014.

Yueran Ma, Tiziano Ropele, David Sraer, and David Thesmar. A quantitative analysis of distortions in managerial forecasts. Technical report, National Bureau of Economic Research, 2020.

David A Matsa and Amalia R Miller. A female style in corporate leadership? Evidence from quotas. *American Economic Journal: Applied Economics*, 5(3):136–69, 2013.

Kevin J Murphy. Executive compensation. *Handbook of Labor Economics*, 3:2485–2563, 1999.

Kristoffer P. Nimark. Man-bites-dog business cycles. *American Economic Review*, 104 (8):2320–67, August 2014.

Kristoffer P. Nimark and Stefan Pitschner. News media and delegated information choice. *Journal of Economic Theory*, 181(C):160–196, 2019.

Knut Nygaard. Forced board changes: Evidence from Norway. *NHH Dept. of Economics Discussion Paper*, (5), 2011.

Rohini Pande and Deanna Ford. Gender quotas and female leadership. 2012.

Robert Parrino. CEO turnover and outside succession a cross-sectional analysis. *Journal of Financial Economics*, 46(2):165–197, 1997.

Maria Petrova. Inequality and media capture. *Journal of Public Economics*, 92(1-2): 183–212, 2008.

Joseph Price. Parent-child quality time: Does birth order matter? *The Journal of Human Resources*, 43(1):240–265, 2008.

Paola Profeta, Livia Amidani Aliberti, Alessandra Casarico, Marilisa D'Amico, and Anna Puccio. *Women directors: The Italian way and beyond*. Springer, 2014.

Garey Ramey and Valerie Ramey. The rug rat race. *Brookings Papers on Economic Activity*, 2010.

Caspar Rose. Does female board representation influence firm performance? The Danish evidence. *Corporate governance: An international review*, 15(2):404–413, 2007.

Fabrizio Rossi and Richard J Cebula. Stock market reactions to announcements of the board of directors: Evidence from italy. *Applied Economics*, 47(20):2102–2118, 2015.

Michelle K Ryan and S Alexander Haslam. The glass cliff: Evidence that women are over-represented in precarious leadership positions. *British Journal of management*, 16(2):81–90, 2005.

Liana C. Sayer, Suzanne M. Bianchi, and John P. Robinson. Are parents investing less in children? Trends in mothers' and fathers' time with children. *American Journal of Sociology*, 110(1):1–43, July 2001.

Miriam Schwartz-Ziv. Gender and board activeness: The role of a critical mass. *Journal of Financial and Quantitative Analysis*, 52(2):751–780, 2017.

Cathrine Seierstad and Tore Opsahl. For the few not the many? The effects of affirmative action on presence, prominence, and social capital of women directors in Norway. *Scandinavian journal of management*, 27(1):44–54, 2011.

Pamela J. Shoemaker and Timothy Vos. *Gatekeeping Theory*. Routledge, 2009.

Silvia Solimene, Daniela Coluccia, and Stefano Fontana. Gender diversity on corporate boards: An empirical investigation of Italian listed companies. *Palgrave Communications*, 3(1):1–7, 2017.

Stuart N. Soroka. The gatekeeping function: Distributions of information in media and the real world. *The Journal of Politics*, 74(2):514–528, 2012.

Annika E Sunden and Brian J Surette. Gender differences in the allocation of assets in retirement savings plans. *The American Economic Review*, 88(2):207–211, 1998.

Lucian A. Taylor. Why are CEOs rarely fired? Evidence from structural estimation. *The Journal of Finance*, 65(6):2051–2087, 2010.

Marko Terviö. Superstars and mediocrities: Market failure in the discovery of talent. *The Review of Economic Studies*, 76(2):829–850, 2009.

Marc Trussler and Stuart Soroka. Consumer demand for cynical and negative news frames. *The International Journal of Press/Politics*, 19(3):360–379, 2014.

Nicholas Van der Walt and Coral Ingley. Board dynamics and the influence of professional background, gender and ethnic diversity of directors. *Corporate governance: An international review*, 11(3):218–234, 2003.

Mingzhu Wang and Elisabeth Kelan. The gender quota and female leadership: Effects of the norwegian gender quota on board chairs and CEOs. *Journal of business ethics*, 117(3):449–466, 2013.

Michael S. Weisbach. Outside directors and CEO turnover. *Journal of Financial Economics*, 20:431–460, 1988.

Justin Wolfers. Diagnosing discrimination: Stock returns and CEO gender. *Journal of the European Economic Association*, 4(2-3):531–541, 2006.