Three Essays in Corporate Finance

by

Emil Lakkis

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business Administration)
in The University of Michigan
2023

Doctoral Committee:

Professor Uday Rajan, Chair
Associate Professor Nadya Malenko
Professor Amiyatosh Purnanandam
Professor Gabriel Rauterberg
Professor Christopher Williams

Emil Lakkis

elakkis@umich.edu

ORCID iD: 0000-0001-9802-6142

ACKNOWLEDGEMENTS

I am grateful to my advisor Uday Rajan for his continuous support throughout the years. He was always generous with his time and was extremely helpful discussing new ideas and existing papers. His guidance helped me find many interesting research questions and provide answers to some of them. Nadya Malenko read countless versions of my drafts and taught me how to clearly explain my thoughts in writing. Her enthusiasm and support helped me tremendously during some of the more difficult years in the program. I am thankful to Amiyatosh Purnanandam for teaching me how to conduct rigorous empirical research. Our conversations helped me improve my work and see the bigger picture. I am also thankful to Gabriel Rauterberg, Christopher Williams, MP Narayanan, Sugato Bhattacharyya, Andrey Malenko, and Mirela Sandulescu. Their feedback and thoughts on my papers helped me tremendously.

I would also like to thank my fellow PhD students, particularly Thomas Flanagan, Jay Kahn, George Malikov, Shodi Saliev, Taeuk Seo, Virginia Traweek, and Alex Wirth. Finally, I would like to thank my family for always supporting me throughout these years.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

In my dissertation, I study how the structure of corporate ownership affects financial policies of firms.

In Chapter 1, I show that labor mobility affects the timing of a firm's decision to go public. Employees of private firms compensated with stock options face high separation costs. They have to exercise or forfeit their option grants shortly after termination, which can force them to stay with the firm until its IPO. As a result, the firm has incentives to delay its IPO to retain the workers. Using state-level changes in the enforceability of non-compete agreements, I show that a decrease in labor mobility pushes firms to go public earlier. The results are stronger for firms with higher levels of stock option compensation and firms whose options are costlier to exercise. Overall, my findings emphasize the importance of frictions in labor contracting in firms' financing decisions.

In Chapter 2, I study whether ESG mutual funds affect environmental policies at their portfolio companies. I use a change in the methodology of the MSCI ESG Leaders Index as a source of quasi-exogenous variation in levels of ESG holdings at firms. I show that firms excluded from the index are more likely to open new plants and increase their toxic releases. I present evidence that the result is driven by lower monitoring efforts by institutional investors. When a fund family's ESG funds divest from a firm, the fund family has lower incentives to influence environmental policies at the firm. The paper highlights that the composition of a firm's ownership has real effects on its environmental policies.

In Chapter 3, I propose a novel mechanism that explains why index funds may have a beneficial effect on corporate governance: Increased index fund holdings increase a fund family's incentives to monitor management. Fund families centralize voting and may have incentives to monitor even when their index funds do not. Increased index fund holdings

facilitate the coordination between funds from the same fund family and lead to fewer pro-management votes cast by funds in the same family. My results demonstrate that, when families offer both index and active funds, the growth of index funds does not lead to less monitoring of the management.

Overall, in my dissertation, I present evidence that the structure of corporate ownership has important effects on the policies of firms. In Chapter 1, I demonstrate that employee stock option holdings affect a firm's IPO decision. In Chapter 2, I show that index funds can alleviate agency frictions at firms they hold. Finally, in Chapter 3, I show that ESG funds can help reduce toxic pollution at the firms they hold in their portfolios.

Chapter 1

Human Capital and IPO Timing

1.1 Introduction

> "Gabriel Cole, who worked in Airbnb's food department, said he had spent his life savings to buy his stock after he left the company in 2015. That incurred a $180,000 tax bill, which he couldn't afford, he said. 'I was returning bottles to buy groceries', he said."
>
> — New York Times, 2019

Two important trends in the US corporate sector in recent decades have been the decline in the number of public firms and the growing importance of human capital-intensive firms. Since the 1990s, the number of publicly listed firms in the US has decreased by 49% (Doidge, Karolyi, and Stulz, 2017), and the number of initial public offerings (IPOs) has decreased almost threefold (Doidge, Karolyi, and Stulz, 2013; Gao, Ritter, and Zhu, 2013). Further, Ewens and Farre-Mensa (2020) document that firms stay private longer: since the 1990s, the median time between the first round of VC funding and an IPO has increased by three years. Over the same period, human capital has become an increasingly important factor of production.

A key feature that distinguishes human capital from physical capital is its inalienability (Hart and Moore, 1994). Employees can leave the firm at any time, and it is difficult to write long-term contracts that ensure human capital availability in the future. When a worker expects to leave the firm, both the worker and the firm have little incentive to invest in firm-specific human capital. Additionally, the firm would have to incur labor market search costs to replace the worker. For these reasons, employee retention is an important concern for firms, and they utilize several devices, such as deferred compensation, stock-based compensation, and non-compete agreements to reduce labor turnover.

In this paper, I propose that delaying the IPO is another tool that a private firm can use to retain its workers. Leaving a private firm can be particularly costly for employees compensated with stock options because they cannot finance the exercise of the options with the sale of the underlying stock. This effectively binds employees to the firm until it goes public. Therefore, especially when labor is highly mobile, the firm has incentives to delay its IPO to retain the workers. In this paper, I investigate this channel in more depth and study how labor mobility affects the timing of a firm's IPO. My main finding is that when labor becomes less mobile, the firm goes public earlier.

The reason why leaving a private firm before the IPO can be very costly is that an employee must exercise or forfeit the stock options within 90 days after leaving the firm. Such exercise can require a large amount of cash since the employee has to pay both the strike price and the associated taxes, which can be substantial.[1] The main difference between an employee at a private versus a public firm is in her ability to finance these costs. If a firm is public, its stock is liquid, and the employee can sell it to finance the costs of option exercise. In contrast, the stock of the private firm is illiquid, and it may be very difficult and sometimes impossible to sell before the firm goes public.[2] Such costs can force the employee to stay with the firm until the IPO. In turn, the firm concerned with employee retention can find it optimal to delay the IPO in order to limit the turnover of its employees. When labor

[1] According to the SECFI Annual Report (2021), the average SECFI's client's cost to exercise stock options is $543K (with 27% of this amount attributed to paying the strike price and 73% to the exercise taxes), while the annual household income is $270K. See also "Inside Airbnb, Employees Eager for Big Payouts Pushed It to Go Public" (New York Times, Sep 20, 2019).

[2] In recent years, a number of marketplaces, such as SharesPost, Equidate, EquityZen, and Nasdaq Private Market have been developed in order to facilitate secondary exchanges in shares of private firms. While these platforms can provide some liquidity to employees of private firms, the firms can still put restrictions on such transactions. See Larcker, Tayan, and Watts (2018) for a detailed discussion.

mobility decreases, the firm can go public sooner since the listing status has less of an effect on the turnover of workers.

Empirically examining the effect of labor mobility on a firm's IPO timing is challenging for two reasons. First, labor mobility is difficult to measure since an employee's outside options and switching costs are not observable. Second, labor mobility and a firm's decision of when to go public can be driven by common confounding factors. In particular, the probability of a firm's success affects the employees' incentives to stay with the firm and the time it takes the firm to go public.

To overcome the above issues, I rely on state-level changes in the enforceability of non-compete agreements as a source of quasi-exogenous variation in labor mobility. Non-compete agreements (NCAs) are contractual clauses in employment contracts that limit an employee's ability to seek employment at competitors or start a new business. Such agreements are common (Starr, Prescott, and Bishara, 2021), and they significantly impede the mobility of workers (Marx, Strumsky, and Fleming, 2009; Jeffers, 2019). Importantly for my analysis, most changes in their enforceability occur as a result of state supreme court rulings and are unlikely to be motivated by the desire to affect business and IPO activity in a state. I identify changes in the enforceability of non-competes in several states and estimate their effect on firms' IPO age in a difference-in-differences framework.

I begin my empirical analysis by comparing IPOs in treated and control states. I document that IPOs in treated and control states are similar on a number of observable characteristics before the treatment. I also show that IPO age evolves similarly in treated and control states before the treatment. This finding provides evidence in support of the parallel trends assumption.

I find that following an increase in the enforceability of non-compete agreements in a state, firms headquartered in that state go public at a younger age. Comparing the age of firms at their IPO in treated and control states before and after the treatment, I find that following an increase in NCA enforceability, treated firms are 3.9 years younger when they go public. This constitutes almost a 26% decline in the IPO age, given a sample mean IPO age of 14.6 years.

The effect of a change in labor mobility on the IPO age should be more pronounced when the firm has more discretion over its IPO timing decision. In particular, firms that receive VC backing are less likely to delay the IPO due to labor retention concerns. Venture capital funds have a limited time to exit from their portfolio companies since they have to return funds to the investors within ten years. As a result, they are less likely to delay an IPO even if it maximizes the value of the firm. Consequently, I expect the age of VC-backed IPOs to be less affected by NCA shocks. I find that the treatment effect estimate is much smaller for VC-backed IPOs, as expected. An effect of NCA shock on IPO age is a 6.1-year reduction for non-VC-backed IPOs and a 2.8-year reduction for VC-backed IPOs.

Next, I perform several empirical tests to examine the two main economic mechanisms that can explain why firms go public sooner when non-competes become more enforceable. First, it is possible that firms delay their IPOs to lock in their employees compensated with stock options. In this case, when NCAs become more enforceable, labor mobility decreases, which makes delaying the IPO a less attractive tool to retain workers. I refer to this mechanism as the *stock option compensation mechanism*. The second mechanism that can explain my results is that an increase in NCA enforceability can have a positive effect on a firm's value and growth rate. This, in turn, can reduce the time it takes the firm to reach its IPO. I refer to this mechanism as the *valuation mechanism*.

I start by examining the *stock option compensation mechanism*. Under this mechanism, I expect the treatment effect to be stronger for firms that heavily use employee stock option compensation. Intuitively, if a firm uses few stock options, the costs of the stock option exercise described above are low for workers leaving the firm. In turn, delaying the IPO is unlikely to increase employee retention. An increase in NCA enforceability should then have little effect on when the firm decides to go public. Directly testing such prediction is challenging since stock option grants for private firms are not observable. To overcome this challenge, I use the number of a firm's stock options as of its first filing after the IPO, which is reported on the firm's 10-K filing. I split the sample into terciles based on the number of stock options outstanding after the IPO and estimate my model on each subsample. I find that the effect of an increase in NCA enforceability is the strongest for the subsample corresponding to the highest tercile of the number of options.

Next, I examine how the treatment effect estimates vary with the costs of option exercise. Under the stock option compensation mechanism, I expect the results to be stronger for firms whose options are costlier to exercise. Intuitively, when the costs of option exercise are low, it is easier for workers to finance the exercise, hence making staying private less effective for worker retention. Empirically, for each firm, I compute the costs of the exercise of its stock options, which consists of the exercise price and the taxes associated with the exercise. I establish two main results: first, the effect of NCA enforceability on IPO age is stronger for firms whose vested options are deep in the money and have high exercise taxes. Second, I find that my main result is driven by the subsample of firms with high total stock option exercise costs.

To provide further support for the stock option compensation mechanism, I show that firms that go public following the treatment compensate for the increased post-IPO liquidity of their shares by subjecting their employees to lock-up provisions. Lock-ups are common provisions that restrict some shareholders from selling a firm's stock for a certain time, usually 180 days, following the IPO. If an employee is subject to a lock-up provision, she may find it difficult to leave the firm immediately after the IPO since she cannot finance the option exercise with the sale of the stock, even though the firm is already public. Interestingly, I find that changes in NCA enforceability affect firms' use of lock-up provisions for their employees. These results provide additional support for the stock option compensation mechanism.

Next, I examine the *valuation mechanism*. Under this theory, when non-compete agreements become more enforceable, a firm experiences an increase in its valuation.[3] In the real option framework, this increase would lead the firm to reach the IPO sooner (Dixit and Pindyck, 1994).

[3]Whether an increase in NCA enforceability has a positive effect on a firm's value is an empirical question. Younge and Marx (2016) find that an increase in NCA enforceability has a positive effect on Tobin's Qs for treated firms. On the other hand, it is possible that a firm's value can decrease if it becomes more difficult to attract new human capital due to stronger NCA enforceability (Stuart and Sorenson, 2003).

To explore this mechanism, I first examine how NCA shocks affect firms' IPO proceeds and post-IPO market valuations. I find that following a strengthening of the enforceability of non-competes, firms raise less financing in their IPOs and are also smaller in size when they go public. Additionally, I document that an increase in the NCA enforceability does not affect the failure rate of VC-backed startups. In combination, these results cast doubt on the ability of the valuation mechanism to explain my findings.

Further, I perform an additional test to examine how NCA shocks affect the valuation of firms. Such a test is difficult to conduct for private firms, as their latest valuations are not observable. Note, however, that under the valuation mechanism, one would also expect the valuations of public firms to increase. I test this prediction by estimating the effect of NCA shocks on Tobin's Q for a sample of public firms. The main empirical challenge with this exercise is that public firms can have offices in multiple states. Thus, even firms headquartered in control states can be affected by changes in non-compete enforceability if they have offices in treated states. To alleviate this concern, I split the sample of public firms into terciles based on the number of employees, expecting that firms in the lower tercile are less likely to have locations in multiple states. I find that public firm valuations decrease after a positive shock to the enforceability of non-competes. This can be explained by the fact that smaller firms find it more difficult to attract employees away from competitors (Stuart and Sorenson, 2003).

To summarize, in a series of tests, I find strong support in favor of the stock option compensation mechanism. A decline in IPO age in response to the treatment is driven by firms with a large number of stock options and firms whose stock options are costlier to exercise. Additionally, firms increase the use of employee lock-up provisions following the IPO. In contrast, I do not find any evidence in support of the valuation mechanism.

I perform several robustness checks. I show that my results are not driven by IPO timing dynamics in any single state. They are also robust to the exclusion of states that experienced NCA changes that were a result of legislative changes, which could be driven by business interest groups. Using a sample of VC-backed startups, I also show that a change in the

NCA enforceability does not affect the startup's probability of going public versus getting acquired or going bankrupt. Additionally, I show that the timing of firms' acquisition and failure exits from VC portfolios are not affected by the non-compete enforceability changes.

The findings in this paper have several important implications. First, my findings provide an additional explanation as to why firms have been staying private longer in recent years (Doidge, Karolyi, and Stulz, 2013, 2017; Ewens and Farre-Mensa, 2020). I show that firms that heavily use employee stock option compensation optimally decide to stay private longer to retain their workers. Over time, the growing fraction of firms relying on human capital and using option compensation can contribute to an increase in the average IPO age. Second, my findings imply that private firms can have a unique advantage in accumulating human capital. Since workers at a private firm are less likely to leave, this may give both workers and firms stronger incentives to invest in firm-specific human capital, positively affecting the firm's productivity.

The paper relates to three strands of literature. First, it contributes to a vast literature documenting a firm's motives to go public. Existing studies show that firms do an IPO to raise new capital, provide liquidity (Chemmanur, He, and Nandy, 2010) and diversification (Pagano, 1993; Bodnaruk et al., 2008; Chod and Lyandres, 2011) for the initial investors, and facilitate acquisitions (Brau and Fawcett, 2006). Additionally, going public can serve as a commitment device for monitoring by public markets (Holmström and Tirole, 1993; Pagano and Röell, 1998) and regulators (Lowry, Michaely, and Volkova, 2020).[4] I contribute to this rich literature by documenting a new motive to stay private: the retention of a firm's employees. Several papers study the optimal timing of the IPO decision. Ibbotson and Jaffe (1975) and Ritter (1984) show that IPOs tend to cluster during periods of high market returns. Benveniste, Busaba, and Wilhelm (2002) and Alti (2005) argue that the IPO timing is determined by information spillovers between early and late-IPO firms. Draho (2000) and Benninga, Helmantel, and Sarig (2005) argue that the optimal IPO timing is determined by

[4]See Lowry, Michaely, and Volkova (2017) for a recent survey of the IPO literature. Brau and Fawcett (2006) provide survey evidence for the reasons firms choose to go public or stay private. Pagano, Panetta, and Zingales (1998) provide an empirical analysis of IPO motives for a sample of Italian firms.

a tradeoff between the private benefits the entrepreneur enjoys by keeping the firm private and the diversification benefits of taking it public. In this paper, I demonstrate how a firm's human capital mobility affects the optimal IPO timing.

Second, I contribute to the literature that studies the connection between the labor market and a firm's financial policies. The existing studies examine how labor mobility affects a firm's capital structure (Sanati, 2018; Ysmailov, 2020) and investment decisions (Garmaise, 2011; Sanati, 2018; Jeffers, 2019).[5] Bernstein (2015) shows that an IPO leads to an outflow of high-skilled inventors from the firm, and Babina, Ouimet, and Zarutskie (2020) show that startup employees leave the firm after it goes public. Gupta (2022) shows that firms capture significant rents when the mobility of immigrant workers declines.

The most closely related paper in this area is Bias (2022), which studies how the going public decision affects a firm's labor mobility and the allocation of talent across firms. Bias (2022) makes a similar argument that stock illiquidity could constrain mobility of employees and finds supporting evidence using stock market shocks during a firm's book-building process. The main difference between Bias (2022) and this paper is that Bias (2022) focuses on the effects of an IPO on a firm's labor force, while I focus on how the timing of a firm's IPO decision is affected by its labor. I provide evidence that firms exploit the lock-in effect of illiquid equity documented by Bias (2022) and postpone their IPO to reduce labor turnover.

Additionally, I contribute to the literature studying the effects of NCA on firms and showing its consequences in financial markets.[6] The existing studies on non-compete agreements focused on their effects on labor mobility (Marx, Strumsky, and Flemming, 2009; Marx, Singh, and Flemming, 2015; Ewens and Marx, 2018), valuations of firms (Younge and Marx, 2015), investment decisions (Garmaise, 2011; Jeffers, 2019) and business creation (Jeffers, 2019). In this paper, I provide new evidence on how non-competes affect the timing of the key financing decision, the IPO.

Finally, the paper sheds new light on the listing gap documented by Doidge, Karolyi, and

[5]See Matsa (2018), Pagano (2020), and Nishesh, Ouimet, and Simintzi (2022) for recent surveys of the large labor and finance literature.

[6]See McAdams (2019) and Starr (2019) for recent surveys of the literature on non-compete agreements.

Stulz (2013, 2017) and Gao, Ritter, and Zhu (2013) and an increase in the average age at which firms go public. The existing studies explain a decline in the number of public firms by the increased regulatory costs of being a public firm (Dambra, Field, and Gustafson, 2015) and an increase in the supply of private financing (Ewens and Farre-Mensa, 2020). In this paper, I propose a new explanation for the increase in the time it takes firms to go public. As more firms in the economy increasingly rely on human capital and use employee stock option compensation, more firms find it optimal to postpone their IPO, reducing the number of IPOs in any given year and increasing the average IPO age.

The rest of the paper is organized as follows. I describe my data sources and sample construction in Section 2. In Section 3, I outline my empirical strategy. I present my main results in Section 4. In Section 5, I explore the economic mechanism behind my findings. Section 6 contains several robustness exercises. Section 7 concludes.

1.2 Data

A. IPO Data

IPO data comes from the SDC New Issues Database. I focus on firms that went public between 2000 and 2020. I limit the sample to IPOs on the NYSE, Amex, and NASDAQ exchanges and exclude REITs, ADRs, unit offerings, and spinoffs. To obtain founding dates, I use the Field-Ritter dataset (Field and Karpoff, 2002; Loughran and Ritter, 2004). I measure a firm's IPO age as the difference between the date of the firm's IPO offering date and its founding date. Since Field-Ritter's data only includes the founding year, I assume that a firm was founded on January 1st. I link the SDC New Issues to CRSP and Compustat using matching by CUSIP and checking the resulting matches manually.

Following the above steps, I construct a sample of 2,055 IPOs between 2000 and 2020. For 1,909 of those offerings, I manage to find a 10-K form filed by a firm in the first year after its IPO. For each IPO, I collect information on the total IPO proceeds, the percent of equity sold to public investors, and the firm's market capitalization following the IPO. Further, I obtain information on a firm's industry and whether the IPO was VC-backed. I follow Loughran and Ritter (2004) in their definition of high technology and biotechnology firms. I trim observations for which the IPO age is above the 97.5th percentile.

Additionally, from the SDC New Issues, I extract information on the use of post-IPO lock-up provisions. I include information on whether such provisions were used and whether the management and the other shareholders, such as employees, were subject to these provisions.

Finally, I read the IPO prospectuses for firms in my sample to obtain information on a firm's office locations. For each firm, I record the state in which it was headquartered at the time of the IPO. In some cases, this information differs from the one reported in the SDC New Issues. I use the information from the prospectus in these cases. Further, for each firm, I collect a list of states in which the firm had offices at the time when it went public. Since I attempt to identify the location of highly skilled human capital, I do not focus on the locations of a firm's stores and warehouse facilities. In some instances, firms do not disclose the exact states in which they have operations but rather specify that they operate in a large number of states. I record such information and use it in my robustness tests.

B. VC Data

The data on VC-backed firms comes from Pitchbook. Pitchbook collects and provides detailed information about private firms, the history of VC financing, and information about key employees and the firm's management. I focus on private firms headquartered in the US that received their first VC financing round between 2000 and 2019.

For each VC-backed firm in the sample, I obtain information on its financing history, including the total number of financing rounds and the total amount of capital raised from VC investors. Additionally, I collect information on the exits of private startups. For each firm, I use the deal information from Pitchbook to identify whether it exited through an acquisition or an IPO or failed. In some cases, Pitchbook contains information on multiple exit deals for the same firm. For example, a firm can be first sold by a VC fund through acquisition and later go public through a spinoff. In such instances, I only use the first exit (acquisition in this example). I also collect the date the firm exited the VC portfolio for each exit.

Metrick and Yasuda (2011) show that for many VC-backed firms sold through acquisitions, the VC investors do not recover the capital they put in those firms, which indicates that ex-post such firms were not successful investments. To classify acquisition exits into successful and unsuccessful ones, I follow a procedure similar to the one described in Ewens and Marx (2018). For each firm, I calculate the total amount of VC financing it received over its lifetime. Then, I take the ratio of the acquisition value over the total amount invested in the firm by the VC investors. I refer to this ratio as the exit multiple. If the exit multiple is above two, I classify this as a successful exit by the VC investors.[7] For firms for which the acquisition deal values are missing, I classify them as unsuccessful acquisitions, as sizes of these deals are likely small (Puri and Zarutskie, 2012). Of 5,341 acquisitions in my sample, 1,223 are classified as successful. In a series of robustness checks, I use alternative thresholds for exit multiples between one and three. I verify that my main results are virtually unchanged if I use these alternative definitions of successful exits, such as those used by Bernstein, Giroud, and Townsend (2016).

C. Firm's Financial Data

For each firm in my sample, I obtain the financials from the first 10-K form filed following the IPO through Compustat. In my sample, the median time gap between the IPO issue and the date of the first 10-K report is 181 days (the mean is 174.9 days). I collect data on the number of stock options outstanding in the year following the IPO. Since 2002, the Securities and Exchange Commission (SEC) has required firms to disclose on their 10-K filings the total number of outstanding and exercisable options at the end of the fiscal year and the number of canceled and exercised options during the year. I also collect information on the weighted average strike price of exercised options. Additionally, I add data on a firm's total assets, market capitalization, stock compensation expenses, headquarters location, and the total number of employees. I use an industry classification based on the SIC codes.

D. Summary Statistics

[7] Bernstein, Giroud, and Townsend (2016) use a slightly different approach and define an acquisition to be a successful exit if the deal value is above $25 million. My results remain qualitatively similar if I follow their approach.

I report the summary statistics for the IPO sample Table 1.1. On average, a firm raises \$184.76 million in an IPO and goes public when it is 14.6 years old. A large fraction, 67% of IPOs, is backed by VC investors. The number of IPO is not evenly distributed across the states. Firms headquartered in California account for 32.6% of all offerings in the sample. A share of IPOs from Massachusetts and Texas are 10.9% and 7.9%, respectively.

In the year following the IPO, firms have a large number of outstanding stock options, to purchase on average 9.7% of the common equity. A large proportion of those outstanding options is exercisable, enough to acquire 6% of the outstanding equity. Firms spend approximately 1.4% of their total assets on stock compensation and 0.2% of assets on deferred compensation.

The majority of firms, 89.8%, impose some type of lock-up provisions. The most frequent group of shareholders subject to lock-ups is the management, in 66.8% of the offerings. Less frequently, in 29.3% of the IPOs, current shareholders, such as employees, are subject to lock-up restrictions.

Technology firms account for a large fraction of all IPOs, 36.1%. Another large group of offerings is from firms in the biotechnology industries, as they comprise 24.3% of all offerings.

Most firms that go public operate in only a few states. On average, a firm that goes public has offices in two states, and the median number of states in which a firm operates is one.

In Table 1.2, I report the summary statistics for a sample of VC-backed startups I use in the paper. I include the startups that received their first VC financing round between 2000 and 2008 in order to be able to observe their exits. 4% of startups went public, while 53% were acquired. 35% failed by either going bankrupt or ceasing operations, and 8% remain private as of 2020. These numbers slightly differ from those reported in earlier studies, such as Puri and Zarutskie (2012), who report a much higher proportion of firms going public and a lower fraction of firms getting acquired. The difference is likely due to the use of a more recent sample period in my paper, over which firms are less likely to go public and, instead, stay private longer (Doidge, Karolyi, and Stulz, 2013; Gao, Ritter, and Zhu, 2013; Ewens and Farre-Mensa, 2020).

On average, a VC-backed startup raises \$35.54 million in VC financing over its lifecycle. The average exit multiple, defined as the exit deal size divided by the amount of VC capital invested in the firm, is 2.1. The distribution of this variable is very skewed, and the median is equal to zero: for most firms, VC funds do not recover their investments. The fraction of VC-backed firms with the exit multiple above two is only 13.1%. The average firm age at exit is 10.64 years, which is much smaller than the average IPO age in the main sample. Such a difference is likely due to the limited lifetime that VC funds have.

1.3 Empirical Framework

In this section, I describe my empirical strategy to estimate the effect of a change in labor mobility at a firm on its IPO timing decision. In an ideal setting, one would run a regression with the age at which the firm goes public as the dependent variable and labor mobility as the independent variable. However, such an approach is not feasible since an employee's outside options and switching costs are not observable. Further, even if the mobility was observed, it is not randomly assigned to firms and correlates with other firm-level outcomes. In particular, more successful firms are more likely to go public soon, while employees are less likely to leave such firms. As a result, the omitted variable bias makes it difficult to interpret the estimates of a simple OLS regression casually.

To overcome the above challenges, I rely on a series of state-level changes in the enforceability of non-compete agreements as a source of quasi-exogenous variation in labor mobility.

1.3 Non-Compete Agreements

Non-compete agreements (NCAs, non-compete covenants, or simply non-competes) are contracts between a firm and an employee that limit the employee's post-employment ability to work at competitors or start a new business. Employers often offer such contractual agreements to their employees to limit labor turnover and incentivize investments in firm-specific human capital (Bishara, 2006).[8] In their survey, Starr, Prescott, and Bishara (2021) find that 18 percent of respondents are bound by non-competes, and 38 percent signed at least one in the past. They also document that non-competes are more common among high-skilled workers.

The enforceability of non-competes in courts is determined at the state level and varies greatly across states (Bishara, 2011). NCAs are enforced to some extent in most states, with a notable exception of California, in which such agreements are not enforceable in any form. To determine whether a given NCA is enforceable, courts seek to decide whether or not it protects legitimate business interests and is reasonably limited in duration and space, and is also consonant with the public interest. Bishara (2011) identifies several important dimensions of the variation in NCA enforceability across states. Those dimensions include whether the state has a statute that governs NCA enforceability, what determines a legitimate employer's protectable interest, what constitutes the sufficient consideration to uphold a non-compete, whether the covenant is enforceable following involuntary termination of employment and whether courts are allowed to modify unenforceable covenants to make them enforceable.

In this paper, I rely on state-level changes in the enforceability of non-compete agreements on the above dimensions as quasi-exogenous shocks to labor mobility. Unequivocally, the large literature empirical literature on non-competes shows that following an increase in the enforceability of the NCA in a state, labor mobility in the state declines (Starr, Prescott, and Bishara, 2021; Jeffers, 2019).[9] Jeffers (2019) estimates that, on average, labor mobility declines by 8.6% after an increase in NCA enforceability in a state.

To identify changes in the enforceability of non-compete agreements, I perform a comprehensive search of state supreme court rulings and law changes that affect the enforceability of non-competes in each of the 50 US states on LexisNexis. Next, I read through the rulings

[8] Another often-cited reason for the use of non-competes is the desire to protect trade secrets. See Bishara (2006) for a discussion.

[9] The existing studies show a negative effect of NCA enforceability on the mobility of inventors (Marx, Strumksy, Fleming, 2009), CEOs (Garmaise, 2011), financial advisors (Gurun, Stoffman, and Yonker, 2021), founders (Ewens and Marx, 2018), and high-technology industry workers (Balasubramanian et al., 2022).

and summaries written by practicing lawyers to determine the nature and significance of a ruling on the enforceability of non-competes in the state. Finally, I combine the above rulings with those used by Ewens and Marx (2018) and Jeffers (2019).

Following these steps, I identify 12 state supreme court rulings that changed the enforceability of non-competes between 2000 and 2020.[10] NCA enforceability increased in seven states: Ohio (2004), Vermont (2005), Wisconsin (2009), Colorado (2011), and Illinois (2011)[11], Texas (2011), and Virginia (2013). Over the same time period, five court rulings made non-competes more difficult to enforce: Louisiana (2001)[12], South Carolina (2010), Montana (2011), Kentucky (2014), and Nevada (2016).[13] Additionally, I identify five state-level legislative changes that affected the enforceability of non-competes. In three states, NCA enforceability increased: Idaho (2008), Georgia (2010), and Arkansas (2015), and in three states, it decreased: Oregon (2008), Utah (2016), and Massachusetts (2018).

These state-level changes typically affect enforceability on one of the dimensions discussed by Bishara (2011). For example, rulings in Ohio (2004), Vermont (2005), and Colorado (2011) clarified that continued employment is a sufficient consideration to uphold a non-compete agreement. The ruling in South Carolina (2010) and the referendum in Georgia (2010) changed the ability of courts to modify the existing non-enforceable contract to make it enforceable.

I show the NCA enforceability changes I use in my main analysis on the map in Figure 1.1. I also report a full list of non-compete changes in my analysis in Table 1.3. Treated states are located in geographically diverse regions and vary both in size and IPO activity.

[10] Over this period, several other states experienced NCA enforceability change, which I do not use. Changes in New Hampshire (2016), New Mexico (2016), Connecticut (2016), and Rhode Island (2016) only affected physicians and are less relevant to my analysis. Change to the Alabama's Statute in 2016 increased NCA enforceability on some dimensions while decreasing it on some others. New Hampshire's non-compete legislation passed in 2012 required employers to disclose non-competes to new and existing employees, which is unlikely to significantly affect enforceability. Also, I do not use a ban on non-competes for tech workers in Hawaii (2015) since it was designed to promote startup performance in that state. In Section 6, I verify that excluding these states from the analysis does not significantly affect my main results.

[11] Illinois state supreme court's 2011 ruling was weakened by its later ruling in 2013.

[12] The change in Louisiana in 2001 was partially undone by a 2003 law.

[13] Nevada's change in 2017 was weakened by a 2017 Assembly Bill 276, which limited the enforceability of NCAs.

To examine which treated states have active IPO markets, I count the number of IPOs by firms headquartered in each state before and after the treatment. I report these numbers in Table 1.4. In total, my sample includes 1,224 IPOs in the control states. Among the treated states, Massachusetts (NCA weakening) has 183 IPOs before, and 11 after, and Texas (NCA strengthening) has 88 before and 59 after the treatment. Among other large treated states is Illinois (52 IPOs before and 20 after), Virginia (38 before and six after), Colorado (31 IPOs before and 16 after), and Georgia (29 IPOs before and 12 after). Other states have fewer IPOs both before and after the treatment. Based on these numbers, I conclude that some treated states, particularly Texas, have sufficient IPO activity both before and after the treatment.

1.3 Empirical Strategy

In this section, I describe my main empirical strategy. To estimate the effect of a change in NCA enforceability on a firm's IPO age, I estimate the following specification:

$$IPOAge_{it} = \mu_{j(i)} + \lambda_t + \delta_{s(i)} + \beta^{+}\text{treated}^{+}_{s(i)} \cdot \text{post}_{s(i),t} + \epsilon_{it} \tag{1.1}$$

where $IPOAge_{it}$ is the age of firm i when it goes public, $s(i)$ is the state that the firm i is headquartered in, t is the year of the IPO, and $j(i)$ is a firm's industry defined by its SIC code. $\text{treated}^{+}_{s(i)}$ is a dummy that equals one for states that experienced an increase in NCA enforceability. $\text{post}_{s(i),t}$ is a dummy that equals one for treated states in the years following the treatment. $\delta_{s(i)}$ and λ_t are the state and time fixed effects, respectively. $\mu_{j(i)}$ are the industry fixed effects. Since the treatment occurs at the state level, I cluster standard errors at the state level.

The coefficient of interest is β^{+}, which captures the weighted average of individual cohort-level treatment effects of NCA enforceability increase on the IPO age of firms headquartered in the treated states.

In my main analysis, I omit states that experienced a decrease in the NCA enforceability for two main reasons. First, most states that experienced a decrease in the enforceability of non-competes have very low levels of IPO activity. Since there are few IPOs in those states, it is more difficult to estimate the effect of the treatment on IPO age.[14] Second, using these states in my main analysis can hurt the interpretability of the estimates, as pointed out by recent literature on two-way fixed effect (TWFE) models. Borusyak, Jaravel, and Spiess (2021), Goodman-Bacon (2021), Callaway and Sant'Anna (2021), and de Chaisemartin and D'Haultfoeuille (2020) find that TWFE models make comparisons between units that receive treatment early in the sample and those that receive treatment late to estimate a weighted average of cohort-level treatment effects. In my settings, the TWFE model would use states that experience a decrease in non-compete enforceability early in the sample as a control group for states that experience an increase in NCA enforceability later on.

An important identifying assumption is the parallel trends assumption, which states that IPO age evolves similarly in the treatment and control groups in the absence of treatment. To provide empirical evidence in support of this assumption, I examine whether the IPO age evolves similarly in both groups before the treatment. In a recent paper, Sun and Abraham (2021) demonstrate that estimating a fully dynamic specification that includes leads and lags of treatment does not allow for performing a valid test for pre-trends. They show that coefficient estimates on the leads are contaminated by treatment effects and propose an alternative estimator that addresses the issue. In this paper, I take a different approach and build on the work of Borusyak, Jaravel, and Spiess (2021). In their paper, they propose to separate the test for the pre-trends from the estimation of treatment effects to avoid the issues discussed in Sun and Abraham (2021). Empirically, I estimate the following specification on a subsample of control and not-yet-treated observations:

$$IPOAge_{it} = \mu_{j(i)} + \lambda_t + \delta_{s(i)} + \sum_{\tau \in [(9,8),(7,6),(5,4),(3,2),(1)]} \beta_\tau \cdot \text{treated}_{s(i)} \mathbf{1}[\tau_1 \leq |t - t^{tr}_{s(i)}| \leq \tau_2] + \epsilon_{it} \tag{1.2}$$

where t^{tr} is the year in which treatment occurs in state $s(i)$ for treated states and zero otherwise. $treated_{s(i)}$ is a dummy that equals one for states with an increase in NCA enforceability

[14] One notable exception is Massachusetts, which passed a law in 2019 that limited the enforceability of non-competes. Unlike most other NCA shocks in the sample, the law was not retroactive, meaning that the enforceability of non-competes already signed at that time did not change.

Since in many treated states there are only a few IPOs in any given year, I group observation in two-year buckets, denoted by $\tau = (\tau_1, \tau_2)$. β_τ then captures the difference in average IPO age between treated and control states τ_1 to τ_2 years before the treatment. In the absence of pre-trends, I expect coefficients β_τ to be close to zero.

I estimate the above specification and plot β_τ estimates and the corresponding 95% confidence intervals on Figure 1.2. I find that none of the coefficient estimates are statistically significant at the 5% level. Based on these results, I conclude that there is no strong evidence of the presence of differential pre-trends in IPO age in the treated group compared to the control group.

One important identifying assumption in my setting is that NCA enforceability changes were not introduced to affect the performance of private firms, startups, as well as local IPO markets. For example, one potential threat to identification is that local VC firms or other business interest groups lobby the state legislature to introduce statutes regulating the use of non-competes in the state to affect startup performance. Similarly, the empirical design would be compromised if judges make rulings taking into account the state of the local economy, particularly the IPO activity of firms located in the state.

I argue that the above scenario is unlikely. First, the reason why enforceability changes in a state is because, in a given court case, the judge decides to make a ruling that differs significantly from the precedent previously established. The rulings are based on the judge's interpretation of the local statutes and not the state of the local economy. Additionally, Jeffers (2019) argues that these rulings are unlikely to be driven by the career concerns of judges and likely stem from their interpretation of the law. Second, in a series of robustness checks, I verify that my results are not driven by legislative changes in the NCA enforceability.

Next, I investigate whether IPOs in treated and control states differ in their observable characteristics before the treatment. I compare average IPO characteristics in both groups before the treatment. I report the results in Table 1.5. I find that IPOs in the treated states are 1.7 years older, and the difference is statistically significant at the 10% level. They raise \$10 million fewer proceeds, and the firms are smaller in size, although these differences are not statistically significant at the 10% level. There are 8% fewer IPOs that are backed by VC investors in the treated states and 6% fewer high-technology firms. Overall, I conclude that IPOs by firms headquartered in treated and control states are fairly similar in their observable characteristics.

Recently, a number of studies have identified and discussed empirical issues associated with the use of two-way fixed effect (TWFE) estimators. In particular, Borusyak, Jaravel, and Spiess (2021), Goodman-Bacon (2021), de Chaisemartin and D'Haultfoeuille (2020), and Sun and Abraham (2021) show that the treatment effect estimate in a TWFE setting is a weighted average of individual treatment effects, and where some of the weights are negative. These papers demonstrate that the problem arises due to the treatment effect heterogeneity across units and time. Such concerns are important in my setting for two reasons. First, as explained in the previous sections, different states experience different types of changes to the enforceability of NCAs, and, hence, their effect on labor mobility. Second, different states have different levels of NCA enforceability prior to the change, making the effect of shocks on labor mobility heterogeneous (Bishara, 2011).

To address such concerns, I perform several additional exercises and robustness checks. I obtain cohort-level treatment effects by estimating my model with just one treated state and a pool of control states. I perform this exercise for each treated state in the sample and report the estimates in Section 6. I find that my results remain both qualitatively and quantitatively similar to the full-sample estimates.

1.4 Results

A. Timing of IPOs

In this section, I use the empirical approach outlined above to estimate the effect of a non-compete shock on the timing of a firm's IPO. I use the firm's age at the time of its IPO as the main dependent variable. The coefficient of interest is on the interaction term between the treatment and a post-treatment dummy variable.

I report the estimates of my main specification in Table 1.6. The coefficient estimate on the interaction term is negative, at -3.94 years, and statistically significant at the 1% level. The estimate indicates that following an increase in the enforceability of NCAs in a state, firms headquartered in the state are, on average, 3.94 years younger than firms headquartered in the control states at the time of their IPO.

The treatment effect estimate is quite large compared to my sample's average IPO age of 15.5 years. To investigate this further, I use the log of a firm's age at its IPO as the main dependent variable. Such transformation would put lower weights on the observations with high IPO age. I report the estimate in Column (2). The estimated effect of NCA enforceability increase on log IPO age is negative and statistically significant at the 5% level. The magnitude of the estimate suggests a $1 - \exp(-0.216) \approx 19\%$ decrease in a firm's age at IPO for treated firms relative to the control group.

Next, I examine how my estimates differ depending on whether an IPO is backed by VC firms. VC funds have a limited time to exit their investment, usually ten years, making it difficult for them to delay exits, such as IPOs. In turn, I perform several additional tests to investigate whether my results differ for VC-backed firms. I create a triple interaction term (treated x post x VC-backed) to investigate whether my treatment effect estimates differ for VC and non-VC-backed firms. I report the estimates in Table 1.7. First, in Column (1), I estimate the same specification as before and add a VC-backed dummy which equals one if VCs back an IPO. I find that my treatment effect estimates remain virtually the same in magnitude, with a four-year decrease in IPO age, statistically significant at the 1% level. VC-backed firms are, on average, 5.8 years younger at their IPO, the estimate is statistically significant at the 1% level. All else equal, VC-backed firms are expected to IPO faster since VC investors have a limited time window during which they must exit their investments. In turn, VC-backed firms should be younger when they go public.

Next, I introduce a triple interaction term, which captures the difference in treatment effect estimates between VC-backed and non-VC-backed IPOs. I report the estimates in Column (2). The estimate on the post-treatment interaction term is negative at -6.1 years and statistically significant at the 1% level. Firms that are not backed by VCs are 6.1 years younger at their IPO following an NCA enforceability increase. In contrast, the estimate on the triple interaction term is positive. The treatment effect estimate for VC-backed firms is 3.26 years smaller in absolute magnitude than for non-VC-backed firms. These estimates imply that a VC-backed firm is $6.143 - 3.26 = 2.883$ years younger at its IPO following the treatment. This treatment effect estimate for VC-backed firms is statistically significant at the 1% level (p-value of 0.007).

I attribute the difference in estimates for VC-backed and non-VC-backed IPOs to the constraints faced by VC funds. Consider a private firm that is considering going public and that is concerned with employee retention. The firm's shareholders select the optimal time to go public that achieves the optimal level of employee turnover and maximizes the firm's value. However, if a firm is backed by VC investors, their preferences can diverge from firm value maximization. In particular, when the optimal IPO time exceeds the VC fund's liquidation date, VC investors would exit earlier than what's optimal from the regular shareholder's point of view. Such conflict would limit a firm's ability to time its IPO in response to labor mobility and lead to smaller effects of NCA shocks on the IPO age.

To summarize, I find that following an increase in the enforceability of NCAs in a state, firms headquartered in the state go public earlier. In the next series of tests, I examine whether this main result is due to firms strategically adjusting their IPO timing or due to a change in the composition of private-to-public exits.

B. Composition and timing of exits of private firms

B.1. Composition of Private Firms' Exits

In the following tests, I examine whether changes in NCA enforceability affect the probability that a private firm is acquired or fails. For example, following a change in the enforceability of NCA, some private firms may find it optimal to get acquired instead of going public, affecting the composition and age of firms that go public. In such a scenario, my results can be explained not by the timing of IPO by firms but rather by a change in the population of firms that decide to do an IPO.

Observing private-to-public exits is challenging due to data limitations. For example, commonly used datasets only track acquisitions of larger private firms, and failures of smaller firms are difficult to observe. Additionally, the majority of small private firms never intend to go public. In this paper, I address those challenges by focusing on private firms that receive venture capital financing. Such firms constitute a large portion of my main sample of IPOs, and the detailed information on the exits of these firms from VC portfolios is observable.

I use the Pitchbook dataset and examine the exits of firms that received their first VC financing between 2000 and 2008. I set the cutoff to 2008 to be able to observe the type of exit made by the firm. For each firm, I obtain its exit information, such as the type (IPO, acquisition, bankruptcy) and the exit date. I classify firms that are headquartered in the treated states and that exit the VC portfolio following the treatment in the state as treated. The rest of the firms are classified into the control group. I then estimate a linear probability model using dummy variables for getting acquired and going bankrupt as the outcome variables. The estimates capture the effect of an increase in NCA enforceability on a firm's propensity to get acquired and its likelihood of failure. Since there are only a few IPOs in this restricted VC simple, I omit them from the analysis.[15]

I report the estimates in Table 1.8. In Column (1), I estimate my main specification with the dummy that equals one if a firm is acquired as the dependent variable. I find that following a strengthening of NCAs in a state, VC-backed firms in those states are 1.28% less likely to get acquired. The estimate, however, is not statistically significant at the 10% level. In Column (2), I use the dummy for a firm's failure as the dependent variable. I find that firms are 3.27% more likely to fail following the treatment. However, the estimate is also not statistically significant at the 10% level.

Based on these results, I do not find evidence of a change in the composition of exits following an increase in NCA enforceability. One potential issue with the above results is that VC funds can sell their portfolio companies through acquisitions to hide failures (Metrick and Yasuda, 2011). In such a scenario, the results above may not be informative about the effect of NCA shocks on the fraction of successful exits by VC-backed firms. To investigate

[15]There are 1,846 IPOs in my main sample, of which 1,135 are VC-backed. Of those, only 421 IPOs are by firms that received their first VC financing round between 2000 and 2008.

this further, I examine the likelihood that a firm will make a successful exit following the NCA enforceability increase. The main dependent variable in Column (3) is a dummy that equals one if the firm's valuation at its exit is at least twice as much as the amount of VC capital invested in the firm. I find that firms are 1.5% more likely to make a successful exit from a VC portfolio, although the estimate is not statistically significant.

To summarize, I fail to find evidence that NCA shocks significantly affect the composition of VC exits. I conclude that my main findings on the effect of NCA enforceability increase on IPO age are unlikely to be driven by a change in the population of firms that decide to go public.

B.2. Timing of Other VC Exits

In the following tests, I examine how NCA shocks affect the timing of other private-to-public exits. In particular, it is possible that an increase in NCA enforceability leads to earlier exit from a VC portfolio through acquisitions and bankruptcies, not just IPOs. To investigate this possibility, I use the same sample of VC-backed firms as in the previous section and estimate the effect of NCA enforceability increase on the age of firms when they get acquired or fail.

I report the estimates in Table 1.9. The dependent variable in all columns is *Exit Age*, where the exit age is measured as the difference between the exit date and the founding date for a firm. The main independent variable is the interaction term between treated and post dummies. All specifications include state and exit year fixed effects and industry fixed effects.

The dependent variable in Column (1) is a firm's age when it exits a VC portfolio through either an IPO, an acquisition, or a failure. The treatment effect estimate is a 0.27 years increase in age, however, it is not statistically significant at the 10% level. Next, in Column (2), I use the age at IPO as the dependent variable. The treatment effect estimate is still negative, although it is not statistically significant. This is likely due to a low power since there are only a few IPOs in treated states in the VC sample. In Column (3), I use the dummy that equals one if a firm fails. An increase in NCA enforceability leads to a 0.32-

year reduction in the time it took those firms to fail, and the estimate is not statistically significant. Finally, in Column (4), I focus on the effect of NCA shocks on the age of VC-backed firms when they get acquired. I find that following the treatment, an average firm's age at acquisition increases by 0.71 years, and the estimate is statistically significant at the 5% level.

Finally, I separately analyze the effect of NCA shocks on the timing of successful and unsuccessful acquisitions. In Column (5), I estimate the effect of NCA shocks on acquisition age for the subsample of successful acquisitions. The treatment effect estimate is not statistically different from zero. In contrast, the estimate for unsuccessful acquisitions in Column (6) is positive, 0.89 years, and it is statistically significant at the 5% level. One possible interpretation of these results can be given in the context of a strategic default model. If strong NCA enforceability benefits firms, they will likely experience an increase in value, increasing the expected time until the firm defaults strategically.

To summarize, I fail to find evidence that the age of VC-backed firms changes when they exit through failures of successful acquisitions. The evidence in this subsection suggests that my main results are IPO-specific, and similar timing trends are not observed for other types of private-to-public exits of firms.

1.5 Mechanism

In this section, I conduct several empirical tests to establish the economic mechanism behind my main result. Two main channels can explain why firms go public earlier when NCAs become more enforceable. The first, which I refer to as *stock option compensation mechanism*, states that the illiquidity of a private firm's stock can reduce the mobility of employees compensated with equity. In turn, firms have incentives to delay IPO to retain workers. The second, which I call *valuation mechanism*, states that if NCA enforceability affects a firm's valuation or growth rate, this can lead firms to reach the IPO sooner.

A. Stock Option Compensation Mechanism

I start by examining the *stock option compensation mechanism.* According to this theory, the mobility of a firm's employees compensated with stock options depends on whether the firm is public or private. Employees holding vested options on a firm's stock have limited time, usually three months, to exercise such options if they decide to leave the firm. Such exercise is costly, as employees need to pay the exercise price and the associated taxes. Unlike public firms, employees of private firms cannot easily finance such exercises with a sale of the underlying stock (Aran, 2018). In turn, they are less likely to leave the firm. Anticipating such dynamics, the firm has incentives to delay its IPO to retain the workers. Once NCAs become more enforceable, the firm can use non-competes for employee retention, making stock options and delaying the IPO less effective. In turn, this leads to an earlier IPO.

This theory generates several empirical predictions. First, I anticipate that the effects of NCA enforceability should be stronger for firms that use stock option compensation more heavily. If a firm does not use any stock options to compensate its employees, whether it is public should have little effect on an employee's cost-benefit analysis when deciding whether to leave the firm. This should hold since an employee can still hold her stock grants even after leaving the private firm. In other words, in the absence of stock option compensation, the decision to leave the firm does not affect the employee's wealth accumulated through equity compensation.

Second, the effect of an increase in NCA enforceability should be stronger for firms whose options are costlier to exercise. When the costs of exercising a firm's stock options are low, employees find it easier to finance them if they decide to leave the firm. This makes staying private less effective at retaining the firm's workers. In such instances, non-compete enforceability should have little effect on the firm's IPO timing.

A.1. Stock Options

According to this mechanism, an increase in the NCA enforceability has a larger impact on IPO timing for firms that heavily use stock options for employee compensation. If a firm does not use any stock options, then employee mobility would not differ depending on the firm's public status. In contrast, if a larger share of an employee's compensation comes from stock options, her decision on whether to stay at the firm depends more on whether the firm is public. As a result, the firms that use a lot of stock options should respond more to the NCA shock in their IPO timing decision.

However, directly testing the above prediction is difficult as researchers do not observe the amount of stock options grants given to employees of private firms, as they are not publicly disclosed. In this paper, I use the data on the number of a firm's stock options reported in its first 10-K filing after the IPO. Such an approach has two potential difficulties: first, some of a firm's outstanding options may be held by investors other than employees, leading to a mismeasurement in the size of employee stock option grants. While such a scenario is possible, I argue that it is unlikely as prior studies show that for firms, most options are given to employees (Oyer and Schaefer, 2005). Second, the number of options outstanding after the IPO can be very different than before the IPO. This is unlikely: as I report in the data section, in my sample, the average time between the IPO and the subsequent 10-K filing is 218 days. On the firm's first 10-K filing, I get to observe the number of options outstanding at the time of the filing and the number of options exercised and granted over the preceding year.

I use the number of outstanding options normalized by the number of common shares outstanding as my measure of stock option compensation. I split the sample into terciles with high, medium, and low numbers of stock options outstanding after the IPO and estimate my main specification for each subsample. I report the results in Table 1.10. The dependent variable in all columns is *IPO Age*, the firm's age at the time of its IPO. I report the estimates for a subsample with a high number of stock options in Column (1). The estimate on the interaction term is negative and statistically significant at the 5% level. Following an increase in the NCA enforceability, firms are, on average, 6.08 years younger when they go public. In Column (2), I report the treatment effect estimates for the subsample with a medium number of options outstanding. The estimate on the interaction term is small in magnitude and not statistically significant at the 10% level. Finally, in Column (3), I report the estimate for the subsample with a low number of options. The coefficient estimate on the interaction term is -3.9 years, but it is not statistically significant.

I conclude that my main findings regarding the effect of NCA enforceability on IPO age are driven by firms with high levels of stock option compensation. This result is consistent with the stock options compensation mechanism, as firms for which public status can affect labor mobility are the ones granting options to employees.

A.2. Stock Option Exercise Costs

In this section, I examine how the effect of an NCA enforceability increase on IPO age varies depending on the costs of stock option exercise. Under the stock option compensation mechanism, I expect that firms whose options are costlier to exercise respond more to the treatment. Intuitively, when the cost of a stock option exercise is low, it is easier for workers to finance the exercise. In turn, staying private is a less effective employee retention device for such firms. As a result, these firms should respond less to a change in the NCA enforceability.

I directly test this prediction by estimating how the treatment effect varies with the costs of stock option exercise. When employees exercise a firm's stock options, they must pay the strike price and the taxes associated with the exercise. The taxes for most stock option grants come in the form of the alternative minimum tax payment (AMT), which depends on the difference between the option's strike price and the market price of shares at the time of the exercise.[16] I refer to this difference as the moneyness of the options.

I start by exploring how the treatment effect depends on the moneyness of the options. I measure moneyness as the percentage difference between the IPO price and the weighted average strike price of options exercised during the previous year. I estimate my main specification on subsamples with the moneyness above and below the median. I report the estimates in Table 1.11. In Column (1), the coefficient estimate on the high-moneyness subsample is large in magnitude and statistically significant at the 5% level. In contrast, for the low-moneyness subsample, the estimate in Column (2) is close to zero in magnitude, and it is not statistically significant at the 10% level. I conclude that my results are driven by the subsample of IPOs for which option exercise leads to a large tax expense due to the

[16] Not all stock option exercises would trigger the AMT payment. AMT is paid only when its amount exceeds the regular income and capital gains taxes that a taxpayer would otherwise pay. Additionally, in this analysis, I assume that an employee exercises the options and holds the underlying shares indefinitely. Selling the shares would trigger additional capital gains taxes.

appreciation of shares relative to the time when the option grants were made.

Although employees pay lower taxes on options that are deep in the money, these options are more likely to have a lower strike price, decreasing the exercise costs. To examine this issue further, I attempt to estimate the total costs of option exercise at any given firm. To do so, I obtain data on the weighted average price and the number of stock options exercised during the year from the first 10-K filing. I then estimate the total exercise costs as

$$TotalExerciseCosts = \frac{K_{t,exercised} \cdot N_{t,exercised} + 0.26 \cdot (P_{IPO} - K_{t,exercised}) \cdot N_{t,exercised}}{P_{IPO} \cdot N_{t,outstanding}} \tag{1.3}$$

where $K_{t,exercised}$ is the weighted average strike of the exercised options in the last year, $N_{t,exercised}$ is the number of such options exercised, $N_{t,outstanding}$ is the number of options outstanding at the beginning of the year, and $K_{t,outstanding}$ is their average strike price. P_{IPO} is the IPO offer price.

The numerator of the above formula includes costs that are due to the strike price and the tax-related costs. I use the 26% tax rate as the rate at which option exercise is usually taxed for incentive stock options.[17] I scale the total costs by the value of shares underlying the outstanding options to make the measure comparable across firms.

I estimate my specification for subsamples that correspond to high and low exercise costs relative to the median. I report the results for the high-costs subsample in Column (3). I find that for firms with high option exercise costs, the effect of an NCA enforceability increase is a five-year decrease in IPO age, statistically significant at the 1% level. In contrast, the estimates for firms with low cost of option exercise are small in magnitude and not statistically significant.

I conclude that my results are driven entirely by firms with high costs of exercising stock options. These results further corroborate the stock option compensation mechanism.

A.3. Lockup Provisions

[17]In this calculation, I am ignoring the regular income and capital gains taxes that an employee would pay in the absence of the AMT. In practice, the tax component of the exercise costs should be adjusted for the fact that when the AMT payment is triggered, the employee does not have to pay the regular income and capital gains taxes.

In the next tests, I examine how changes in the NCA enforceability affect the use of IPO lock-up provisions by firms. Such provisions are commonly used in IPOs, and they prohibit certain shareholders from selling their shares for a period of time after the firm goes public, typically 180 days (Field and Hanka, 2001). Lock-ups can limit employee mobility because workers cannot sell their stock in the firm during the lock-up period. In turn, they cannot finance the stock option exercise discussed above, bounding them to the firm.

NCAs can have two effects on the use of lock-up provisions. Consider a firm that experiences an increase in the NCA enforceability and, in turn, decides to go public earlier. The decision to do the IPO earlier has a positive effect on labor mobility, which is the opposite of the effect of an NCA increase. Which of the two forces dominates is an empirical question. When the effect of an increase in NCA enforceability on labor mobility is too weak to counteract the effect of the IPO, firms have incentives to bridge that gap by using lock-ups more aggressively. Alternatively, firms may decide to limit the use of lock-ups when non-competes significantly reduce the mobility of labor, and earlier IPO does not counteract that effect.

Empirically, I examine how non-compete enforceability increases affect the firm's propensity to subject its investors to lock-up provisions on the IPO. Two groups of investors most commonly subject to lock-ups are the management and the existing shareholders, including the employees. I create dummy variables that equal one if the firm uses lock-ups for each type of investor and use them as dependent variables in my main specification.

I report the estimates in Table 1.12. In Column (1), the dependent variable is a dummy that equals one if a firm uses any lock-ups in its IPO. The estimate on the interaction term is close to zero, and it is not statistically significant. Following an increase in the NCA enforceability, firms that go public do not change their propensity to use lock-up provisions. The dependent variable in Column (2) is a dummy that equals one if a firm subjects its management to lock-up provisions. The estimate on the interaction term is 0.8%, and it is not statistically significant at the 10% level. In contrast, the estimate in Column (3) is positive and statistically significant at the 10% level. Following an increase in the NCA enforceability, firms are 9% more likely to subject their current shareholders, including their employees, to lock-up provisions.

The above findings have two major implications. First, I find evidence that firms utilize additional stock-based employee retention devices following the NCA shocks. This provides additional evidence in favor of the stock compensation channel. Second, the results here suggest that my main findings are likely driven by rank-and-file employees, not the management.

B. Valuation Mechanism

The *valuation mechanism* can explain the effect of the NCA enforceability increase on IPO timing through the effect non-competes have on firm value. Stronger non-competes can lower the turnover of employees at the firm and increase its cash flows and value (Younge and Marx, 2016).[18] In the real options framework, this can lead the firm to reach the IPO stage sooner (Dixit and Pindyck, 1994). Draho (2000) and Benninga, Helmantel, and Sarig (2005) show that firms are more likely to go public earlier when they have higher cash flows and market valuations.

I start by examining the effect of an increase in NCA enforceability on the size of a firm's IPO and report the results of this test in Table 1.13. In Column (1), the dependent variable is the logarithm of the IPO proceeds. Counter to the above predictions, following a strengthening of non-competes in a state, IPOs raise 23.1% fewer proceeds in the offering. The estimate is statistically significant at the 5% level. Next, I examine the effect of the NCA shocks on the post-IPO market value of the firm's equity in Column (2). The estimate for the treatment effect of an increase in non-compete enforceability is negative and statistically significant at the 5% level. Following a strengthening of non-competes in a state, the post-IPO market value of firms is 26.4% smaller. Finally, in Column (3), I examine whether firms change the fraction of equity they offer to IPO investors. The estimate on the interaction term is small in magnitude, and it is not statistically significant at the 10% level.

To summarize, I find that following a strengthening of non-competes in a state, IPOs

[18] Whether non-competes increase firm value is an empirical question. Theoretically, they can also decrease firm value, as they make it more difficult for small firms to attract labor (Stuart and Sorenson, 2003). I focus on the positive effect of NCAs on firm value since a negative effect cannot explain my main findings.

raise fewer proceeds, and firms are smaller in size after the offering.

Next, to further examine the valuation story, I test the predictions made by Benninga, Helmantel, and Sarig (2005) that use the real options framework to show that firms are more likely to go public earlier if they have higher cash flows and valuations. The main empirical challenge with this test is the unobservability of valuations for private firms. However, the valuation mechanism would also predict that NCA enforceability increases benefits public firms and their value, and market prices are observable for these firms. In the following test, I examine the effect of NCA shocks on Tobin's Q of public firms in the empirical framework outlined above.

I report the estimates in Table 1.14. The dependent variable in all columns is the log of a firm's Tobin's Q. In Column (1), I report the estimate for all public firms in my sample. I find that following an increase in the enforceability of NCAs, a firm's Q decreases by 6.6%. The estimate is statistically significant at the 5% level. Another empirical challenge with such a test is the fact that public firms are larger than private firms, and they usually have multiple locations in different states. As a result, firms headquartered in control states can be exposed to treatment through their offices in other states. To ameliorate such concerns, I split the sample into terciles based on the number of a firm's employees and estimate the treatment effect for each subsample. Firms with fewer employees are less likely to have locations in multiple states.

I report the estimates for the subsample of firms with a low number of employees in Column (2). The coefficient estimate on the interaction term is negative, a 14.3% decrease, and it is statistically significant at the 1% level. In contrast, the estimates for firms with medium and high levels of employment, in Columns (3) and (4), respectively, are much smaller in magnitude, and they are not statistically significant.

To summarize, I find that valuations decrease for low-employment public firms. Such a finding is in contrast with Younge and Marx (2016). It can be explained by the fact that stronger NCAs make it more difficult for firms to hire workers from other firms, lowering their valuations (Stuart and Sorenson, 2003). Based on the above results, I do not find strong support for the valuation mechanism. IPOs are smaller in size, and the effect on the valuation of public firms is the opposite of what the valuation channel predicts.

1.6 Robustness

A. Single-State Dependence

I start by investigating whether the results are driven by IPO timing patterns in any single state. According to Table 1.1, 32.6% of firms in my main sample is headquartered in California and 10.9% in Massachusetts. In the following set of tests, I investigate whether my results are sensitive to the IPO dynamics in any single state. I omit each of the largest states from the sample and reestimate my main specification. I report the estimates in Table 1.15. The estimates on a subsample that excludes the IPO of firms headquartered in California are in Column (1). The estimate on the interaction term is -2.83 years, close to the full sample estimate. Similarly, in Column (2), I omit Massachusetts from the sample and find that the treatment effect estimates remain similar.

Next, I investigate whether my results are driven by Texas alone. Among states with an increase in NCA enforceability, Texas has the most IPOs, with 7.9% of firms headquartered there. I exclude Texas from the sample and report the estimates in Column (3). The coefficient estimate on the interaction term is -3.27 years, close to the full sample estimates.

B. Court Rulings and Legislative Changes

In my main analysis, I use state-level changes in the enforceability of non-competes driven by either court rulings or legislative changes. In practice, there are several important differences between these two types of enforceability changes. First, court rulings affect the enforceability of existing and new employment contracts. In contrast, some of the legislative changes are not retroactive and only apply to new agreements. Second, legislative changes can be driven by business interest groups to promote the performance of firms in the state. In contrast, most court rulings in the sample are unanticipated in that they reversed the rulings of lower courts (Jeffers, 2019).

I remove firms with legislative changes to the NCA enforceability and estimate my main specification. I report the estimates in Column (6). The estimate of the effect of an increase in NCA enforceability on the IPO age is -3.79 years, close to the previous results. I conclude that lobbying by interest groups is unlikely to explain my main findings.

C. States with reversals and other NCA changes

NCA enforceability changes in Louisiana, Illinois, and Nevada were fully or partially reversed later in the sample. To check whether this affects my main results, I exclude those states from the sample and repeat my analysis. I report the estimates in Column (7). The estimated treatment effect remains quantitatively similar.

Additionally, several states had NCA changes that targeted only specific groups of workers and thus are not classified as treated in my main analysis. I repeat my estimation after excluding those states: New Hampshire, New Mexico, Connecticut, Rhode Island, and Hawaii, and I report the estimates in Column (8). The estimated treatment effect remains virtually the same both in terms of its magnitude and statistical significance.

D. Staggered Nature of Shocks

In my main analysis, I rely on the two-way fixed effects (TWFE) model to identify the effect of an increase in the NCA enforceability on the IPO age of firms. More recently, several papers pointed out that TWFE generally fails to uncover a meaningful weighted average of individual cohort-level treatment effects (Borusyak, Jaravel, and Spiess, 2021; Goodman-Bacon, 2021; de Chaisemartin and D'Haultfoeuille, 2020; Sun and Abraham, 2021; Callaway and Sant'Anna, 2021). The issue arises because the TWFE model uses units treated early as a part of the control group for units that are treated later in the sample. When treatment effects are heterogeneous, such comparisons do not allow for uncovering the underlying treatment effects.

In this robustness exercise, I examine this issue in more detail by estimating individual state-level treatment effects for treated states. For each treated state with an increase in the NCA enforceability, I estimate my main specification on the sample that includes this state and all control states. Such an exercise is similar to the first step of the estimator proposed by Callaway and Sant'Anna (2021).

I plot the estimates and the corresponding 95% confidence intervals in Figure 1.3. First, the estimates of a treatment effect of NCA enforceability increase on the IPO age are negative for Colorado, Texas, Virginia, Illinois, and Georgia. The magnitudes are very similar as well: between five and six years. One exception is Colorado, for which the estimate is smaller in magnitude and not statistically significant at the 5% level.

Additionally, I report the estimates for the two treated states with the most IPOs: Texas and Virginia, in Columns (4) and (5) of Table 1.15. In both columns, the estimate on the interaction term is negative and statistically significant at the 1% level. Moreover, the estimates are close in magnitude: a 5-year reduction in IPO age for firms in Texas and a 5.5-year reduction for firms in Virginia.

Based on these results, I conclude that issues associated with using TWFE models are not driving my main results.

E. Office Locations

One of the drawbacks of using a state-level variation in labor mobility is that some firms can have offices in multiple states. For example, assume that a firm headquartered in California has additional offices in Texas. Even though non-compete enforceability did not change in California, a change in Texas affects the mobility of the firm's workers. To address this challenge, I use IPO prospectuses to collect detailed information on the locations of offices for firms in my sample. For each firm in the sample, I identify states where it had offices at the time of the IPO. Then I examine how sensitive my results are to the exclusion of multi-state firms.

In Column (1) of Table 1.16 I report the estimates from my baseline specification. In Column (2), I restrict the sample to firms whose offices are located in control states and firms whose only location in treated states is their headquarters. The estimate remains similar in magnitude, with a decrease in the IPO age by 3.3 years. In Column (3), I restrict the sample to only firms with offices in one state. The estimate becomes smaller but remains negative. Notably, it is no longer statistically significant at the 10% level, most likely due to a loss of power due to a smaller sample size.

1.7 Conclusion

In this paper, I present novel evidence on the role of labor market frictions in a firm's IPO timing decision. Following a decrease in the mobility of its labor, a firm goes public at a younger age. To explain these results, I propose a stock option compensation channel, which states that workers compensated with options on the illiquid stock of a private firm are less likely to leave the firm before its IPO. Firms hold up their employees compensated with stock option grants and delay their IPO to retain their employees.

The findings have several interesting implications. First, I document a novel motive for the firms to stay private. While the existing literature focuses on the need for liquidity as one reason firms go public, I show that the reverse can also be the case. If a firm has a strong motive to retain its workers, delaying the IPO can be one effective way of doing so. As a result, illiquidity can serve as a benefit of being private. Second, my findings suggest a new explanation for why firms stay private longer than in the past decades. As the composition of firms in the economy has changed over the last 20 years, more firms heavily rely on human capital and use stock option compensation for their employees. Consequently, staying private longer is important for these firms. In turn, this can increase the average age of firms at the IPO, which has been a puzzle in the literature. Empirically quantifying how important the labor-market explanation is for the listing gap is an interesting avenue for future research.

1.8 Tables and Figures

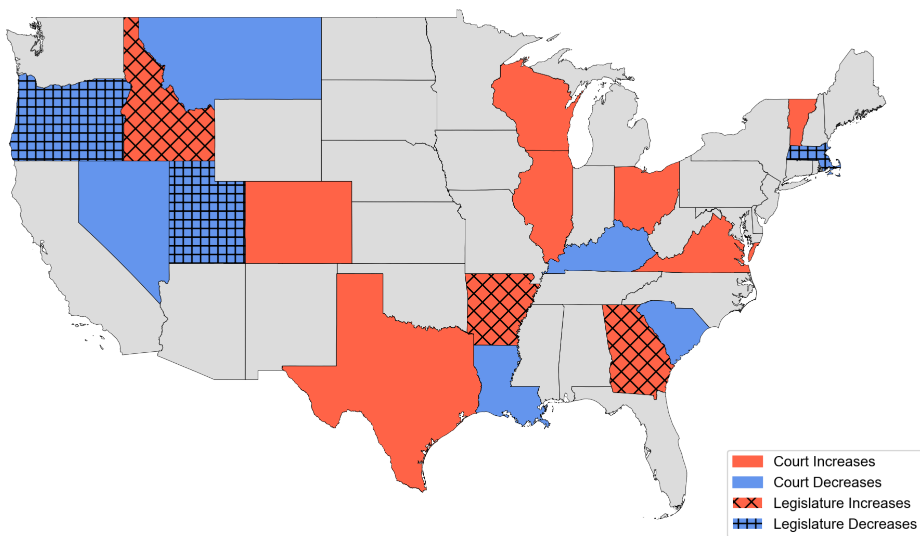

The map shows major state-level changes to the enforceability of non-compete agreements between 2000 and 2020. Court Increase (Decrease) is an increase in the enforceability caused by a state's Supreme Court ruling. Legislative Increase (Decrease) is a change in the enforceability of non-compete agreements as a result of a legislative action.

Figure 1.1. Changes in the Enforceability of Non-compete Agreements

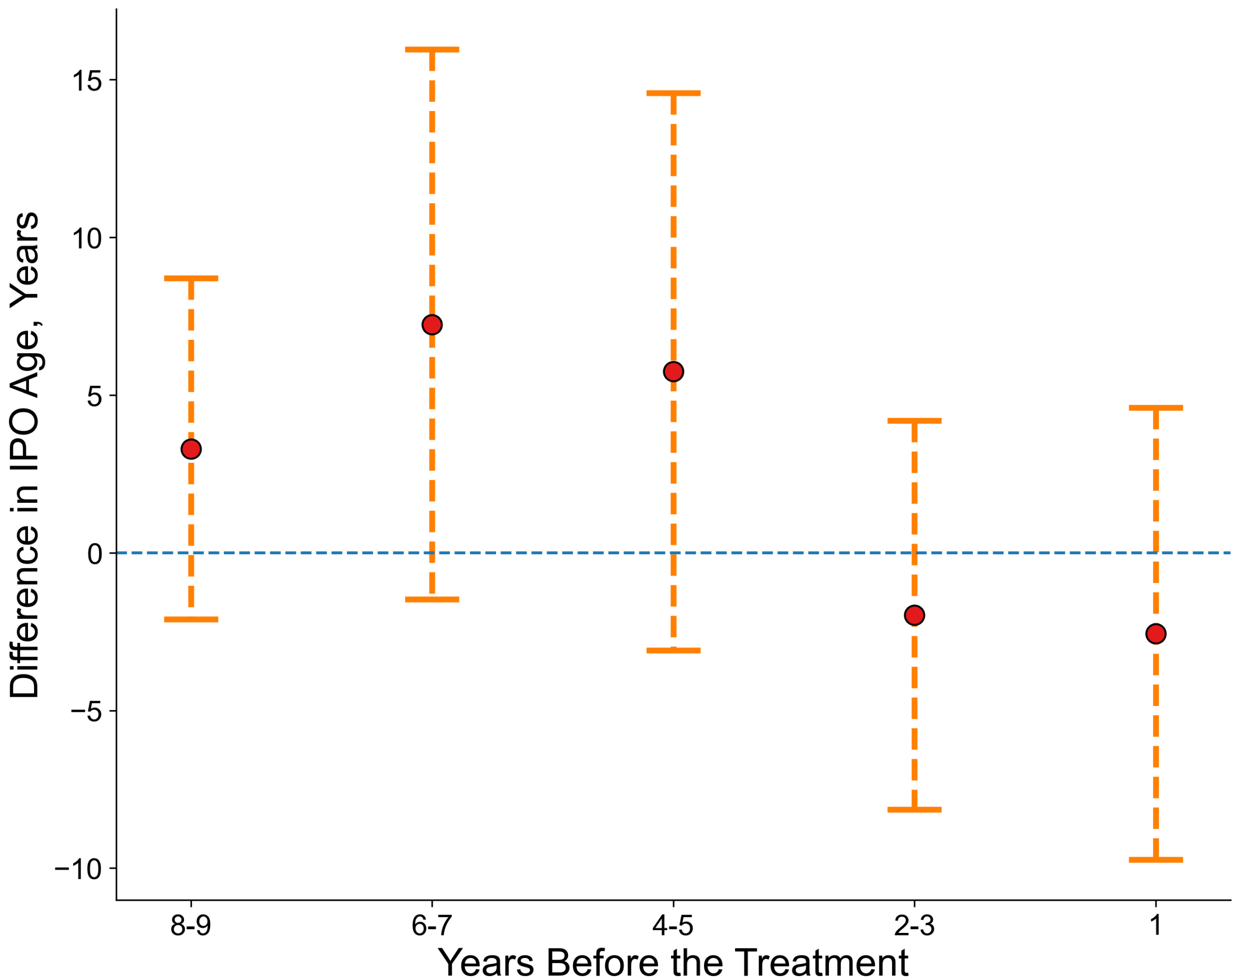

The figure plots the differences in the average IPO age between treated and control states before the treatment. On the vertical axis, I show the estimated difference at a horizon shown on the horizontal axis.

Figure 1.2. Pre-Trends in IPO Age

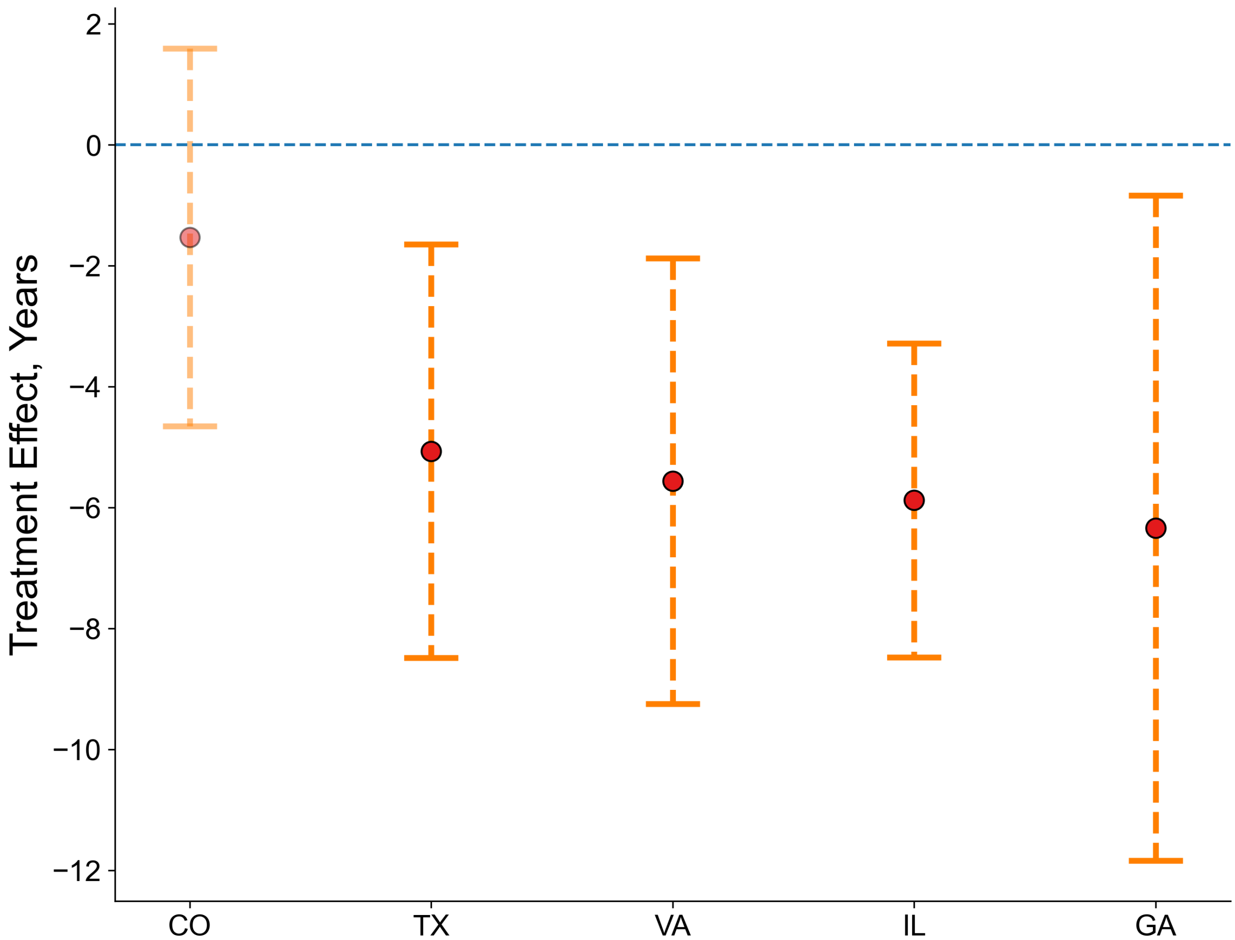

The figure reports the estimates from the following specification:

$$IPOAge_{it} = \mu_{j(i)} + \lambda_t + \delta_{s(i)} + \beta^{+}\text{treated}^{+}_{s(i)} \cdot \text{post}_{s(i),t} + \epsilon_{it} \tag{1.4}$$

where i is a firm, t is the IPO year, $s(i)$ is the state in which a firm is headquartered, $j(i)$ is firm i's industry. On the vertical axis, I plot the β^{+} estimate for each state. I estimate these for each state shown on the horizontal axis by running the above regression on the sample consisting of that state and all control states in my sample.

Figure 1.3. State-by-State Estimates

TABLE 1.1. Summary Statistics: Initial Public Offerings

The table reports summary statistics for the sample of 1,846 IPOs by US firms between 2000 and 2020. *Proceeds* is the total amount of financing raised in an IPO. *IPO Age* is the difference between a firm's IPO date and the founding date. *VC-backed* is a dummy that equals one if the IPO is VC-backed. *CA, MA, TX* are dummies equal to one if a firm is headquartered in CA, MA, and TX, respectively. *Options Outstanding* (*Options Exercisable*) is the number of outstanding stock options (eligible for exercise) divided by the number of common shares outstanding as reported on the first 10-K filing after the IPO. *Options Exercised* is the number of stock options exercised during the year prior to the first 10-K filing, scaled by the number of common shares outstanding. *Stock Compensation Expense* (*Deferred Compensation Expense*) is the total stock (deferred) compensation expense in the year of IPO, divided by the market value of the firm's equity. *Total Assets* is the book value of the firm's assets as reported on its first 10-K filing. *Market Capitalization* is the market value of the firm's equity as of its IPO date. *Lockup Provision* is a dummy that equals one if lockup provisions are included. *Current Shareholders Lockup* and *Management Lockup* are dummies that indicate the presence of lockup provisions for the existing shareholders and the management, respectively. *Technology Industry* (*Biotechnology Industry*) is a dummy that equals one if a firm is in a high technology (biotechnology) industry, as defined by Loughran and Ritter (2004). *N States with Offices* is the number of states in which the firm has offices at the time of its IPO.

	Mean	std	10%	50%	90%	N Observations
Proceeds ($M)	184.76	270.32	38.00	103.50	381.35	1,846
IPO Age, Years	14.56	15.49	3.74	9.32	31.47	1,846
VC-backed IPO	0.67	0.57	0.00	1.00	1.00	1,846
CA	0.326	0.469	0.000	0.000	1.000	1,846
MA	0.109	0.312	0.000	0.000	1.000	1,846
TX	0.079	0.269	0.000	0.000	0.000	1,846
Options Outstanding	0.097	0.147	0.000	0.071	0.201	1,846
Options Exercisable	0.060	0.082	0.000	0.041	0.133	1,846
Options Exercised	0.014	0.063	0.000	0.005	0.033	1,846
Stock Compensation Expense	0.014	0.021	0.001	0.007	0.030	1,709
Deferred Compensation Expense	0.002	0.006	0.000	0.000	0.006	1,656
Total Assets ($M)	742.68	2,403.33	47.60	177.62	1,569.89	1,846
Market Capitalization ($M)	1,009.81	1,854.29	75.36	414.14	2,299.48	1,844
Lockup Provision	0.898	0.303	0.000	1.000	1.000	1,846
Management Lockup	0.668	0.471	0.000	1.000	1.000	1,846
Current Shareholders Lockup	0.293	0.455	0.000	0.000	1.000	1,846
Technology Industry	0.361	0.481	0.000	0.000	1.000	1,846
Biotechnology Industry	0.243	0.429	0.000	0.000	1.000	1,846
N States with Offices	2.025	2.041	1.000	1.000	4.000	1,543

TABLE 1.2. Summary Statistics: VC-backed Firms

The table reports summary statistics for the sample of VC-backed firms. The sample consists of 9,717 US firms that raised their first round of VC financing between 2000 and 2008. *Year First VC* is the year in which a firm received its first VC financing round. *IPO* is a dummy that equals one if a firm went public, *Acquired* is a dummy that equals one if a firm was acquired, and *Bankrupt* is a dummy that equals one if a firm failed. *Still Private* is a dummy that equals one if the firm is still private as of December 31st, 2019. *Total VC Capital Raised* is the total capital raised by a firm from VC investors. *Exit Multiple* is the ratio of proceeds generated by VC investors in the exit divided by the total amount of funding invested in the firm by VC investors. *Successful Exit* is a dummy that equals one if the exit multiple is greater or equal to two.

	Mean	std	10%	50%	90%	N Observations
Year First VC	2004.23	2.78	2000	2005	2008	10,114
IPO	0.04	0.20	0.00	0.00	0.00	10,114
Acquired	0.53	0.49	0.00	1.00	1.00	10,114
Bankrupt	0.35	0.48	0.00	0.00	1.00	10,114
Still Private	0.08	0.26	0.00	0.00	0.00	10,114
Total VC Capital Raised ($M)	35.54	56.25	1.00	15.20	90.68	10,114
Exit Multiple	2.14	5.71	0.00	0.00	6.02	6,133
Firm Age at Exit	10.64	5.64	4.25	9.70	18.01	10,114
Successful Exit	0.13	0.34	0.00	0.00	1.00	10,114

TABLE 1.3. State-level changes in the enforceability of NCAs

The table lists major changes to the enforceability of non-compete covenants in the United States between 2000 and 2020. Panel A lists states' supreme court rulings that changed the NCA enforceability in those states. Panel B lists legislative changes that changed NCA enforceability in a state. The date column reports the date of a ruling or when a new piece of legislation goes into effect. The last column reports the direction of the change, + denotes an increase in the enforceability of non-compete covenants, while − stands for a decrease in NCA enforceability.

State	Date	Court Case	Direction
		Panel A: Court Rulings	
Lousiana	6/29/2001	*Shreveport Bossier, Inc. v. Bond*	-
Ohio	3/10/2004	*Lake Land Employment Group v. Columber*	+
Vermont	8/19/2005	*Summits 7 v. Kelly*	+
Wisconsin	7/14/2009	*Star Direct v. Dal Pra*	+
South Carolina	5/24/2010	*Poynter Investments v. Century Builders of Piedmont*	-
Colorado	6/11/2011	*Luncht's Concrete Pumping v. Horner*	+
Montana	11/22/2011	*Wrigg v. Junkermier, Clark, Campanella*	-
Illinois	12/1/2011	*Reliable Fire Equipment v. Arredondo*	+
Texas	12/16/2011	*Marsh v. Cook*	+
Virginia	9/12/2013	*Assurance Data v. Malyevac*	+
Kentucky	6/19/2014	*Creech v. Brown*	-
Nevada	7/21/2016	*Golden Road Motor Inn v. Islam*	-
		Panel B: Legislative Changes	
Idaho	7/1/2008		+
Oregon	1/1/2008		-
Georgia	11/3/2010		+
Arkansas	8/6/2015		+
Utah	5/10/2016		-
Massachusetts	10/1/2018		-

TABLE 1.4. Number of IPOs in Treated and Control States

This table reports the number of IPOs in control and treated states before and after the treatment. Control states include all US states not listed in Table 1.3. Other States include UT, WI, NV, LA, KY, OR, SC, ID, MT, AR, OH, and VT.

State	N IPO Pre	N IPO Post
Control States:	1,221	
Treated States (Enforceability Increases):		
TX	77	68
IL	40	18
VA	37	7
CO	26	16
GA	24	13
Treated States (Enforceability Decreases):		
MA	146	56
Other States	50	47

TABLE 1.5. IPO Characteristics in Treated and Control States

This table reports the average characteristics of IPOs in control and treated states before the treatment. The first column reports the averages for 400 IPOs in treated states. The second column reports the average characteristics for 1,221 IPO in control states. The third column reports the difference between the two and results of a two-sided t-test for equal means. *IPO Age* is the difference between a firm's IPO and its founding date. *IPO Proceeds* is the total amount of funding raised in an IPO. *Lockup Provision* is a dummy that equals one the firm uses lockup provisions in its IPO. *Market Capitalization* is the market value of a firm's equity as of the date of the IPO. *VC-backed* is a dummy that equals one if the IPO is VC-backed. *Technology Industry* (*Biotechnology Industry*) is a dummy that equals one if a firm is in a high technology (biotechnology) industry, as defined by Loughran and Ritter (2004). *Stock Options Outstanding* is the number of stock options outstanding divided by the number of common shares. *Stock Options Exercisable* is the number of stock options eligible for exercise as reported on the first 10-K filing, divided by the number of common shares. *Stock Options Exercised* is the number of stock options exercised during the year before the first 10-K filing, scaled by the number of common shares outstanding. *Stock Compensation/Market Cap* is the stock-based compensation expense scaled by the firm's market capitalization. *, **, and *** indicate that the difference is statistically significant at the 10%, 5%, and 1% level, respectively.

	Treated	Control	Difference
IPO Age, Years	15.57	13.85	1.71*
IPO Proceeds ($M)	172.68	182.70	-10.02
Lockup Provision	0.86	0.89	-0.03
Market Capitalization, ($M)	628.68	844.49	-215.81
VC-backed	0.61	0.69	-0.08**
Technology Industry	0.34	0.40	-0.06**
Biotechnology Industry	0.21	0.24	-0.03
Stock Options Outstanding	0.095	0.103	-0.007
Stock Options Exercisable	0.058	0.064	-0.006
Stock Options Exercised	0.012	0.016	-0.004*
Stock Compensation/Market Cap	0.010	0.014	-0.004

TABLE 1.6. Non-Compete Shocks and IPO Age

The table reports the estimates from the following specification:

$$y_{it} = \mu_{j(i)} + \lambda_t + \delta_{s(i)} + \beta^{+}\text{treated}^{+}_{s(i)} \cdot \text{post}_{s(i),t} + \epsilon_{it} \tag{1.5}$$

where i is a firm, t is the IPO year, $s(i)$ is the state in which a firm is headquartered, $j(i)$ is firm i's industry. The sample consists of 1,846 US IPOs between 2000 and 2020. y_{it} is the dependent variable of interest, $treated^{+}_{s(i)}$ is a dummy that equals one for IPOs in states that had an increase in the NCA enforceability. $post_{s(i),t}$ is a dummy that equals one for treated states following the treatment. The dependent variable in Column (1) is *IPO Age*, the difference between a firm's IPO date and its founding date. In Column (2), the dependent variable is log *IPO Age*, the logarithm of the IPO age of a firm. Standard errors are clustered at the state level. *, **, and *** indicate that a coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1) IPO Age	(2) log IPO Age
treated · post	-3.935***	-0.216**
	(0.989)	(0.100)
Constant	13.982***	2.283***
	(0.086)	(0.009)
Observations	1,444	1,444
R-squared	0.391	0.331
Year FE	Yes	Yes
State FE	Yes	Yes
Industry FE	Yes	Yes

TABLE 1.7. Non-Compete Shocks and IPO Age, VC-backed Firms

The table reports the estimates from the following specification:

$$IPOAge_{it} = \delta_{s(i)} + \mu_{j(i)} + \lambda_t + \rho VC_i + \beta^{+}\text{treated}^{+}_{s(i)} \cdot \text{post}_{s(i),t} + \\ + \gamma^{+}\text{treated}^{+}_{s(i)} \cdot \text{post}_{s(i),t} \cdot VC_i + \epsilon_{it}$$

where i is a firm, t is the IPO year, $s(i)$ is the state in which firm i is headquartered in, $j(i)$ is firm i's industry. The sample consists of 1,846 US IPOs between 2000 and 2020. The dependent variable in all columns is *IPO Age*, the difference between a firm's IPO date and its founding date. $treated^{+}_{s(i)}$ is a dummy that equals one for IPOs in states that had an increase in the NCA enforceability. $post_{s(i),t}$ is a dummy that equals one for treated states following the treatment. *VC* is a dummy that equals one if an IPO is backed by VC investors. The dependent variable in all columns is *IPO Age*, the difference between the IPO date and the founding date for a firm. Standard errors are clustered at the state level. *, **, and *** indicate that a coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1) IPO Age	(2) IPO Age
treated · post	-4.019***	-6.143***
	(0.938)	(1.229)
VC	-5.802***	-6.273***
	(0.581)	(0.722)
treated · post · VC		3.260**
		(1.214)
Constant	17.935***	18.267***
	(0.404)	(0.486)
Observations	1,444	1,444
R-squared	0.427	0.429
Year FE	Yes	Yes
State FE	Yes	Yes
Industry FE	Yes	Yes

TABLE 1.8. Non-Compete Shocks and VC Exits

The table reports the estimates from the following specification:

$$y_{it} = \mu_{j(i)} + \lambda_t + \delta_{s(i)} + \beta^{+}\text{treated}^{+}_{s(i)} \cdot \text{post}_{s(i),t} + \epsilon_{it}$$

where i is a firm, t is the IPO year, $s(i)$ is the state in which firm i is headquartered in, $j(i)$ is firm i's industry. y_{it} is the dependent variable of interest. The sample consists of 9,717 private US firms that received at least one round of VC financing between 2000 and 2008. $treated^{+}_{s(i)}$ is a dummy that equals one for exits from a VC portfolio by firms headquartered in a treated state. $post_{s(i),t}$ is a dummy that equals one for treated states following the treatment. The dependent variables in Columns (1), (2), and (3) are dummies that equal one if a firm exited a VC's portfolio through an IPO, an acquisition, and bankruptcy, respectively. The dependent variable in Column (4) is *Still Private*, a dummy that equals one if the firm did not exit a VC's portfolio by the end of 2019. The dependent variable in Column (5) is a dummy that equals one if a firm's exit valuation divided by the total amount of VC funding it received while private is greater than two. Standard errors are clustered at the state level. *, **, and *** indicate that a coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)
	Acquired	Bankrupt	Successful Exit
treated · post	-1.279	3.266	1.532
	(2.228)	(2.075)	(2.095)
Constant	52.546***	35.593***	12.929***
	(0.235)	(0.219)	(0.221)
Observations	8,754	8,754	8,754
R-squared	0.153	0.110	0.043
State FE	Yes	Yes	Yes
Vintage FE	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes

TABLE 1.9. Non-Compete Shocks and Timing of VC Exits

The table reports the estimates from the following specification:

$$ExitAge_{it} = \mu_{j(i)} + \lambda_t + \delta_{s(i)} + \beta^{+}\text{treated}^{+}_{s(i)} \cdot \text{post}_{s(i),t} + \epsilon_{it}$$

where i is a firm, t is the IPO year, $s(i)$ is the state in which firm i is headquartered in, $j(i)$ is firm i's industry. The sample consists of 9,717 private US firms that received at least one round of VC financing between 2000 and 2008. $treated^{+}_{s(i)}$ is a dummy that equals one for firms headquartered in states that had an increase in the NCA enforceability. $post_{s(i),t}$ is a dummy that equals one for treated states following the treatment. The dependent variable in all columns is *Exit Age*, the firm's age, measured as the difference between the exit date and the founding date. Column (1) contains the results for the whole sample. Columns (2), (3), and (4) contain the estimates for subsamples of IPO, Bankruptcy, and Acquisition exits, respectively. Columns (5) and (6) contain estimates for the sample of successful and unsuccessful acquisition exits, respectively. An acquisition is classified as good if its transaction value is at least twice as large as the amount of capital invested in the firm by VC investors. Standard errors are clustered at the state level. *, **, and *** indicate that a coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	Exit Age					
	All Exits	IPO	Bankruptcy	Acquisition	Acquisition Good	Acquisition Bad
treated · post	0.269	-0.380	-0.323	0.710**	0.355	0.886**
	(0.207)	(1.027)	(0.646)	(0.289)	(0.400)	(0.337)
Constant	10.558***	9.720***	10.356***	9.858***	9.563***	9.927***
	(0.022)	(0.057)	(0.075)	(0.026)	(0.031)	(0.031)
Observations	8,754	343	3,146	4,585	1,058	3,515
R-squared	0.425	0.495	0.392	0.388	0.412	0.396
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Vintage FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes

TABLE 1.10. Mechanism - Stock Option Compensation

The table reports the estimates from the following specification:

$$IPOAge_{it} = \mu_{j(i)} + \lambda_t + \delta_{s(i)} + \beta^{+}\text{treated}^{+}_{s(i)} \cdot \text{post}_{s(i),t} + \epsilon_{it}$$

where i is a firm, t is the IPO year, $s(i)$ is the state in which a firm is headquartered, $j(i)$ is firm i's industry. The sample consists of 1,886 US IPOs between 2000 and 2020. $treated^{+}_{s(i)}$ is a dummy that equals one for IPOs in states that had an increase in the NCA enforceability. $post_{s(i),t}$ is a dummy that equals one for treated states following the treatment. The dependent variable in all columns is *IPO Age*, the difference between a firm's IPO date and its founding date. Columns (1), (2), and (3) correspond to high, medium, and low terciles based on the number of stock options outstanding at the time of the first 10-K filing, scaled by the number of shares outstanding. Standard errors are clustered at the state level. *, **, and *** indicate that a coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)
	IPO Age		
	High Options	Medium Options	Low Options
treated · post	-6.075**	-0.130	-3.898
	(2.641)	(3.024)	(2.886)
Constant	12.317***	12.572***	15.685***
	(0.139)	(0.347)	(0.435)
Observations	361	349	305
R-squared	0.592	0.545	0.468
Year FE	Yes	Yes	Yes
State FE	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes

TABLE 1.11. Mechanism - Stock Option Exercise Costs

The table reports the estimates from the following specification:

$$IPOAge_{it} = \mu_{j(i)} + \lambda_t + \delta_{s(i)} + \beta^{+}\text{treated}^{+}_{s(i)} \cdot \text{post}_{s(i),t} + \epsilon_{it}$$

where i is a firm, t is the IPO year, $s(i)$ is the state in which a firm is headquartered, $j(i)$ is firm i's industry. The sample consists of 1,886 US IPOs between 2000 and 2020. y_{it} is the dependent variable of interest, $treated^{+}_{s(i)}$ is a dummy that equals one for IPOs in states that had an increase in the NCA enforceability. $post_{s(i),t}$ is a dummy that equals one for treated states following the treatment. The dependent variable in all columns is *IPO Age*, a firm's age at the time it goes public. Columns (1) and (2) correspond to subsamples of IPOs of firms whose options have high and low moneyness, respectively. I define moneyness as the difference between the IPO offer price and the average weighted strike price of the options exercised in the first year following the IPO, scaled by the IPO offer price. Columns (3) and (4) correspond to subsamples with high and low stock option exercise costs, respectively. The details on the calculation of these costs are in the main text. Standard errors are clustered at the state level. *, **, and *** indicate that a coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
	Moneyness		Option Exercise Costs	
	High	Low	High	Low
treated · post	-7.835**	-0.566	-5.077***	0.246
	(3.405)	(2.069)	(1.821)	(3.097)
Constant	11.723***	15.160***	12.862***	12.630***
	(0.188)	(0.239)	(0.112)	(0.305)
Observations	560	546	552	507
R-squared	0.457	0.469	0.465	0.525
State FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes

TABLE 1.12. Mechanism - Lockups

The table reports the estimates from the following specification:

$$y_{it} = \mu_{j(i)} + \lambda_t + \delta_{s(i)} + \beta^+ \cdot \text{treated}^+_{s(i)} \cdot \text{post}_{s(i),t} + \epsilon_{it}$$

where i is a firm, t is the IPO year, $s(i)$ is the state in which a firm is headquartered, $j(i)$ is firm i's industry. The sample consists of 1,886 US IPOs between 2000 and 2020. y_{it} is the dependent variable of interest, $treated^+_{s(i)}$ is a dummy that equals one for IPOs in states that had an increase in the NCA enforceability. $post_{s(i),t}$ is a dummy that equals one for treated states following the treatment. The dependent variable in Column (1) is a dummy that equals one if a firm uses lock-up provisions in its IPO. The dependent variable in Column (2) is a dummy that equals one if a firm uses lock-up for the management. The dependent variable in Column (3) is a dummy that equals one if a firm uses lock-up provisions for its current shareholders, including the employees. Standard errors are clustered at the state level. *, **, and *** indicate that a coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1) Lockup Dummy	(2) Management Lockup	(3) Current Shareholders Lockup
treated · post	0.001	0.008	0.090*
	(0.037)	(0.035)	(0.045)
Constant	0.893***	0.654***	0.287***
	(0.003)	(0.003)	(0.004)
Observations	1,444	1,444	1,444
R-squared	0.511	0.856	0.735
Year FE	Yes	Yes	Yes
State FE	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes

TABLE 1.13. Mechanism - IPO Size following NCA shocks

The table reports the estimates from the following specification:

$$y_{it} = \mu_{j(i)} + \lambda_t + \delta_{s(i)} + \beta^+ \text{treated}^+_{s(i)} \cdot \text{post}_{s(i),t} + \epsilon_{it} \tag{1.6}$$

where i is a firm, t is the IPO year, $s(i)$ is the state in which firm i is headquartered in, $j(i)$ is firm i's industry. y_{it} is the dependent variable of interest, $treated^+_{s(i)}$ is a dummy that equals one for IPOs in states that had an increase in the NCA enforceability. $post_{s(i),t}$ is a dummy that equals one for treated states following the treatment. The dependent variable in Column (1) is log *IPO Proceeds*. The dependent variable in Column (2) is log *post-IPO Market Value*, where market capitalization is measured using the IPO offer price. The dependent variable in Column (3) is *IPO % Equity Sold*, which is the number of shares offered in the IPO divided by the number of shares outstanding after the IPO. Standard errors are clustered at the state level. *, **, and *** indicate that a coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1) log IPO Proceeds	(2) log post-IPO Market Value	(3) IPO % Sold
treated · post	-0.263**	-0.307**	-1.103
	(0.114)	(0.116)	(1.980)
Constant	4.720***	6.023***	27.550***
	(0.010)	(0.010)	(0.169)
Observations	1,444	1,372	1,372
R-squared	0.334	0.344	0.264
Year FE	Yes	Yes	Yes
State FE	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes

TABLE 1.14. Mechanism - NCA and valuation of public firms

The table reports the estimates from the following specification:

$$\log Q_{it} = \mu_{j(i)} + \lambda_t + \delta_{s(i)} + \beta^{+}\text{treated}^{+}_{s(i)} \cdot \text{post}_{s(i),t} + \epsilon_{it}$$

where i is a firm, t is the IPO year, s is the state in which a firm is headquartered, $j(i)$ is firm i's industry. $treated^{+}_{s(i)}$ is a dummy that equals one for IPOs in states that had an increase in the NCA enforceability. $post_{s(i),t}$ is a dummy that equals one for treated states following the treatment. The dependent variable in all columns is log of a firm's Tobin's Q, measured as the ratio of a firm's value to the book value of assets. The sample consists of US public firms covered by Compustat. Column (1) contains the estimates for the entire sample. Columns (2), (3), and (4) contain the estimates for firms with low, medium, and high number of employees, respectively. Standard errors are clustered at the state level. *, **, and *** indicate that a coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
	log Q			
	All Firms	Low Employment	Medium Employment	High Employment
treated · post	-0.066**	-0.143***	-0.019	-0.033
	(0.027)	(0.034)	(0.032)	(0.022)
Constant	0.209***	0.605***	-0.044***	0.065***
	(0.004)	(0.004)	(0.004)	(0.004)
Observations	95,878	31,290	29,184	35,364
R-squared	0.415	0.352	0.581	0.450
State FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes

TABLE 1.15. Robustness

The table reports the estimates from the following specification:

$$IPOAge_{it} = \mu_{j(i)} + \lambda_t + \delta_{s(i)} + \beta^{+}\text{treated}^{+}_{s(i)} \cdot \text{post}_{s(i),t} + \epsilon_{it}$$

where i is a firm, t is the IPO year, $s(i)$ is the state in which firm i is headquartered in, $j(i)$ is firm i's industry. $treated^{+}_{s(i)}$ is a dummy that equals one for IPOs in states that had an increase in the NCA enforceability. $post_{s(i),t}$ is a dummy that equals one for treated states following the treatment. The dependent variable in all columns is *IPO Age*, the difference between a firm's IPO date and its founding date. Columns (1), (2), and (3) report the estimates on subsamples that exclude all IPOs in California, Massachusetts, and Texas, respectively. Estimates in Columns (4) are obtained by excluding all treated states in which treatment was due to a legislative change. Columns (5) and (6) report the estimates on a subsample consisting of all control states and Texas and Virginia, respectively. Column (7) reports the estimates obtained by excluding states with partial or full treatment reversal: Louisiana, Illinois, and Nevada. Column (8) reports the estimates obtained by excluding the states which had NCA enforceability changes for a narrow set of workers: New Hampshire, New Mexico, Connecticut, Rhode Island, and Hawaii. Standard errors are clustered at the state level. *, **, and *** indicate that a coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1) No CA	(2) No MA	(3) No TX	(4) Only TX Treated	(5) Only VA Treated	(6) Only Court Rulings	(7) No LA, IL, NV	(8) No NH, NM, CT, RI, HI
treated · post	-2.832**	-3.935***	-3.268**	-5.065***	-5.558***	-3.790***	-4.060***	-3.970***
	(1.173)	(0.989)	(1.194)	(1.744)	(1.880)	(1.187)	(1.096)	(0.969)
Constant	16.079***	13.982***	13.435***	13.160***	12.603***	13.859***	13.782***	14.034***
	(0.168)	(0.086)	(0.058)	(0.087)	(0.008)	(0.097)	(0.087)	(0.085)
Observations	840	1,444	1,301	1,245	1,142	1,406	1,395	1,409
R-squared	0.399	0.391	0.396	0.417	0.424	0.398	0.393	0.393
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

TABLE 1.16. Robustness: Office Locations

The table reports the estimates from the following specification:

$$IPOAge_{it} = \mu_{j(i)} + \lambda_t + \delta_{s(i)} + \beta^{+}\text{treated}^{+}_{s(i)} \cdot \text{post}_{s(i),t} + \epsilon_{it}$$

where i is a firm, t is the IPO year, $s(i)$ is the state in which firm i is headquartered in, $j(i)$ is firm i's industry. $treated^{+}_{s(i)}$ is a dummy that equals one for IPOs in states that had an increase in the NCA enforceability. $post_{s(i),t}$ is a dummy that equals one for treated states following the treatment. The dependent variable in all columns is *IPO Age*, the difference between the firm's IPO date and its founding date. In Column (1), I report the full sample estimates. In Column (2), I focus on firms that have none of their offices in the treated states or those whose only office in the treated states is the headquarters. In Column (3), I focus on a subsample of firms that only have offices in a single state. Standard errors are clustered at the state level. *, **, and *** indicate that a coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1) All	(2) Only HQ in treated states	(3) Single-State Firms
treated · post	-3.935***	-3.327**	-1.430
	(0.989)	(1.359)	(1.336)
Constant	13.982***	12.264***	11.360***
	(0.086)	(0.122)	(0.106)
Observations	1,444	1,058	869
R-squared	0.391	0.431	0.446
Year FE	Yes	Yes	Yes
State FE	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes

Chapter 2

Real Effects of ESG Investing

2.1 Introduction

Economists have long recognized that in the presence of externalities, markets generally fail to achieve socially efficient outcomes (Coase, 1960). For example, a production quantity selected by the owners of the factory can maximize the present value of their profits but at the same time result in a suboptimally high level of pollution from a social standpoint. Traditionally, government regulation was viewed as the solution (Pigou, 1920). More recently, investment in funds that incorporate Environmental, Social, and Governance (ESG) criteria in their investment decisions has been proposed as a market-based solution to the externalities (Hart and Zingales, 2017).

Growth in the popularity of ESG funds has been astounding. In 2019 alone, the number of ESG funds grew from 511 to 592, and the total assets under their management increased from \$321 billion to \$465 billion (Investment Company Institute, 2021). At the same time, ESG funds charge higher fees than non-ESG funds: 0.61% vs. 0.41% (Morningstar, 2020). While the difference in fees may be justified by the premium investors pay to reduce the externalities, whether the funds affect the ESG performance of their portfolio companies is unclear.

In this paper, I investigate this question and examine whether ESG mutual funds improve the ESG performance of their portfolio companies. To answer this question, I focus on the environmental aspect of the ESG, using the Toxic Release Inventory (TRI) data published by Environmental Protection Agency (EPA). I use this data with a novel shock that generates a quasi-random divestment from some firms by a subset of ESG funds.

Given that the holdings of ESG funds are relatively small, it is unclear whether they can affect firms' policies. On the one hand, ESG funds can exercise their control rights (Shleifer and Vishny, 1986), engage with the management, and improve the environmental performance of the firms. Additionally, the funds can use a threat of divestment to improve ESG performance at firms (Broccardo, Hart, and Zingales, 2020). On the other hand, ESG funds can lack the resources and incentives to impact corporate policies at portfolio companies. First, they may hold small stakes, which makes the exercise of voice or exit channels of governance ineffective. Second, they may suffer from conflicts of interest and engage in greenwashing. In both cases, ESG funds would have little impact on corporate outcomes, as they capture rent in the form of fees. Ultimately, whether ESG funds improve corporate performance on ESG dimensions is an empirical question.

Estimating the causal effect of ESG investors on a firm's environmental policies is challenging. The fraction of equity held by ESG funds is not randomly assigned across firms, as the funds incorporate the environmental performance of the firms they choose to hold. In particular, ESG funds are more likely to select firms with better ESG performance (Friedman and Heinle, 2021; Heath et al., 2021). As a result, firms with high ESG ownership can be cleaner due to such selection, as opposed to the effect of the ESG funds on firms.

To circumvent the above challenges and estimate the treatment effect of ESG investment, I rely on a quasi-random shock to the fraction of firms' equity held by ESG funds. In 2018, MSCI, a leading provider of ESG ratings and data in the US, changed the methodology for one of its ESG indexes, the MSCI ESG Leaders index. Before May 2018, the MSCI required firms to have an MSCI ESG rating of at least 'B' to be eligible for exclusion from the index. After May 2018, the MSCI requires firms to have MSCI ESG ratings of at least 'BB' to be eligible for index inclusion. I use this change in methodology as a source of variation in the supply of ESG capital. Treated firms whose MSCI ESG rating was 'B' in 2018 got excluded from the MSCI ESG Leaders index. I use this shock in a difference-in-difference setup and compare the environmental performance of treated and control firms.

To verify that the shock impacted the ownership structure of treated firms, I study the effect of the change in MSCI's methodology on the fraction of a firm's equity owned by ESG funds. I find that following the exclusion from the MSCI ESG Leaders index, a firm is almost 20% less likely to be held by a fund that tracks one of the MSCI ESG indexes. Moreover, I find a negative and statistically significant effect of index exclusion on the fraction of a firm's equity held by ESG index funds. I verify that this result is driven by a reduction in the fraction of a firm's equity held by ESG index funds, particularly those using MSCI indexes as the benchmark.

Importantly, I find that the shock did not affect the level of institutional ownership. Moreover, I find that not only the total institutional ownership but also the ownership by index funds does not change following the shock. These findings indicate that the MSCI ESG index exclusion affected the composition but not the level of institutional ownership. Such observation is important since the shock does not affect environmental outcomes through its effect on the total level of institutional ownership.

After establishing that the shock does indeed lower the amount of ESG capital invested in treated firms, I turn to analyze the effect of index exclusion on environmental outcomes. As my primary outcome variable, I use the total amount of toxic chemicals released by all plants owned by a firm. The main result is that treated firms have between 34% and 41% higher toxic emissions, depending on the specification. I find that this is driven by higher air and ground releases.

Next, I attempt to uncover the source of such an increase. First, it is possible that following the index exclusion, a firm is less likely to invest in pollution abatement. I empirically investigate this possibility using EPA's pollution prevention (P2) dataset. I find no evidence that firms respond to the index exclusion in their abatement activities.

Another possible source of reduction can be a change in a firm's productivity. On the internal margin, a treated firm may increase production at its existing plants following the treatment, which would cause higher levels of pollution. On the external margin, a treated firm can open new plants, which would increase total toxic releases.

I investigate each of those possibilities by estimating the effect of the index exclusion on productivity and the propensity to open or close plants. I find that there is no significant effect of the index exclusion on the productivity of a firm's plants. At the same time, treated firms are 8.3% more likely to open new plants than the control firms following the treatment. I conclude that an increase in pollution comes from an increase in the size of a firm's operations.

Next, I analyze the potential mechanisms behind this finding. There are two possible reasons why after a firm is excluded from the ESG index, it increases its toxic releases. First, it is possible that after the index exclusion, the firm faces a higher cost of capital. Thus it may find it more difficult to invest in pollution-reducing activities. Second, it is possible that after ESG investors divest, firms are less monitored in their environmental activities.

To examine whether the divestment affected the cost of capital, I examine whether firms' capital structure, valuations, and financial performance were affected by the divestiture. I find no evidence of a change in a firm's leverage, cash holdings, market-to-book ratios, or total assets. Based on this evidence, it is unlikely that firms pollute more because of their cost of capital. Such results are in line with Berk and van Binsbergen (2021), who fail to find evidence of the impact of divestment on the cost of capital. Similar to their paper, I conclude that the divestment was too small in magnitude to affect the cost of capital of firms.

Next, I examine the monitoring channel. I note that holdings of ESG funds are relatively small, in my sample, ESG funds hold, on average, only 0.4% of a firm's equity. It is thus unlikely that those funds can influence the management using their stakes alone. However, ESG funds are sponsored by large fund families that centralize voting and can use votes of non-ESG funds to induce managers to behave in a socially conscious way. When an ESG fund from a family exits its position in a firm, the fund family can have fewer incentives to affect the firm's environmental policies. This mechanism is similar to the one studied in Lakkis (2021) in the context of the index and active funds.

To examine this channel, I note that BlackRock is the only fund family from the 'Big Three': Vanguard, BlackRock, and SSGA, which offers the MSCI ESG Leaders index fund. Consequently, under the fund family monitoring channel, I expect the effect of divestment on pollution to be the strongest for firms in which BlackRock had large total family holdings. Intuitively, BlackRock is more likely to successfully promote the ESG agenda at a firm when its total holdings in the firm are large.

Using a triple-difference analysis, I find that my results are driven by firms for which BlackRock had large holdings. On the other hand, when I compare the treatment effects for firms with high and low holdings for Vanguard and SSGA, I find no difference in the estimates. Based on this evidence, I conclude that the exclusion from MSCI ESG Leaders worsened the environmental performance of firms as BlackRock reduced its monitoring efforts at those firms.

To summarize, I show that ESG mutual funds affect environmental policies at their portfolio companies. Such evidence sheds light on the welfare effects of ESG investing. In particular, it casts some doubt on the greenwashing concern. Interestingly, my results demonstrate that ESG funds do not need to be large in size to have a meaningful effect on their portfolio companies. As long as these funds are offered by large fund families, those fund families can use holdings of their other funds to promote the ESG agenda at portfolio companies.

My paper relates to three strands of the literature. First, the paper contributes to the emerging literature that studies the effects of ESG investing on portfolio companies.[1] Theoretically, the literature has identified two main channels through which investors can affect corporate policies: exit and voice (see Edmans and Holdernes (2017) for a comprehensive survey of the theoretical literature on voice and exit in the corporate governance context). In

[1]There is a rich body of literature studying Corporate Social Responsibility (CSR) policies adopted by firms, measurement and reporting of ESG activities, and the relation between CSR and a firm's performance. For comprehensive surveys on these topics, see Liang, and Renneboog (2020), Liang and Renneboog (2020), Matos (2020), Grewal and Serafeim (2021), and Gillan, Koch, and Starks (2021).

the context of ESG investing, Heinkel, Kraus, and Zechner (2001), Pastor, Stambaugh, and Taylor (2021), Landier and Lovo (2020), and Broccardo, Hart, and Zingales (2020) theoretically examine the effectiveness of ESG investing on corporate policies. Broccardo, Hart, and Zingales (2020) conclude that voice is a more effective mechanism for inducing ESG-friendly behavior at firms than exit is.

Empirically, the evidence of the impact of ESG funds on firms' ESG policies has been sparse. Berk and van Binsbergen (2021) develop and calibrate a model of ESG investing. They demonstrate that the impact of divestiture on the cost of capital is negligible, given the current state of the ESG market. Furthermore, Berk and van Binsbergen (2021) use exclusions from the FTSE USA 4Good index and show that they do not materially affect the cost of capital of the excluded firms. In closely related papers, Akey and Appel (2019) and Naaraayanan, Sachdeva, and Sharma (2021) find that shareholder activism leads firms to reduce toxic releases at their plants and invest in new abatement initiatives.[2] More broadly, there are several papers that document the effect of large institutional investors on E, S, and G outcomes. Gormley et al. (2021) find that campaigns by the 'Big Three' fund families succeed in increasing the gender diversity of corporate boards. Azar et al. (2021) study engagements of the 'Big Three' with portfolio firms and conclude that the investors are effective at reducing carbon emissions at firms. Dikolli et al. (2021) and Michaely, Ordonez-Calafi, and Rubio (2021) document that ESG mutual funds are more likely to support shareholder-sponsored ES proposals than non-ESG mutual funds.[3] My main contribution to this rich literature is to show that ESG mutual funds can improve the environmental performance of their portfolio companies. While Akey and Appel (2019) and Naaraayanan, Sachdeva, and Sharma (2021) provide similar evidence for activist investors, to the best of my knowledge, this paper is the first to show this result for ESG mutual funds, which are usually thought to be more passive in their engagement with firms than the activist investors.

[2] It is worth noting that Naaraayanan, Sachdeva, and Sharma (2021) focus exclusively on environmental activism, while Akey and Appel (2019) study the environmental consequences of hedge fund activist campaigns in general. In another paper, Barko, Cremers, and Renneboog (2021) show the positive effects of behind-the-scenes engagements on the ESG performance of targeted firms.

[3] Michaely, Ordonez-Calafi, and Rubio (2021) also find that ES mutual funds vote strategically, and they are less likely to support close proposals if they belong to a non-ES mutual fund family.

Second, I contribute to the literature that studies the return to ESG and impact investing. Several studies document that investors have a preference for investment in 'clean' firms (Hong and Kacpercyzk, 2009; Hartzmark and Sussman, 2019). A rich literature studies the risk and returns of ESG and impact funds. Renneboog, Ter Horst, and Zhang (2008) document lower returns of ESG funds than non-ESG funds and show that after performing the risk adjustment, the difference becomes small. Bialkowski and Starks (2016) show that flow-performance sensitivity is lower for ESG than for non-ESG funds. Similarly, Barber, Morse, and Yasuda (2021) show that investors derive non-pecuniary utility from investing in an impact fund. Most papers in this space focus on pecuniary returns to investors and their willingness to pay for the investment in 'green' funds. At the same time, little is known about whether this difference in performance between ESG and non-ESG funds reflects a premium investors pay to reduce an externality or a rent that is captured by fund managers. I contribute to this literature by showing that investors indeed succeed in reducing the externality by investing in ESG mutual funds.

Finally, I contribute to the literature on ESG ratings. Existing studies explain the factors behind the ratings and their discrepancies. In particular, Chatterji et al. (2016) and Gibson, Krueger, and Schmidt (2021) document a significant disagreement across providers of ESG scores. Berg, Koelbel, and Rigobon (2020) demonstrate that such divergence is mainly driven by different scopes and measurements of categories used by different providers. Liang and Renneboog (2017) show that common-law origin is the first-order explanation of a firm's CSR rating. I contribute to this literature by showing that ESG scores not only reflect a firm's performance on relevant E, S, and G dimensions but also affect firm-level outcomes on those dimensions. ESG scores affect the amount of ESG capital invested in a firm through ESG index funds and, in turn, corporate policies.

The rest of the paper proceeds as follows. Section 2 describes the data. Section 3 describes the main empirical strategy. Section 4 contains the main results. Section 5 provides robustness checks. Section 6 concludes.

2.2 Data

A. MSCI Ratings

MSCI is one of the leading providers of ESG ratings for US firms. They use a combination of publically available and proprietary data to evaluate the performance of firms on 35 key issues on the environmental, social, and governance dimensions. MSCI then aggregates the performance on these issues into a single letter grade such: 'CCC', 'B', 'BB', 'BBB', 'A', 'AA', and 'A'.

I obtain a complete list of ESG reports from MSCI ESG Direct. I record the MSCI ESG rating for each US-incorporated firm as of May 2018. In my sample of 1,911 firms, 497 have a rating of 'B', 534 are rated 'BB', 461 have a rating of 'BBB', and 262 are rated 'A'. Other letter grades are rarer.

I merge the data on MSCI ratings to CRSP-Computstat (CCM) merged database, using a combination of matching algorithms based on tickers and firm names. I manually cross-check all the matches. I can find a matching firm in CCM for 1,800 firms.

B. Toxic Releases Data

To measure environmental externalities generated by firms, I rely on the EPA's Toxic Release Inventory (TRI) dataset. TRI contains information on the toxic chemicals released by plants that satisfy the reporting criteria.[4] The data is self-reported by firms, and the firms have incentives to report their releases truthfully (Akey and Appel, 2021).

The data contains the amount of each toxic chemical released in a given year for each facility. Release amounts are categorized into releases into the air, ground, and water. For each plant, the data reports the parent firm that owns at least 50% of the interest in the plant.

I use the TRI for the years between 2011 and 2021. I identify a matching firm in Compustat for each parent company using matching by name and manually checking the resulting matches. As many of the plants in the TRI data are owned by privately-held firms, there

[4]Facilities are subject to reporting requirements if they have ten or more employees, have a NAICS industry code covered by TRI, and manufacture or otherwise process one of the TRI-covered chemicals in an amount that is above a pre-specified threshold.

is not always a match in Compustat. I can find TRI records for 702 firms covered by Compustat. I complement the matches with the ones from Xiong and Png (2019). I aggregate toxic releases across all plants and chemicals for each firm-year. Additionally, I keep separate records of pollution released in the air, water, and ground.

To measure a firm's investments in pollution prevention, I rely on the EPA's Pollution Prevention Dataset (P2). This dataset records pollution-preventing activities undertaken by firms.

For each firm-year, I collect all pollution prevention activities that a firm's plants implemented over the year. Following Akey and Appel (2021), I focus on two common types of abatement activities: good operating practices and process modifications. The former refers to improved maintenance and scheduling, while the latter usually involves installing new equipment. For each firm-year, I create two dummy variables equal to one if any of the firm's plants implemented one of the two abatement activities described above.

C. Mutual fund holdings

The data on mutual fund holdings comes from CRSP Mutual Fund Database. US mutual funds are required by the Securities and Exchange Commission (SEC) to disclose their holdings quarterly using N-Q and N-CSR forms. I focus on all equity mutual funds covered by CRSP Mutual Fund Database. I keep its holdings as of the latest report date for each fund in any given year. Additionally, I aggregate holdings across multiple share classes of the same fund. For each fund in the sample, I retrieve information on whether the fund is an index fund and the fund family that sponsors the fund.

To identify ESG mutual funds, I use MorningStar's list of ESG funds, which I link to the CRSP Mutual Fund holdings database by a fund's ticker and name. I identify 354 ESG equity funds in the sample and 79 index funds. Eight of those are benchmarked against the MSCI ESG Leaders index, and 13 are benchmarked against other MSCI ESG indexes.

For each firm in the sample, I compute the total number of the firm's shares held by all mutual funds combined. I divide it by the firm's total number of outstanding shares from Compustat to compute the percentage of its outstanding equity held by mutual funds.

Similarly, I compute the fraction of a firm's equity held by ESG funds and ESG funds benchmarked against an MSCI ESG Leaders index. Finally, for each of the 'Big Three' fund families: BlackRock, Vanguard, and State Street Global Advisors (SSGA), I compute the percentage of a firm's equity held by the fund family. Finally, to merge these data to Compustat, I use CRSP-Compustat merged database.

D. Firm's characteristics

For each firm in the sample, I obtain financial information from Compustat. I retain firms with at least $100,000 in total assets in the sample. I obtain information on the firm's total assets, investment, Tobin's Q, sales, leverage, and cash balances.

E. Summary Statistics

My final sample consists of almost 1,800 firms over the 2011-2021 period. I report the summary statistics in Table 3.1. An average firm in my sample has 10.4 plants and releases 1.8 million pounds of toxic releases per year. Note that this number is much larger than in other studies that use EPA TRI data, such as Akey and Appel (2019, 2021), because I analyze firms' environmental performance at the firm, not the plant level. Most of the releases occur at the firm's facilities (as opposed to offsite facilities), 1.4 million pounds per year. The average firm releases 78.9 thousand pounds of toxic chemicals in the water, 484.5 thousand in the air, and 540 thousand in the ground.

On average, 18% of the firms close a plant in a given year, and 12.9% open a new one. The average production growth for a firm is 5% per year. Firms frequently abate releases through a change in the operating practices (19.9% of the time) and a change in processing (17.3% of the time).

The average firm has a market capitalization of $13.7 billion and a leverage of 27.6%. Its cash holdings are 10.5% of its assets. Tobin's Q for the average firm is 1.88.

To better understand the ownership structure of firms in my sample, I compute the percentage of a firm's equity held by mutual funds. On average, firms have 35.7% of their equity held by mutual funds. As index funds grew in their size over the recent decades, they hold, on average, 15.9% of a firm's equity. Average holdings of the 'Big Three' fund families:

BlackRock, Vanguard, and SSGA are 1.2%, 6.3%, and 0.86%, respectively. The fraction of equity held by ESG index funds is still modest: on average, a firm in the sample only has 0.021% of its equity held by ESG index funds.

2.3 Empirical Strategy

Identifying the effect of a firm's ESG ownership on its policies is empirically challenging. ESG funds strategically choose their holdings, and they may select firms that have better performance on E, S, and G dimensions (Heath et al., 2021). Additionally, ESG funds may affect a firm's policies through exit or voice channels. The main identification challenge is disentangling these two effects, selection, and treatment. In this paper, I use a quasi-exogenous shock that affected portfolios of certain ESG funds to solve this identification issue. In the following sections, I describe the nature of the shock and the main identifying assumptions behind my empirical strategy.

A. MSCI ESG Leaders Index

MSCI is a leading provider of ESG ratings for US firms. Using rich data on ESG-related firm-level outcomes, MSCI rates firms and assigns them a letter grade that captures their ESG performance. In addition to ESG ratings, MSCI also constructs several ESG stock-market indexes that are used as a benchmark by institutional investors. In particular, BlackRock offers several ESG index funds that track the MSCI ESG indexes.

In May 2018, MSCI modified the methodology for one of its ESG indexes, the MSCI ESG Leaders Index. Prior to that, the universe of firms eligible for index inclusion consisted of all firms with an MSCI ESG rating of 'B' or higher. In May 2018, MSCI changed the eligibility rule requiring that firms have an MSCI ESG rating of at least 'BB' to be eligible to be included in its MSCI ESG Leaders. The change forced 'B'-rated firms to be excluded from the index. In turn, ESG index funds that track indexes based on the MSCI ESG Leaders index are likely to exit their position from such firms.

I use this change as a quasi-exogenous shock to a firm's ESG ownership. In particular, following the change, firms that were rated as 'B' in the May of 2018 are excluded from the index, and, as a result, index funds that invested in those firms are forced to divest from them. In contrast, firms that had ESG ratings other than 'B' do not experience such exclusion.

Empirically, to identify the effect of the ESG index exclusion on a firm's policy, I utilize a difference-in-differences framework by comparing the evolution of toxic releases at treated and control firms before and after the treatment. The treated firms include those that had an MSCI ESG rating of 'B' in May 2018. When selecting a set of appropriate control firms, it is important to ensure that they have similar evolution of toxic releases in the absence of the treatment. By performing several tests, I conclude that firms that had a rating of 'BB' serve as the most appropriate control group. They are both similar to treated firms in their observable characteristics and had a very similar evolution of toxic pollution prior to 2018.

To estimate the effect of index exclusion on the levels of toxic pollution, I estimate the following specification:

$$y_{it} = \mu_{j(i)} + \lambda \cdot post_t + \gamma \cdot treated_i + \beta \cdot treated_i \cdot post_t + \epsilon_{it} \tag{2.1}$$

, where y_{it} is the outcome of interest, $\mu_{j(i)}$ are the industry-fixed effects defined by NAICS code. $post_t$ is a dummy that equals one for 2018 and the following years. $treated_i$ is a dummy that equals one if a firm had an MSCI ESG rating of 'B' in May 2018. ϵ_{it} is the error term. The β is the coefficient of interest, as it captures the average treatment effect on treated under the identifying assumptions discussed below.

For the estimate of β from the above specification to capture the effect of index exclusion on pollution, one must carefully select the control group. One possibility would be to select all firms that were not rated 'B' in May 2018 as the control group. The downside of such an approach is that some of the firms in the control group would be very different from those in the treatment group on unobservable dimensions. For example, firms that had an MSCI ESG rating of 'AAA' in May 2018 may have taken specific actions to reduce their level of pollution. Thus the evolution of their pollution over time would be very different from the 'B'-rated group in the absence of treatment. For these reasons, in my main analysis, I use firms with ratings of ‘BB’ as the control group.

B. Identifying Assumptions

To identify the causal effect of a firm’s exclusion from the ESG index on its environmental policies, the shock must be orthogonal to the error term in the above specification. Put differently, one must argue that the MSCI’s decision to divest from ‘B’-rated firms was not caused by factors affecting the firm’s pollution decisions.

The above assumptions may be violated if the decision to modify the MSCI ESG Leaders index methodology comes as a result of pressure from firms. To investigate this possibility, I perform a LexisNexis search on the MSCI’s decision to modify its MSCI ESG Leaders Index methodology. I fail to find any mentions of it in the press prior to 2018. To investigate further, I perform a search of press releases by MSCI. I find that MSCI reached out to its clients (institutional investors) to provide input on the construction of its indexes. However, I do not find evidence that it asked its clients about the eligibility criteria.

Finally, the outcomes for the treated and control groups must satisfy the parallel trends assumption. The assumption states that in the absence of treatment evolution of outcome variables for both groups should be similar. To provide some evidence on this issue, I estimate a fully dynamic specification that includes both leads and lags of the treatment at each time horizon and then test for the presence of pre-trends. I report the findings in the Results section, where I show that there is no evidence that trends in toxic releases are different for the two groups before the treatment.

C. Treated and Control Firms

Next, I investigate whether treated and control firms are similar in the observable dimensions before the treatment. One possible concern is that treated and control group firms are different on observable variables, which may, in turn, drive the difference in outcomes. To investigate this possibility, I report the means of several variables for both groups in Table 2.2. For each group, I take the cross-section of firms from the group prior to the treatment (in years before 2018). I compute the average values of several variables within each group. Additionally, I perform a t-test for the equality of means across the groups.

Several interesting patterns emerge. First, I find that firms in the treated group pollute more than the ones in the control group. 'B'-rated firms release, on average, 567,000 more pounds of toxic chemicals than firms in the control group. This difference is statistically significant at the 5% level. Such a result is to be expected since MSCI takes the levels of toxic releases as one of the inputs to its scoring methodology. As a result, firms that pollute more are more likely to have lower ESG scores.

Next, I investigate whether firms in the two groups differ in their operating activities and capital structure. I find that firms in the treated group have slightly lower cash-to-asset ratios, although the difference is not statistically significant at the 10% level. They also have higher leverage, but this difference is also not statistically significant. In terms of their size, treated firms are smaller on average: their market value of equity is lower by $6.6 billion, and they have fewer total assets. Both of these differences are statistically significant at the 1% level.

Additionally, firms in the treated group have lower levels of institutional ownership. They have 2.05% lower institutional ownership, and 0.08% lower ownership by ESG funds. Both of these differences are statistically significant at the 1% level.

Finally, firms in the treated group have, on average, 2.2 fewer plants. However, there is no statistically significant difference in their productivity growth or their propensity to open or shut down plants.

To summarize, I conclude that treated firms are smaller on average and have lower levels of institutional ownership. However, they are similar across other observable dimensions.

2.4 Results

A. Firm's Ownership

I start the analysis by examining the effect of the MSCI ESG Leaders index exclusion on a firm's ownership structure. The primary purpose is this exercise is twofold. First, I attempt to verify whether the index exclusion caused ESG index funds that track the index to divest from the affected firms. Second, if the index exclusion caused the above divestment, it is important to understand how it affected the ownership of firms by other types of shareholders. In particular, the interpretation of my main results would depend on whether the index exclusion affected the total levels of institutional ownership.

I report the estimates in Table 2.3. The dependent variable in Column (1) is the log of the fraction of a firm's equity held by ESG index funds. The estimate on the interaction terms is negative and statistically significant at the 5% level. The exclusion of a firm from the MSCI ESG Leaders index led to a 19.6% reduction in the size of equity stakes held by ESG index funds. Interestingly, the estimates on the post dummy indicate that the size of holdings of ESG index funds has increased remarkably since 2018.

To further investigate the source of such reduction in ownership by ESG index funds, I consider whether firms are less likely to be held by ESG index funds that track one of the MSCI ESG indexes. In Column (2), the dependent variable is a dummy that equals one if any such fund holds a firm's equity. I use the dummy variable instead of the size of the holdings since, following the treatment, a firm may no longer be held by MSCI ESG index funds, making using logs problematic. The coefficient on the interaction term is negative and statistically significant at the 1% level. Index exclusion leads to a 20% lower probability that a firm is held by an ESG index fund that tracks one of the MSCI ESG indexes. In the full sample, the mean of this dummy variable is 33.7%, meaning that the effect is large in relative terms.

Additionally, I examine how the index exclusion affects the likelihood that a firm is held by an ESG index fund that tracks the MSCI ESG Leaders index. I find that, following the exclusion, treated firms are 4.7% less likely to be held by funds that track the MSCI ESG Leaders index. The estimate is statistically significant at the 1% level. In the full sample, only 8.6% of the firms are held by such indexes.

To summarize, MSCI ESG Leaders index exclusion led to a significant decrease in ownership of firms by ESG index funds. Next, I examine how the treatment affected levels of ownership by other types of institutional investors. This exercise is important since index exclusion can affect not only the composition but also the levels of institutional ownership.

If it not only lowered the ownership by ESG index funds but also by other institutions, it is more challenging to identify which of the two effects drives my main findings.

I report the estimates in Table 2.4. In Column (1), I use the log of the fraction of a firm's equity held by all mutual funds. The estimate on the interaction term is 1.8%, and it is not statistically significant at the 10% level. This result indicates that index exclusion did not significantly change the total fraction of a firm's equity held by mutual funds. Next, in Column (2), I use the log of the fraction of a firm's equity held by index funds. The estimate on the interaction term is 2.5%, also not statistically significant at the 10% level.

The results above indicate that total levels of institutional ownership and ownership by index funds did not significantly change after the treatment. To provide additional evidence, I examine how ownership by the 'Big Three: BlackRock, Vanguard, and State Street Global Advisors (SSGA) respond to the index exclusion. In Columns (3), (4), and (5), the dependent variables are the fractions of a firm's equity held by funds sponsored by BlackRock, Vanguard, and SSGA, respectively. The estimate on the interaction term for BlackRock in Column (3) is close to zero and not statistically significant. For Vanguard in Column (4), the estimate is -0.2% and statistically significant at the 10% level. For SSGA in Column (5), the estimate on the interaction term is 0.1%, not statistically significant at the 10% level. In sum, I do not find evidence that index exclusion had a meaningful impact on the sizes of stakes held by the Big Three fund families.

To summarize, I find that the MSCI's change in its ESG Leaders index methodology had a meaningful impact on the ownership structure of firms. Following the shock, treated firms had lower ownership by ESG funds, particularly ESG index funds that track MSCI ESG indexes. At the same time, the shock did not affect the total level of institutional ownership, the total level of ownership by index funds, or the ownership for any of the 'Big Three' fund families. In sum, these findings indicate that the exclusion from the MSCI ESG Leaders index affected the composition of a firm's ownership but not the total fraction of the firm's equity held by mutual funds. The next section examines how the index exclusion affects a firm's toxic releases.

B. Toxic Pollution

In this subsection, I examine the impact of the ESG index exclusion on firms' toxic pollution. Theoretically, the effect of an ESG index exclusion on a firm's toxic pollution is unclear. On the one hand, since index exclusion reduces the amount of ESG capital invested in a firm, the firm faces less monitoring by environmentally-conscious investors. In this case, index exclusion would lead to more toxic pollution. On the other hand, if index exclusion increases the firm's cost of capital, the firms may have incentives to reduce their toxic releases to get included back in the index. Ultimately, which of these two effects is stronger is an empirical question.

As the main outcome variable, I use the log of one plus the total toxic releases from a firm's plants, scaled by the firm's total assets. In the robustness section, I verify that the results remain similar when scaling by lagged assets or costs of goods sold. Within the empirical framework described in the previous section, I compare the total amount of toxic releases by firms in the treatment and control groups before and after the treatment. First, I need to verify that in the pre-treatment period, total toxic releases evolved similarly in both groups. To do so, I estimate the following specification:

$$\log\left(1+\frac{ToxicReleases_{it}}{TotalAssets_{it}}\right) = \sum_{k=-6;k\neq -1}^{3} \beta_k treated_i \cdot \mathbf{1}[t-2018=k] + \\ + \gamma \cdot treated_i + \lambda_t + \alpha_{j(i)} + \epsilon_{it} \tag{2.2}$$

, where $ToxicReleases_{it}$ are the total toxic releases of firm i's plant in year t, and $TotalAssets_{it}$ are the firm's total assets. $\alpha_{j(i)}$ and λ_t are the industry and time fixed effects. $treated_i$ is a dummy that equals one for firms whose MSCI ESG rating is 'B' in May 2018. The sample consists of all firms with a rating of 'B' or 'BB' in May 2018.

The estimates $\hat{\beta}_k$ from the above regression yield the difference in the evolution of the outcome variable between the treated and the control groups k years following the treatment. Estimates with $k < 0$ allow me to test whether the toxic releases evolved differently before the treatment. Estimates with $k \geq 0$ capture treatment effect k periods after the treatment. I normalize the coefficient β_{-1} to zero.

I plot the estimates $\hat{\beta}_k$ on Figure 2.1. For each k, I plot the corresponding $\hat{\beta}_k$ and its 95% confidence interval. First, the estimates with $k < 0$ are small in magnitude and are not statistically significant. This result highlights that I can not reject the null hypothesis that any of the $\hat{\beta}_k$ with $k < 0$ is statistically different from zero. I perform the F-test for the joint significance of such coefficients. The p-value of the test of 0.97, meaning that I fail to find evidence of pre-trends in the outcome variable for the treatment and the control group.

Next, I examine how toxic releases evolve following the treatment. The estimates for β_1 and β_2 are both positive and statistically significant at the 1% level. The magnitude of the estimates is also very large: index exclusion caused a 0.3 increase in the outcome variable, which corresponds to a roughly 35% increase. The estimates decline somewhat in 2021.

To provide additional estimates, I estimate my main specification with the log of one plus scaled releases as my outcome variable. I report the estimates in Table 2.5. The estimate on the interaction term in Column (1) is 0.253, statistically significant at the 5% level. This number corresponds to a 28.7% increase in the scaled toxic releases for treated firms relative to the firms in the control group.

To get a more detailed view of the effect of a shock on a firm's polluting activities, I study how the index exclusion affected different types of releases, namely releases into the air, ground, and water. The dependent variables in Columns (2), (3), and (4) of Table 2.5 are the log of one plus scaled air, ground, and water releases, respectively. Interestingly, the estimates on the interaction terms in Columns (3) and (4) are not statistically significant at the 10% level and are small in magnitude. In contrast, the estimate in Column (2) is statistically significant at the 5% level, and its magnitude is close to the one in Column (1). These results indicate that the major source of an increase in toxic releases from treated firms is higher releases into the air. Such findings can be explained by the fact that MSCI heavily weighs carbon emissions in their ESG index construction, which can force firms held by funds that track the index to lower their air emissions.

To summarize, I find that a quasi-exogenous exclusion of firms from the ESG index increases the total amount of toxic releases by firms. Estimated treatment effects are both large in magnitude and statistically significant. In the following section, I investigate the possible mechanism behind this result.

C. Sources of an Increase in Toxic Releases

1. Abatement Activities In the next set of results, I attempt to identify the source of an increase in toxic releases for treated firms. There are several possible explanations. First, it is possible that firms excluded from the index decreased their abatement activities compared to the firms in the control group. This can happen if firms held by ESG investors face the pressure to invest in cleaner production. Second, the effect can be explained by a change in productivity. For instance, firms monitored by ESG investors can face pressure to reduce pollution, which may cause them to reduce output. I examine each of the above possibilities below.

I start by examining the effect of the index exclusion on a firm's abatement activity. I focus on abatement activities reported to the EPA's pollution prevention (P2) database. Following Akey and Appel (2021), I focus on two main types of abatement: changes in operating practices and process improvements. The former can include reusing materials for longer periods of time to reduce waste. The latter can be implemented by improving control over leaks and modifying production to produce less waste.

For each firm in the sample, I create dummy variables equal to one if the firm implemented abatement at any of its plants. I report the estimates in Table 2.6. In Column (1), I focus on operating abatement. The coefficient on the interaction term is positive but statistically insignificant at the 10% level. Similarly, the estimate of the treatment effect of index exclusion on process abatement is not statistically significant in Column (2). Based on the above evidence, it is unlikely that treated firms increase their toxic releases due to fewer abatement activities. Another possibility is that firms shift toxic waste to offsite facilities. To investigate this, I examine how the offsite releases change following the treatment. Based on the estimate in Column (3), I fail to find any evidence that firms change their offsite releases following index exclusion.

2. Productivity

Another explanation for my main results is a change in productivity. In particular, plants owned by treated firms may increase their output following the index exclusion, which leads to higher toxic releases.

I examine this possibility using the production ratio reported in the EPA TRI database. TRI reports the ratio of a plant's output in a given year relative to the previous year. Since I conduct my analysis at the firm level, for each firm, I aggregate productivity growth across all its plants in a given year. I weigh each plant's productivity ratio by its total releases and average it for each firm. I conduct such weighting since plants with higher releases are likely to produce more.

Additionally, I analyze the cumulative production ratio used in Akey and Appel (2021). For each firm, I set its productivity level to one in the first year of my sample, which is 2011. Then for each year, I multiply the firm-level production growth defined above in prior years:

$$CumulativeProduction_{it} = \prod_{\tau=2012}^{t} ProductionGrowth_{i\tau} \tag{2.3}$$

I report the estimates of the effect of index exclusion on productivity in Table 2.7. In Column (1), the dependent variable is production growth. The estimate on the interaction term is a 1.3% reduction in productivity growth, and it is not statistically significant at the 10% level. The estimate of the effect of index exclusion on the cumulative productivity in Column (2) is positive but also not statistically different from zero at the 10% level. Based on these results, I conclude that index exclusion did not significantly affect a firm's plant productivity.

3. Change in the Number of Plants

Finally, I investigate whether there was a change in the number of plants following the MSCI ESG Leaders index exclusion. It is possible that following the treatment, firms change their investment strategy and open new plants or close the existing ones. For example, if the firm faced pressure from the shareholders regarding its toxic releases, the management may have been reluctant to extend operations if it meant higher pollution levels.

To investigate this further, I create two dummy variables that capture a change in the number of facilities reporting toxic releases for each firm. I report the estimates in Table 2.8. The dependent variable in Column (1) is a dummy that equals one if a firm has more reporting plants in the current year compared to the previous year. Based on the coefficient on the interaction term, treated firms are 8.3% more likely to expand their operations by increasing the number of polluting factories. The estimate is statistically significant at the 1% level. In Column (2), the dependent variable is a dummy if a firm has fewer reporting facilities than the previous year. The estimate of the treatment effect is a 5.3% reduction in the propensity to reduce the number of plants, although it is not statistically significant.

Based on the above evidence, the increase in the toxic releases for the treated firms is explained on the extensive margin by an increase in the number of their facilities. While firms do not change their abatement activities and their productivity remains similar, they open new plants, which increases the total releases.

2.5 Mechanism

A. Cost of Capital Channel

In the results above, I establish that firms that are excluded from the ESG index increase their toxic releases. In this subsection, I investigate the possible mechanisms behind this finding.

The theory literature establishes several possible mechanisms that can explain the above finding. First, the index exclusion may affect the firms' pollution activities through the cost of capital. In particular, if divestment increases the cost of capital for the firms, they may find it difficult to invest in new projects, particularly those that lead to lower pollution. Another possible explanation for the increase in pollution after ESG index exclusion is a decrease in shareholder monitoring. If environmentally-conscious investors sell their shares in a firm, the firm can lower its investment in green projects, leading to higher pollution. I examine each of the above channels below.

To investigate the cost of capital channel, I study how the index exclusion affects the valuation of firms. If the cost of capital for firms changed, it is likely that market to book ratio also changed. In Column (1) of Table 2.9 I estimate the effect of index exclusion on Q. I find that Tobin's Q increases by 0.024, but the estimate is not statistically significant. To investigate further, I focus on the firm's PPE and total assets. One of the findings in the results section is that treated firms increase the number of polluting plants, which causes the total toxic releases to go up as well. In Columns (2) and (3), I use the logs of PPE and total assets as the dependent variables. I fail to find evidence that assets or the book value property, plants, and equipment change significantly following the index exclusion.

Finally, I evaluate the effect of index exclusion on operating performance. If it indeed caused the cost of capital to go up, this may end up having an adverse effect on a firm's performance. In particular, it may be more costly to raise financing for future projects, leading to lower output. In Column (4), I use the log of sales over assets as the main dependent variable. I find that following the index exclusion, sales increase slightly, but the effect is not statistically significant at the 10% level. Additionally, in Column (5), I document that there is no significant effect of index exclusion on leverage, implying that firms did not change their funding source in response to the shock.

To summarize, I fail to find empirical support for the hypothesis that the ESG index exclusion affected firms' policies through the cost of capital channel. These findings align with the main results in a recent paper by Berk and van Binsbergen (2021). They examine how exclusion from another popular ESG index, FTSE USA 4Good, affects the cost of capital of excluded firms. They find no such effect in their paper. Intuitively, they argue that given the current amount of capital invested in ESG funds, the effects of divestment by ESG funds on the cost of capital should be negligible.

B. Shareholder Monitoring Channel

An alternative explanation for my findings is that following index exclusion, firms' managers face less monitoring from shareholders with ESG preferences. As a result, managers shift away from green projects, and their firms have higher levels of toxic releases.

While such a mechanism is well-documented in various settings in corporate governance literature (Yermack, 2010), it does not seem plausible in my setting. In particular, the effect of the ESG index exclusion on the fraction of the firm's equity held by ESG funds is relatively small in magnitude. Additionally, ESG funds hold a relatively small stake in their portfolio firms. For example, they are unlikely to be pivotal and affect the outcome of shareholder proposals.

To explain why even small ESG funds may affect the pollution of their portfolio companies, I build on the insight from Lakkis (2021). Lakkis (2021) shows that due to coordination in voting within fund families (Morgan et al., 2011), holdings of a family's index funds affect how other funds from the family vote. In the context of this paper, holdings of ESG funds at the fund family level may affect the incentives of the family to monitor the management on the E and S dimensions.[5] While holdings of ESG funds are still moderate, many of them are sponsored by large fund families, and those families can hold enough shares to affect corporate policies through the voice channel.[6]

Since the amount of ESG capital has increased substantially in recent years, investment managers have incentives to attract future flows into ESG funds. One of the ways they can do so is by affecting the ESG performance of their portfolio companies, which they can then advertise to investors. Once a firm is excluded from the ESG index, it is less likely to be held by a family's ESG funds. In turn, the family has fewer incentives to use holdings of its non-ESG funds to promote the ESG agenda at the firm.

To test the above channel, I look for a cross-sectional variation in the exposure of large fund families to this shock. Interestingly, I find that BlackRock is the only of the Big Three fund families that sponsor funds that track MSCI ESG indexes. Building on this finding, I argue that the MSCI ESG Leaders index exclusion should affect the incentives of BlackRock to engage on ESG issues with treated firms. ESG funds sponsored by BlackRock had to divest from treated firms, lowering the fraction of within-family ESG capital at these firms. At the same time, the shock should not affect the incentives of other large fund families since

[5] Michaely, Ordonez-Calafi, and Rubio (2021) show that ESG funds vote strategically and incorporate the ESG preferences of their fund families in their votes on close proposals.

[6] For example, Gormley et al. (2021) show that 'Big Three' fund families successfully increase corporate board gender diversity at target firms. Azar et al. (2021) document that the 'Big Three' can force firms to reduce carbon emissions through private engagements.

the index exclusions did not change their portfolios.

I also note that institutional investors have more incentives to engage and monitor when their holdings in a firm are larger (Iliev, and Lowry, 2015). Taken combined, the two observations above lead to a hypothesis that the effect of ESG index exclusion on toxic releases should be stronger for firms that had high ownership by BlackRock's funds. At the same time, I do not expect to find differences in the treatment effects across high and low levels of ownership by Vanguard and SSGA.

To empirically test the above hypothesis, I utilize the triple-difference regression. More specifically, I compare the magnitude of treatment effects between high and low levels of a fund family's ownership in a firm. I define the ownership of a family to be high if the holdings of the family in a firm are above the sample median.

I report the estimates of a triple-difference regression in Table 2.10. In Column (1), I perform a triple-difference analysis using levels of ownership stakes held by BlackRock. I find that the estimate on the interaction treated · post · high is positive and statistically significant at the 5% level. Firms with a high level of ownership by BlackRock funds have a treatment effect that is higher by 0.667. In Column (2), I repeat the analysis for Vanguard. I find that the coefficient estimate on the triple interaction term is negative but not statistically significant. Such findings show that there is no statistically significant difference in the magnitude of treatment effect between high and low levels of ownership by Vanguard. Similarly, in Column (3), I repeat the analysis for the SSGA and do not find any meaningful difference in treatment effect estimates for high and low levels of ownership by SSGA.

The above findings provide evidence that firms increase toxic releases since they face less monitoring by BlackRock. To provide further insight into how this happens, I examine how shareholder votes at treated firms change. Suppose there is a significant effect of index exclusion on the levels of support the management receives in voting. In that case, my main finding can be explained by observable efforts on behalf of BlackRock. Alternatively, if I fail to find any effect on votes cast, this would suggest that most of such engagement happens behind the scenes (McCahery, Sautner, and Starks, 2016).

To summarize, I find that the effect of the MSCI ESG Leaders index exclusion on a firm's toxic releases is entirely driven by a subsample of firms that had a high level of BlackRock ownership. Triple-difference estimates using the size of ownership by Vanguard and SSGA do not uncover any meaningful variation in the treatment effect estimates. Based on these results, I conclude that there is empirical support for a fund family-based explanation of my main result. As a firm is excluded from the ESG index, BlackRock's ESG funds have lower levels of ESG ownership in a firm. In turn, other funds in BlackRock have fewer incentives to monitor the firm's environmental policies, particularly toxic pollution.

2.6 Robustness

In this section, I verify that my results are robust to alternative scaling of the toxic releases. In my main analysis, I scale the total releases by total assets. One potential concern is that total assets can also respond to the index exclusion, affecting my main results. To investigate this possibility further by running the analysis using the lagged total assets as the scaling factor. I report the estimates in Table 2.12. In Column (1), I report the results from the main specification as a reference. In Column (2), I use the lagged total assets as the scaling factor. The estimate remains almost identical both in terms of its magnitude and statistical significance.

As an alternative scaling, I also use the costs of goods sold. Unlike assets, costs of goods sold can be thought of as a measure of production activity, which makes it attractive for use as a scaling factor. I scale releases by both costs of goods sold and their lag and report the estimates in Columns (3) and (4), respectively. I find that the estimates remain very similar in their magnitude.

2.7 Conclusion

In this paper, I investigate the effect of ESG mutual funds' investment on toxic pollution by firms. I rely on a quasi-exogenous exclusion from the MSCI ESG Leaders index as a source of variation in holdings by ESG funds. I find that firms open new plants and increase their toxic releases following the index exclusion. I explain this result by lower incentives of fund families to monitor the environmental performance of firms after a decrease in holdings of ESG funds in the families.

The paper's main results contribute to a growing body of literature studying the real effects of ESG investment. While several papers argue that the amount of ESG capital is still moderate in size to have a meaningful impact on firms' policies, I can still detect the effect of index exclusion on firms' pollution. I propose a novel explanation for why even a slight decrease in the amount of ESG capital invested in a firm can affect a firm's policies. As the amount of a fund family's ESG holdings in a firm decreases, other funds from that family have fewer incentives to monitor the firm's environmental performance.

While I primarily focus on the environmental outcomes in this paper, it is unclear whether firms change their policies with respect to the S and G dimensions of the ESG following their exclusion from an ESG index. While those dimensions are harder to measure, studying the effects of ESG investments on social outcomes can be a fruitful avenue for future research.

2.8 Tables and Figures

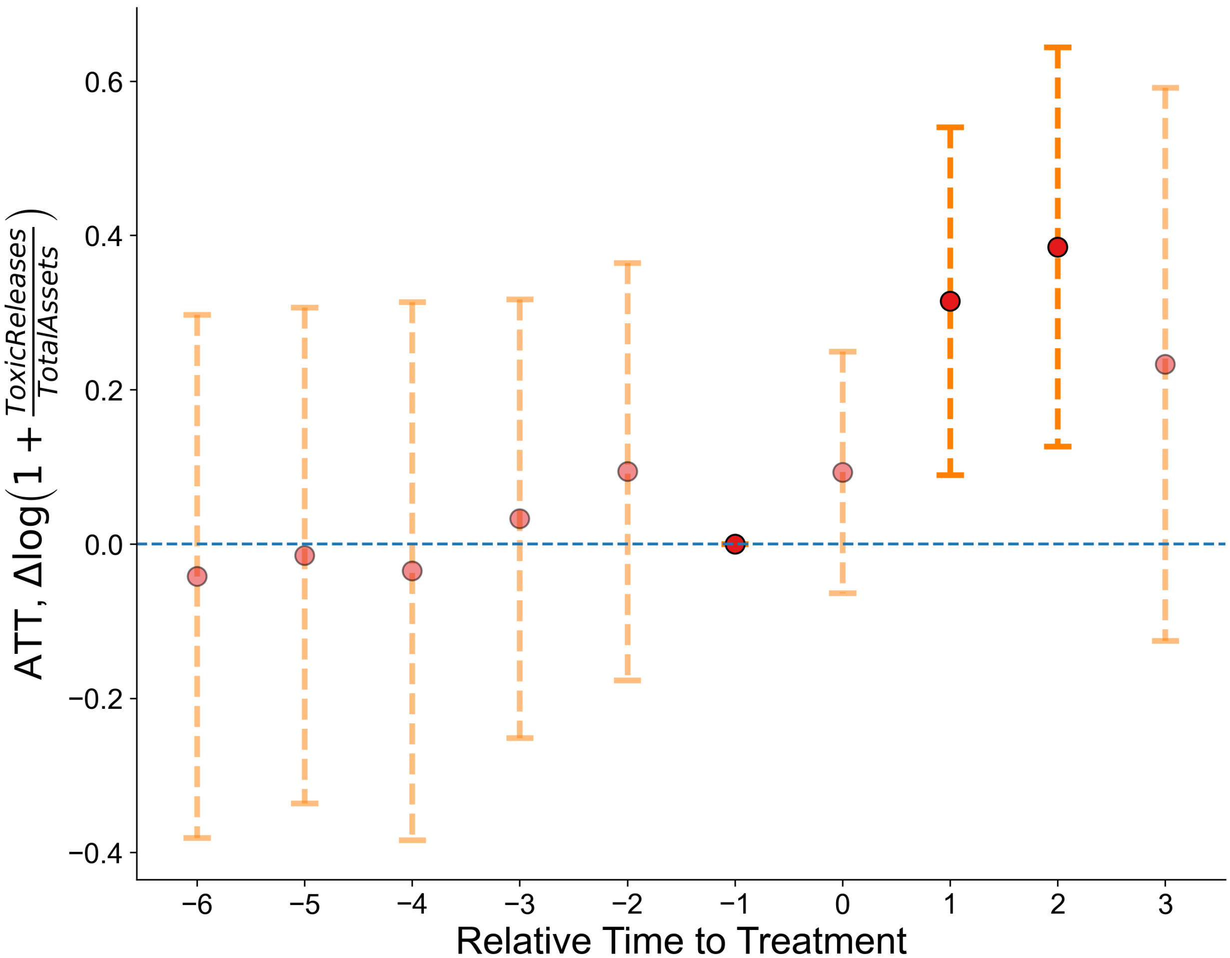

The figure shows the dynamic treatment effect of the exclusion from the MSCI ESG Leaders Index on toxic releases estimated from the following specification:

$$\log\left(1+\frac{ToxicReleases_{it}}{TotalAssets_{it}}\right) = \sum_{k=-6;k\neq-1}^{3} \beta_k treated_i \cdot \mathbf{1}[t-2018=k]+ \\ + \gamma \cdot treated_i + \lambda_t + \alpha_{j(i)} + \epsilon_{it} \tag{2.4}$$

, where $ToxicReleases_{it}$ are the total toxic releases of firm i's plant in year t, and $TotalAssets_{it}$ is the firm's total assets. $\alpha_{j(i)}$ and λ_t are the industry and time fixed effects. $treated_i$ is a dummy that equals one for firms whose MSCI ESG rating is 'B' in May 2018. Vertical bars show 95% confidence intervals. The errors are clustered at the industry level.

Figure 2.1. Treatment Effect of ESG Index Exclusion on Toxic Releases

TABLE 2.1. Summary Statistics

The table reports summary statistics. The sample consists of 1,752 firms between the years 2011 and 2021. *N Plants* is the number of a firm's plants that report toxic releases in a given year. *Closed a Plant* (*Opened a Plant*) is a dummy that equals one if a firm has more (less) plants that report toxic releases in the current year compared to the previous year. *Total Releases* is the total amount of toxic chemicals released by all plants owned by a firm. *Onsite Releases* and *Offsite Releases* are combined releases at the firm's plants and shipped outside those plants, respectively. *Air Releases*, *Water Releases*, and *Ground Releases* are the total amount of chemicals released in the air, water, and the ground, respectively. *Production Growth* is the percentage growth in a firm's production weighted by releases across its plants. *Cumulative Production Growth* is the production growth accumulated from 2011. *Abatement Operating* (*Abatement Processing*) is the dummy variable that equals one if one of a firm's plants implemented an operating (processing) abatement in a given year. *Cash/Assets*, *Leverage/Assets*, and *Investment/Assets* are the firm's cash holdings, total debt, and capital expenditures scaled by total assets. *Tobin's Q* is the firm's equity value plus the value of debt, divided by total assets. *Institutional Ownership* is the percentage of a firm's outstanding equity that is held by mutual funds. *ESG Funds Ownership* is the percentage of a firm's equity that is held by ESG mutual funds. *BlackRock Ownership*, *Vanguard Ownership*, and *SSGA Ownership* are the fraction of a firm's equity that is held by funds sponsored by BlackRock, Vanguard, and SSGA, respectively. *Index Funds Ownership*, *ESG Index Funds Ownership*, and *MSCI ESG Index Funds Ownership* are the percentage of a firm's equity held by index funds, ESG index Funds, and MSCI ESG index funds, respectively.

	N	mean	std	10%	50%	90%
N Plants	4,385	10.384	17.812	1.0	5.0	24.0
Closed a Plant	4,385	0.181	0.385	0.0	0.0	1.0
Opened a Plant	4,385	0.129	0.336	0.0	0.0	1.0
Total Releases, '000 lbs	4,385	1,807.9	4,906.1	0.12	75.8	5,143.8
Onsite Releases, '000 lbs	4,385	1,422.1	3,980.7	0.002	33.2	3,727.5
Offsite Releases, '000 lbs	4,385	213.5	599.3	0.0	6.16	541.9
Air Releases, '000 lbs	4,385	484.5	1,259.9	0.001	22.5	1,351.4
Water Releases, '000 lbs	4,385	78.9	307.4	0.0	0.0	79.5
Ground Releases, '000 lbs	4,385	540.5	1,955.3	0.0	0.0	731.2
Production Growth	4,385	1.05	0.28	0.81	1.001	1.29
Cumulative Production Growth	4,385	1.593	1.834	0.502	1.013	2.99
Abatement, Operating	4,385	0.199	0.399	0.0	0.0	1.0
Abatement, Processing	4,385	0.173	0.378	0.0	0.0	1.0
Cash/Assets	19,448	0.105	0.119	0.007	0.061	0.267
Leverage/Assets	19,753	0.276	0.214	0.005	0.255	0.577
Market Value of Equity, $B	18,806	13.67	57.2	0.51	2.60	25.19
Total Assets, $B	19,810	21.58	122.3	0.424	3.26	31.44
Investment/Assets	19,735	0.032	0.036	0.0	0.021	0.08
Tobin's Q	18,753	1.878	2.264	0.349	1.305	3.764
Institutional Ownership, %	18,293	35.72	12.93	17.13	36.74	52.07
ESG Funds Ownership, %	18,293	0.403	0.629	0.003	0.149	1.148
BlackRock Ownership, %	18,293	1.208	2.623	0.0	0.073	4.141
Vanguard Ownership, %	18,293	6.361	3.799	0.0	6.366	11.163
SSGA Ownership, %	18,293	0.862	0.911	0.0	0.572	2.117
Index Funds Ownership, %	18,293	12.469	7.49	0.07	12.34	22.67
ESG Index Funds Ownership, %	18,293	0.021	0.047	0.0	0.002	0.055
MSCI ESG Index Funds Ownership, %	18,293	0.008	0.019	0.0	0.0	0.025

TABLE 2.2. Treated and Control Groups

The table compares the means of several observable variables for firms in treated and control groups before the treatment. The last column reports the difference between the means and the result of a t-test of the equality of the means between the treatment and the control group. The treatment group consists of firms that had the MSCI ESG rating of 'B' in May 2018. Firms in the control group had a rating of 'BB'. *Total Releases* is the total amount of toxic chemicals that a firm's plants release, as reported to the EPA. *Cash/Assets* is the firm's cash holdings divided by its total assets. *Leverage/Assets* is the firm's total debt liabilities divided by the total assets. *Market Value of Equity* is the firm's stock price multiplied by the number of common shares outstanding. *Investment/Assets* is the firm's capital expenditures divided by total assets. *Tobin's Q* is the firm's equity value plus the value of debt, divided by total assets. *Institutional Ownership*, *ESG Funds Ownership*, and *MSCI ESG Index Funds Ownership* are the percentages of a firm's equity held by mutual funds, ESG funds, and MSCI ESG index funds, respectively. *Production Growth* is the growth in a firm's production, weighted by toxic releases across its plants. *Number Plants* is the number of a firm's plants reporting releases in a given year. *Closed a Plant* (*Opened a Plant*) is a dummy that equals one if a firm has more (less) reporting facilities compared to the previous year. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	Treated	Control	diff
Total Releases, '000 lbs	2,309.8	1,741.8	567.0**
Cash/Assets	0.105	0.109	-0.004
Leverage/Assets	0.258	0.254	0.004
Market Value of Equity, $B	6.58	8.59	-2.02***
Total Assets, $B	11.8	21.4	-9.6***
Investment/Assets	0.033	0.033	-0.001
Tobin's Q	1.69	1.83	-0.14***
Institutional Ownership, %	32.25	34.30	-2.05***
ESG Funds Ownership, %	0.28	0.36	-0.08***
MSCI ESG Index Funds Ownership, %	0.0	0.001	-0.001***
Production Growth	1.08	1.06	0.02
Number Plants	8.99	11.24	-2.25***
Closed a Plant	0.144	0.156	-0.012
Opened a Plant	0.111	0.135	-0.024

TABLE 2.3. ESG Index Exclusion and Index Fund Ownership

The table reports the estimates of the treatment effect of MSCI ESG index exclusion on ownership of firms by ESG index funds. The dependent variable in Column (1) is the logarithm of the percentage of a firm's equity held by ESG index funds. The dependent variables in Columns (2) and (3) are dummy variables that equal one if a firm is held by MSCI ESG Index funds and MSCI ESG Leaders index funds, respectively. *treated* is a dummy that equals one if a firm had the the MSCI ESG rating of 'B' in May 2018. *post* is a dummy that equals one for 2018 and the following years. *MCap* is a firm's market capitalization. The sample covers 1,752 firms between the years 2011 and 2021. Standard errors are clustered at the industry level. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1) log ESG Index Fund Ownership	(2) Held by MSCI ESG Index Funds	(3) Held by MSCI ESG Leaders Index Funds
treated · post	-0.219**	-0.209***	-0.047***
	(0.106)	(0.023)	(0.011)
post	1.809***	0.321***	-0.045***
	(0.092)	(0.029)	(0.010)
treated	-0.279**	-0.034*	-0.000
	(0.130)	(0.018)	(0.008)
log MCap	0.520***	0.122***	-0.007***
	(0.056)	(0.007)	(0.003)
Constant	-14.775***	-0.673***	0.166***
	(0.449)	(0.057)	(0.022)
Observations	5,461	9,753	9,753
R-squared	0.500	0.333	0.074
Industry FE	Yes	Yes	Yes

TABLE 2.4. ESG Index Exclusion and Institutional Ownership

The table reports the estimates of the treatment effect of MSCI ESG index exclusion on a firm's institutional ownership. The dependent variables in Columns (1) and (2) are the logarithms of the fraction of equity held by mutual funds and index funds, respectively. The dependent variables in Columns (3), (4), and (5) are the percentages of a firm's equity held by funds sponsored by BlackRock, Vanguard, and SSGA, respectively. *treated* is a dummy that equals one if a firm had the MSCI ESG rating of 'B' in May 2018. *post* is a dummy that equals one for 2018 and the following years. *MCap* is a firm's market capitalization. The sample covers 1,752 firms between the years 2011 and 2021. Standard errors are clustered at the industry level. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1) log Institutional Ownership	(2) log Index Fund Ownership	(3) % BlackRock	(4) % Vanguard	(5) % SSGA
treated · post	0.018	0.025	-0.000	-0.002*	0.001
	(0.025)	(0.038)	(0.001)	(0.001)	(0.000)
post	0.147***	0.454***	0.028***	0.002	0.001
	(0.018)	(0.029)	(0.001)	(0.001)	(0.001)
treated	-0.075	-0.070	-0.001*	-0.002	-0.002**
	(0.047)	(0.073)	(0.000)	(0.002)	(0.001)
log MCap	0.085***	0.116***	-0.002***	0.006***	0.002***
	(0.015)	(0.018)	(0.000)	(0.001)	(0.000)
Constant	-1.836***	-3.217***	0.014***	0.012***	-0.010***
	(0.106)	(0.127)	(0.003)	(0.004)	(0.002)
Observations	9,017	8,087	9,017	9,017	9,017
R-squared	0.387	0.386	0.266	0.269	0.333
Industry FE	Yes	Yes	Yes	Yes	Yes

TABLE 2.5. ESG Index Exclusion and Toxic Pollution

The table reports the estimates of the treatment effect of MSCI ESG index exclusion on a firm's toxic releases. The dependent variable in Column (1) is the logarithm of one plus a firm's total toxic releases scaled by its total assets. The dependent variables in Columns (2), (3), and (4) are the logarithms of one plus a firm's releases into air, ground, and water, respectively, scaled by the firm's total assets. *treated* is a dummy that equals one if a firm had the MSCI ESG rating of 'B' in May 2018. *post* is a dummy that equals one for 2018 and the following years. *MCap* is a firm's market capitalization. The sample covers 965 firms that had the MSCI ESG rating of 'B' or 'BB' as of May 2018. Firms rated 'B' are treated as they were excluded from the MSCI index. Standard errors are clustered at the industry level. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1) log(1 + Total Releases/Total Assets)	(2) log(1 + Air Releases/Total Assets)	(3) log(1 + Ground Releases/Total Assets)	(4) log(1 + Water Releases/Total Assets)
treated · post	0.253**	0.253**	-0.071	-0.000
	(0.118)	(0.113)	(0.079)	(0.000)
post	-0.497***	-0.462***	-0.029	-0.000**
	(0.090)	(0.080)	(0.053)	(0.000)
treated	-0.171	0.017	0.244	0.000
	(0.421)	(0.422)	(0.324)	(0.000)
Constant	3.610***	2.470***	0.827***	0.001***
	(0.221)	(0.220)	(0.164)	(0.000)
Observations	1,974	1,974	1,974	1,974
R-squared	0.794	0.735	0.817	0.727
Industry FE	Yes	Yes	Yes	Yes

TABLE 2.6. ESG Index Exclusion and Abatement Activity

The table reports the estimates of the treatment effect of MSCI ESG index exclusion on a firm's abatement activities at its plants. The dependent variable in Column (1) is a dummy variable that equals one if a firm implemented operating abatement at any of its plants in a given year. Similarly, the dependent variable in Column (2) is a dummy that equals one if a firm implemented processing abatement. The dependent variable in Column (3) is the logarithm of one plus a firm's offsite releases scaled by total assets. *treated* is a dummy that equals one if a firm had the MSCI ESG rating of 'B' in May 2018. *post* is a dummy that equals one for 2018 and the following years. *MCap* is a firm's market capitalization. The sample covers 965 firms that had the MSCI ESG rating of 'B' or 'BB' as of May 2018. Firms rated 'B' are treated as they were excluded from the MSCI index. Standard errors are clustered at the industry level. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1) Operating Abatement	(2) Process Abatement	(3) log(1 + Offsite Releases/TotalAssets)
treated · post	0.013	0.028	0.040
	(0.039)	(0.039)	(0.135)
post	-0.140***	-0.118***	-0.221**
	(0.027)	(0.026)	(0.099)
treated	-0.054	-0.032	-0.454*
	(0.057)	(0.043)	(0.262)
Constant	0.248***	0.211***	1.979***
	(0.030)	(0.025)	(0.138)
Observations	1,974	1,974	1,974
R-squared	0.320	0.272	0.718
Industry FE	Yes	Yes	Yes

TABLE 2.7. ESG Index Exclusion and Plants Productivity

The table reports the estimates of the treatment effect of MSCI ESG index exclusion on the productivity of a firm's plants. The dependent variable in Column (1) is the total release-weighted average of production growth across the firm's plants. The dependent variable in Column (2) is the cumulative productivity growth for a firm relative to 2011. *treated* is a dummy that equals one if a firm had the MSCI ESG rating of 'B' in May 2018. *post* is a dummy that equals one for 2018 and the following years. *MCap* is a firm's market capitalization. The sample covers 965 firms that had the MSCI ESG rating of 'B' or 'BB' as of May 2018. Firms rated 'B' are treated as they were excluded from the MSCI index. Standard errors are clustered at the industry level. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1) Production	(2) Cumulative Production
treated · post	-0.013	0.072
	(0.026)	(0.212)
post	-0.022	0.448***
	(0.019)	(0.149)
treated	-0.025	0.029
	(0.027)	(0.444)
Constant	1.084***	1.466***
	(0.015)	(0.224)
Observations	1,974	1,974
R-squared	0.184	0.446
Industry FE	Yes	Yes

TABLE 2.8. ESG Index Exclusion and Closure of Plants

The table reports the estimates of the treatment effect of MSCI ESG index exclusion on a firm's decision to open or close new plants. The dependent variable in Column (1) is a dummy variable that equals one if a firm has more plant reporting releases compared to the previous year. The dependent variable in Column (2) is a dummy variable that equals one if a firm has fewer plants reporting releases compared to the previous year. *treated* is a dummy that equals one if a firm had the MSCI ESG rating of 'B' in May 2018. *post* is a dummy that equals one for 2018 and the following years. *MCap* is a firm's market capitalization. The sample covers 965 firms that had the MSCI ESG rating of 'B' or 'BB' as of May 2018. Firms rated 'B' are treated as they were excluded from the MSCI index. Standard errors are clustered at the industry level. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1) Open Plant	(2) Close Plant
treated · post	0.083***	-0.053
	(0.024)	(0.036)
post	-0.051***	0.065**
	(0.019)	(0.029)
treated	-0.011	0.005
	(0.036)	(0.051)
Constant	0.140***	0.138***
	(0.021)	(0.027)
Observations	1,974	1,974
R-squared	0.156	0.119
Industry FE	Yes	Yes

TABLE 2.9. ESG Index Exclusion and a Firm's Assets and Performance

The table reports the estimates of the treatment effect of MSCI ESG index exclusion on a firm's assets, leverage, and operating performance. The dependent variable in Column(1) is Q, the ratio of a firm's total market value to total assets. The dependent variables in Columns (2) and (3) are log of the value of property, plant, and equipment, and the log of total assets, respectively. The dependent variable in Column (4) is log of a firm's sales scaled by its total assets. The dependent variable in Column (5) is a firm's leverage, the ratio of a firm's total liabilities over the total assets. *treated* is a dummy that equals one if a firm had the MSCI ESG rating of 'B' in May 2018. *post* is a dummy that equals one for 2018 and the following years. The sample covers 965 firms that had the MSCI ESG rating of 'B' or 'BB' as of May 2018. Firms rated 'B' are treated as they were excluded from the MSCI index. Standard errors are clustered at the industry level. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1) Q	(2) log PPE	(3) log Assets	(4) log Sales/Assets	(5) Leverage
treated · post	0.024	-0.003	0.014	0.031	-0.003
	(0.121)	(0.073)	(0.050)	(0.024)	(0.010)
post	0.044	0.527***	0.474***	-0.117***	0.049***
	(0.103)	(0.053)	(0.039)	(0.024)	(0.013)
treated	-0.178	-0.015	-0.230	-0.010	-0.009
	(0.142)	(0.201)	(0.154)	(0.041)	(0.009)
Constant	1.871***	6.221***	7.893***	-0.966***	0.267***
	(0.072)	(0.100)	(0.076)	(0.021)	(0.005)
Observations	9,721	7,935	10,506	10,458	10,473
R-squared	0.352	0.616	0.498	0.798	0.548
Industry FE	Yes	Yes	Yes	Yes	Yes

TABLE 2.10. ESG Index Exclusion and a Firm's toxic releases: Effect of Fund Family Ownership

The table reports the estimates from a triple-difference regression:

$$\log\left(1+\frac{TotalReleases_{it}}{Assets_{it}}\right) = \alpha_1 post_t + \alpha_2 treated_i + \alpha_3 high_{it} + \gamma_1 post_t \cdot high_{it} + \gamma_2 post_t \cdot treated_i + \gamma_3 treated_i \cdot post_t + + \beta post_t \cdot treated_i \cdot high_{it} + \mu_{j(i)} + \epsilon_{it} \tag{2.5}$$

where *treated* is a dummy that equals one if a firm had the MSCI ESG rating of 'B' in May 2018. *post* is a dummy that equals one for 2018 and the following years. *high* is a dummy that equals one if a given fund family has holdings in the firm that is above the sample median. $\mu_{j(i)}$ are the industry-fixed effects. The estimates in Columns (1), (2), and (3) are using the *high* variable for BlackRock, Vanguard, and SSGA, respectively. The sample covers 965 firms that had the MSCI ESG rating of 'B' or 'BB' as of May 2018. Firms rated 'B' are treated as they were excluded from the MSCI index. Standard errors are clustered at the industry level. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1) BlackRock	(2) Vanguard	(3) SSGA
treated · post · high	0.677**	-0.279	0.086
	(0.322)	(0.196)	(0.279)
treated · post	-0.054	0.401**	0.247
	(0.195)	(0.183)	(0.200)
high · post	-0.203	0.469***	0.189
	(0.207)	(0.157)	(0.213)
treated · high	-0.472*	0.208	-0.071
	(0.284)	(0.236)	(0.292)
high	0.221	-0.317*	-0.010
	(0.186)	(0.173)	(0.218)
post	-0.422***	-0.725***	-0.601***
	(0.144)	(0.146)	(0.158)
treated	0.012	-0.276	-0.144
	(0.470)	(0.469)	(0.473)
Constant	3.517***	3.746***	3.613***
	(0.253)	(0.259)	(0.246)
Observations	1,974	1,974	1,974
R-squared	0.795	0.796	0.795
Industry FE	Yes	Yes	Yes

TABLE 2.11. ESG Index Exclusion and Shareholder Voting

The table reports the estimates of the treatment effect of MSCI ESG index exclusion on proxy voting on a firm's proposals. The dependent variable in Column (1) is the fraction of votes cast in favor of the management for the management-sponsored proposals. The dependent variable in Column (2) is the fraction of votes cast in favor of shareholder-sponsored proposals. The dependent variable in Column (3) is a dummy that equals one if a firm had an environmental proposal submitted at its shareholder meetings in a given year. *treated* is a dummy that equals one if a firm had the MSCI ESG rating of 'B' in May 2018. *post* is a dummy that equals one for 2018 and the following years. The sample covers 965 firms that had the MSCI ESG rating of 'B' or 'BB' as of May 2018. Firms rated 'B' are treated as they were excluded from the MSCI index. Standard errors are clustered at the industry level. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1) Fraction For Management Proposals	(2) Fraction For Shareholder Proposals	(3) Dummy Environmental Proposals
post · treated	-0.000	-0.006	0.001
	(0.003)	(0.039)	(0.011)
post	-0.000	0.010	-0.046***
	(0.002)	(0.026)	(0.009)
treated	0.000	0.008	0.007
	(0.003)	(0.032)	(0.014)
Constant	0.945***	0.361***	0.084***
	(0.002)	(0.016)	(0.007)
Observations	9,919	1,069	10,507
R-squared	0.142	0.350	0.246
Industry FE	Yes	Yes	Yes

TABLE 2.12. Robustness: Alternative Scaling of Releases

The table reports the estimates of the effect of ESG index exclusion on a firm's toxic releases under various scaling approaches. In Columns (1) and (2), the total releases are scaled by total assets and lagged total assets. In Columns (3) and (4), the total releases are scaled by the cost of goods sold and lagged cost of goods sold, respectively. *treated* is a dummy that equals one if a firm had the MSCI ESG rating of 'B' in May 2018. *post* is a dummy that equals one for 2018 and the following years. The sample covers 965 firms that had the MSCI ESG rating of 'B' or 'BB' as of May 2018. Firms rated 'B' are treated as they were excluded from the MSCI index. Standard errors are clustered at the industry level. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
	$\log\left(1+\frac{TotalReleases\_i,t}{Assets_{i,t}}\right)$	$\log\left(1+\frac{TotalReleases\_i,t}{Assets_{i,t-1}}\right)$	$\log\left(1+\frac{TotalReleases\_i,t}{COGSi,t}\right)$	$\log\left(1+\frac{TotalReleases\_i,t}{COGSi,t-1}\right)$
treated · post	0.253**	0.242**	0.235*	0.228*
	(0.118)	(0.120)	(0.125)	(0.127)
post	-0.497***	-0.505***	-0.379***	-0.365***
	(0.090)	(0.096)	(0.094)	(0.099)
treated	-0.171	-0.176	-0.015	-0.001
	(0.421)	(0.427)	(0.447)	(0.448)
Constant	3.610***	3.677***	3.925***	3.958***
	(0.221)	(0.226)	(0.233)	(0.235)
Observations	1,974	1,974	1,974	1,974
R-squared	0.794	0.791	0.789	0.787
Industry FE	Yes	Yes	Yes	Yes

Chapter 3

Fund Family Matters: How Index Funds Improve Corporate Governance

3.1 Introduction

The ownership structure of U.S. firms has experienced a dramatic change in the last few decades. Index funds are becoming increasingly large owners of public firms, and the share of U.S. equity they hold increased from 3.5% in 2000 to 17.2% in 2018 (Bebchuk and Hirst, 2019). Although indexing benefits individual investors by reducing the cost of holding a diversified portfolio, it may also bring a social cost. More specifically, index funds may lack incentives and resources to exercise governance in their portfolio companies, resulting in inefficiently low levels of monitoring from a societal perspective.[1]

Academic research on the corporate governance implications of indexing has led to mixed conclusions. On one side of the debate, Appel, Gormley, and Keim (2016) establish a positive relationship between the fraction of a firm's equity held by index funds and variables typically

[1] Among others, William Ackman writes in his 2016 annual letter to the shareholders of Pershing Square Capital Management that index funds lack both incentives and resources to cast meaningful votes on the hundreds of stocks in their portfolios.

associated with better governance, such as more independent corporate boards, removal of takeover defenses, and more equal voting rights. However, examining the proxy votes cast by mutual funds, Heath et al. (2021) and Brav et al. (2021) find that index funds side with the management more often than active funds do. Such conflicting evidence is puzzling and the role of index funds in corporate governance remains an open question.

In this paper, I propose a novel mechanism that explains the role of index funds in corporate governance. Because of their large size, index funds increase the voting power of fund families they are sponsored by, strengthening the voice mechanism of governance (Shleifer and Vishny, 1986). With control over more votes, fund families can invest in costly monitoring, particularly proxy voting, and alleviate agency conflicts at their portfolio companies.

Fund families centralize proxy voting (Rothberg and Lilien, 2006; Morgan et al., 2011; Choi, Fisch, and Kahan, 2013) because they are better equipped to monitor a firm's management than their individual funds. First, delegating costly monitoring to the fund family avoids the duplication of effort by individual funds and reduces the total cost of monitoring incurred by the funds.[2] Moreover, families are more likely to have superior information processing capacity compared to individual funds. Large fund families have stewardship teams whose primary job is to oversee the engagement with the management and proxy votes cast by the family. Second, a fund family has higher voting power than its funds. Voting all the shares owned by the family as a block increases the probability of being pivotal and, in turn, the expected benefits of voting.

I argue that index funds are beneficial for corporate governance because their growth increases their fund families' voting power, which the family can use to monitor. In particular, index funds increase the number of votes that the family casts, making it more likely that the family is a pivotal voter.[3]

The mechanism I propose generates two empirical hypotheses about index funds' role

[2]In the context of proxy voting, such costs arise because institutional shareholders need to collect and process information for a large number of proposals.

[3]In addition, large index funds increase the fees collected by the family, which can expand the stewardship budget.

in a fund family's proxy voting decisions. First, all else equal, as index funds from a fund family increase their holdings in a firm, active and index funds from the family should coordinate their votes on the firm's proposals more often. The coordination between fund managers is costly because funds may have different objectives. Additionally, centralizing voting can adversely affect the incentives of individual fund managers to acquire proposal-specific information. As index fund holdings increase, the fund family becomes more likely to be a pivotal voter (Levit and Malenko, 2011), increasing the expected benefits of coordinating the votes of its funds.

The second hypothesis is that, all else equal, increasing index holdings increases the fraction of votes that the family casts against the management. Voting against the management is costly because it may compromise business ties (Davis and Kim, 2007; Cvijanović, Dasgupta, and Zachariadis, 2016) and the relationship with the firm's management (Matvos and Ostrovsky, 2010; McCahery, Sautner, and Starks, 2016). In the face of such costs, a fund family votes against the management only when benefits of such voting are sufficiently high, particularly when the family is more likely to affect the outcome of the vote. Because large index holdings can increase the probability of being pivotal, the family is more likely to vote against the management.[4]

Empirically testing these hypotheses is challenging. Both the size of index and active funds holdings in a firm and votes cast on that firm's proposals are driven by unobservable firm characteristics such as the firm's governance quality. Well-governed firms are more likely to have higher valuations, a larger weight in stock indices, and higher holdings by mutual funds. At the same time, they are more likely to receive more votes in favor of the management. OLS estimates are therefore biased and do not identify the causal effect of holdings on voting.

To circumvent these issues, I propose an instrumental variable approach. I use aggregate fund flows to index and active funds in a given family as instruments for the size of the index

[4]It is worth noting that both predictions also apply to an increase in holdings of actively managed funds sponsored by a fund family. In this paper, my main focus is on index funds because their role in corporate governance is a source of controversy, unlike that of active funds. Additionally, I discuss the difference in the role of active and index funds in Section 2.

and active holdings of the family's funds in a given firm. The flows have a substantial effect on holdings, and I argue that they do not affect family-level voting except through their impact on holdings. Family-level flows are both orthogonal to characteristics of portfolio companies and are driven by economic forces that are not related to voting decisions. In particular, fund flows arise in equilibrium either as investors learn of the fund manager's stock-picking skill or because of behavioral biases (Christoffersen, Musto, and Wermers, 2014). On the other hand, voting on a proposal is driven by a firm and proposal-specific characteristics.

In this paper, I show three main results. First, larger index fund holdings facilitate the family-level coordination in voting. As holdings of the family's index funds in a firm increase, funds from the family vote in lockstep more often. The effect is particularly significant in magnitude for a subset of contentious proposals for which the management and the Institutional Shareholder Services (ISS) issue conflicting voting recommendations.

Second, I find that following an increase in index fund holdings, funds within a family vote against the management more often. This finding is consistent with the hypothesis that index funds change the voting strategy of their families by making it pivotal. I find that the effect is stronger for contentious proposals and those sponsored by the shareholders. To further examine whether such voting has further real outcomes, I examine the effect of index funds on votes of their families for close proposals. I find that my results are the strongest for very close proposals.

Finally, I attempt to examine the impact of the growth in indexing on governance over time. Although I demonstrate that larger index funds can improve voting if active holdings are fixed, Corum, Malenko, and Malenko (2021) show that they can be detrimental for governance if they crowd out investments in actively managed funds. To examine this possibility, I estimate a counterfactual frequency with which fund families would vote against the management if all the flows into index funds over the last ten years went into active funds instead. Because not all index fund flows came at the expense of active funds, such counterfactual exercise provides a worst-case scenario bound to the negative impact of indexing. In the exercise, I fail to reject the null that the counterfactual frequency of opposing the management is the same as the one observed in the data. This result is particularly interesting for two reasons. First, it suggests that, because of the fund family structure, index funds and active funds are close substitutes for governance purposes. Second, it contributes to the heated debate on the effect of indexing on governance. Although most current papers focus on the degree on incentives of index funds to engage, the existing studies mostly ignore what the counterfactual outcome is. I show that fund families did not become more passive in voting as a result of a growth in index investing.

To summarize, I find that index funds' holdings positively affect the fund family's monitoring efforts. Index funds enhance family-level coordination in voting and lead to fewer pro-management votes. In a series of robustness checks, I show that these results hold for subsamples of proposals sponsored by shareholders and the management, and across several common proposal types. I find that my results are particularly strong for proposals that change the proxy voting rules at the firm, giving shareholders more voting power in the future.

My findings provide novel evidence on the role of index funds in corporate governance. The existing literature engages in a substantial amount of debate about the extent to which low-fee index funds have incentives to invest in stewardship (Bebchuk and Hirst, 2019). My paper shows that index funds can be beneficial for corporate governance even when they have no direct financial incentives to monitor. When a family offers them together with large actively managed funds, index funds facilitate monitoring on behalf of the active funds.[5]

The paper contributes to three strands of literature. First, I contribute to the growing body of research studying the corporate governance implications of the rise in index investing. Appel, Gormley, and Keim (2016) find a positive relationship between a fraction of a firm's equity held by index funds and positive governance outcomes. In contrast, Heath et al. (2021) and Brav et al. (2021) find that index funds vote with management more often than active funds do. They also show evidence that families that have a larger proportion of index funds are more management-friendly in voting.

[5] The largest investment managers, including Vanguard, State Street Global Advisors, and BlackRock, all offer actively managed funds along with index funds.

Importantly, the findings of this paper are complementary to those in Brav et al. (2021) and Heath et al. (2021). To explain why, in Table 3.2 I report how often index and active funds from the same family vote with the management. Within most fund families, index funds oppose the management as often as active funds. As a result, most of the differences between votes of active and index funds stem from differences in votes of fund families as a whole. To that end, Brav et al. (2021) and Heath et al. (2021) compare votes of predominantly passive and active fund families.[6] They find that fund families that have more index funds are more likely to side with the management. They explain these results by lower incentives of such funds families to monitor the management.[7] In this paper, I take a further step and examine how size of index and active holdings affect the family's voting. Whereas Brav et al. (2021) and Heath et al. (2021) exploit *across-family* variation in the family's index fund share, I utilize additional, *within-family* variation in the size of a family's holdings in firms in its portfolio. I find that when a family's index holdings in a firm increase, the family is more likely to vote against the firm's management. Families with large holdings challenge the management on close proposals that are ultimately more likely to have a significant affect on a firm's value. In sum, my paper helps reconcile evidence in Appel, Gormley, and Keim (2016) and in Heath et al. (2021) and Brav et al. (2021) by explaining why index funds are beneficial for governance even though they do not vote against the management more often.

In a recent paper, Filali Adib (2019) shows that index funds increase shareholder value by changing the outcome of close votes. My paper differs in its focus on fund family-level coordination and voting. Filali Adib (2019) primarily considers the market reaction to passage of close proposals depending on the votes of index funds, while I examine how those votes are determined at the fund family level. Corum, Malenko, and Malenko (2021) analyze the role of index funds in governance theoretically. They argue that the overall impact of

[6] Heath et al. (2021) examine votes as a function of the fraction of a family's assets managed by index funds. Brav et al. (2021) compare votes across families with different fraction of index funds within these families.

[7] Bolton et al. (2020) and Bubb and Catan (2018) provide an alternative explanation: the heterogeneity in fund family's preferences that is not captured by their size or other observable characteristics.

growth in indexing on governance depends on whether index funds primarily crowd out active fund holdings or private savings. In the latter case, the overall governance improves even when index funds charge fees that are lower than fees of active funds.

The paper also contributes to the literature on proxy voting. A large body of research investigates the incentives of individual funds to cast votes (Davis and Kim, 2007; Matvos and Ostrovsky, 2008; Iliev and Lowry, 2015), as well as the consequences of those votes for the firm-level outcomes (Cai, Garner, and Walkling, 2009; Fos, Li, and Tsoutsoura, 2018; Cuñat, Gine, and Guadalupe, 2012; Ertimur, Ferri, and Oesch, 2013 and 2018).[8] The central insight from my paper is that incentives to vote and, in particular, to oppose the management are shaped at the fund family level. Therefore, it is crucial to consider the fund family structure when studying the effects of mutual funds' proxy votes on firm-level policies and outcomes.

Finally, the paper relates to the literature on the organizational structure of mutual fund families. Nanda, Wang, and Zheng (2004) find that a well-performing fund attracts flows to other funds in the same family. Fund families allocate resources to produce star funds (Gaspar, Massa, and Matos, 2006) and strategically time a fund's initial offering to investors (Evans, 2010). In a more general context, Stein (1997) studies a model of a firm's headquarters that picks the winners and reallocates the resources among the firm's divisions. In my paper, I find that fund families reallocate the voting rights across their funds and vote for the benefits of some of their funds. To the best of my knowledge, this is the first paper testing this mechanism in such a context.

The paper proceeds as follows. I develop the main empirical hypotheses in Section 2. I describe the data and provide some descriptive evidence in Section 3. In Section 4, I discuss my empirical strategy. I report the main results in Section 5 and robustness exercises and extensions in Section 6. I conclude in Section 7.

[8] See Yermack (2010) for a detailed survey on the proxy voting literature.

3.2 Hypothesis Development

Consider a fund family that consists of an index fund and an actively managed fund. Funds vote on a proposal that increases the firm's value if passed. Both the index fund and the active funds have incentives to increase the firm's value, although the extent of the incentives of the latter may be larger (Bebchuk and Hirst, 2019). In this framework, I study the effect the index fund has on the fund family's voting decisions.

First, consider how the size of the index fund holdings affects coordination in voting between the family's funds. Voting in lockstep is beneficial because it increases the probability of being pivotal and changes the likelihood that the proposal passes. Such coordination, however, can also be costly for two reasons. First, in some cases, individual funds can have different objectives, which requires them to vote differently. Second, centralizing voting decisions at the family level can hurt fund managers' incentives to engage in costly information acquisition, leading to less informed voting.

In the context of this cost-benefit analysis, consider the effect of an increase in index fund holdings while keeping the active fund's holdings fixed. As index holdings increase, the family is more likely to be pivotal, increasing the benefit of coordinated voting. If such an increase outweighs the costs of coordination, funds will vote in lockstep. This reasoning leads to the following hypothesis:

Hypothesis 1. *All else equal, an increase in the size of index fund holdings in a firm leads to more coordination among funds from the same family voting on the firm's proposal.*

If the active and index fund coordinate, the next question is how they use their combined voting power. Mutual funds decide whether to vote against the management based on the associated costs and benefits (Iliev and Lowry, 2015). On the benefits side, a fund manager may want to vote against the management when it increases the firm value, as he receives a fraction of the increase.

Voting against the management is costly for two reasons. First, fund managers need to engage in costly information acquisition to determine which proposals to oppose the management on. Voting with the management constitutes a cheaper default option. Second, opposing the management can result in the loss of access (McCahery, Sautner, and Starks, 2016) and business opportunities, such as managing a firm's retirement accounts (Davis and Kim, 2007; Cvijanović, Dasgupta, and Zachariadis, 2016).

Index funds alter this cost-benefit analysis of their fund families. When a family has an index fund with larger holdings in a firm, the family is more likely to be pivotal on the proposal, increasing the expected benefits of voting against the management. Intuitively, as the family becomes pivotal, it will oppose proposals that it did not before since it could not affect their passage.

Hypothesis 2. *All else equal, if index fund holdings in a firm increase, funds from the same family, including active funds, are more likely to vote against the management on the firm's proposals.*

Finally, it is important to discuss the different roles that active and index funds play in the fund families. Although both index and active funds have the same effect on the family's voting power, there are important differences between the two. In particular, index and active funds differ in their incentives to vote against the management and in how costly it is to coordinate their votes. Such heterogeneity generates different roles for active and index funds in their coordination and voting against management. I discuss these differences in more detail in the Results section.

3.3 Data

A. Proxy Voting Data

Since 2003, the Securities and Exchange Commission (SEC) requires mutual funds to disclose their proxy voting records using the N-PX forms. I use these records compiled by the ISS in the Voting Analytics database. The data contains the votes cast by equity mutual funds from the largest fund families during shareholder meetings at Russell 3000 firms. Voting Analytics also provides proposal-specific information, voting recommendations of the ISS, and vote outcomes. I focus on votes cast on proposals at the U.S. firms between 2007 and 2016. I drop proposals that ask the shareholders to determine the frequency of advisory votes on the executive's compensation.

B. Mutual Fund Holdings

I use S12 mutual fund holdings data compiled by the Center for Research in Security Prices (CRSP) in its CRSP Mutual Fund Database. Since 2004, the SEC requires mutual funds to disclose their holdings quarterly using N-Q and N-CSR forms. Unfortunately, ISS and CRSP databases do not share common mutual fund identifiers. To merge the ISS and CRSP datasets, I obtain the identifier (SeriesID) assigned to each mutual fund by Electronic Data Gathering, Analysis, and Retrieval (EDGAR). Using the N-PX form number in the Voting Analytics database, I obtain a set of potential SeriesIDs for each fund in the Voting Analytics and perform matching by a fund name. Finally, I retrieve fund tickers from EDGAR for each SeriesID and use it to match voting records to holdings from CRSP.[9] Following this matching procedure, I obtain a sample of 5,948 mutual funds from 409 mutual fund families.

I match firms covered by the Voting Analytics database with CRSP by ticker-year and then cross-check resulting matches by firm names. Additionally, I perform a manual search to fill the missing observations. Finally, I match each fund's voting record on a firm's proposal to the fund's last reported holdings in that firm before the voting record day. Mutual funds report holdings quarterly, which causes a gap between the report date and the voting record date. I require that gap to be no longer than 60 days.

For each mutual fund covered by CRSP, I compute its monthly fund flows before the record date. In particular, if the holdings for fund k are obtained at month t, I compute the monthly flow as $f_{k,t} = (TNA_{k,t} - TNA_{k,t-1}(1 + r_{k,t}))/TNA_{k,t-1}$, where $r_{k,t}$ is the fund's return between $t-1$ and t, TNA is the fund's total net assets.

I classify each fund as either actively managed or an index fund. A fund is considered to be an index fund if CRSP classifies it as such (Flag 'D') or the fund's name contains one of the words common for index funds.[10] I drop funds that CRSP classifies as levered

[9] If a fund offers several classes of shares, I keep only one match because CRSP links those classes to the same portfolio.

[10] This approach follows Iliev and Lowry (2015) and Appel, Gormley, and Keim (2016, 2019). I classify the fund as an index fund if its name contains one of the following strings: 'Index', 'ETF', 'Russell', 'S&P',

exchange-traded funds (ETFs). Following these steps results in 1,098 index funds and 4,644 actively managed funds.

C. Family-Level Data

In the empirical analysis, I study family-level voting decisions.[11] I focus on two measures: the degree of coordination between funds from the same family, and the degree of support for the management. I define a dummy variable *Disagreement* to be equal to zero if all the funds within a family vote on a given proposal identically. I use the fraction of funds that vote for the management on a given proposal to capture the degree of support for the management.[12] Similarly, I calculate the measure for active and index funds separately.

I aggregate fund flows and holdings of individual funds at the family level. I sum the holdings of index funds and holdings of active funds from the same family in a firm at a given date. Because my analysis focuses on the coordination between active and index funds in the same family, I discard the observations for which either active or index holdings are zero. When aggregating flows, I compute the total dollar flow for index (active) funds and normalize it by the previous quarter's total size of index (active) funds from a given fund family.[13] Finally, I censor the family-level flows and holdings at the 1st and 99th percentile. Additionally, I remove families with less than 100 observations and firms with fewer than 50 observations.

D. Summary Statistics

Table 1 reports the summary statistics. The sample consists of votes cast by funds from

'SP', 'SNP', '100', '500', '1000', '3000'. I manually cross-check the results and manually check with the prospectuses when in doubt.

[11] Some studies, in particular Bubb and Catan (2020), aggregate voting at the investment advisor, not the fund family level. I focus on fund families for two reasons. First, it is empirically challenging to identify whether a sub-advisor exercises the voting authority. For example, while Wellington Management served as a sub-advisor for some of the Vanguard's funds, they did not vote any of the Vanguard's proxies until 2019, when Vanguard decided to delegate some of the voting rights to Wellington Management. Second, focusing on fund families will generally understate the effect of index funds on coordination.

[12] The fund sides for the management if it votes according to the recommendation issued by the management. This includes voting in favor of the proposals sponsored by the management and supporting shareholder-sponsored proposals when the management recommends doing so.

[13] In some cases, the latest holdings report for different funds in the same family do not correspond to the same date. For example, if the record date is February 15th, one fund may report its latest holdings on February 1st, while the other one reports on January 5th. I discard such observations.

61 fund families on 126,270 proposals for 1,935 different firms. Panel A provides summary statistics for individual funds. An average fund manages \$2.2 billion in assets and has an average holding of \$14.29 in a firm. Its average monthly flow before the vote's record day is 2%, with the median at zero. The fund often votes in favor of the management, 92% of the time.

Panel B provides summary statistics for the sample aggregated at the family level. On average, 2.7 index and 2.9 active funds from the same fund family vote on a proposal. The median numbers are slightly smaller. On average, a family's holdings in a firm are \$55 million through its index funds and \$73 million through its active funds. Importantly, holdings of the index and active funds at the family level are comparable.

On average, 92.5% of a family's funds vote with the management. Interestingly, the numbers are very similar for active funds (92%) and index funds (92.7%). Funds within the same family rarely disagree in their votes, only in 4.1% of cases.

E. Descriptive Evidence

In this section, I provide descriptive evidence on voting coordination within fund families. I start with Table 3.2, which documents the voting patterns for some of the largest mutual fund families. Fund families differ in the extent to which they coordinate the voting of their funds. Funds from some families, such as Vanguard, Rydex, and TIAA-CREF, almost always fully coordinate their votes. Funds from others, for example, State Street Global Advisors and Fidelity, often disagree with each other.[14]

In the last two Columns of Table 3.2, I report the average fractions of active and index funds voting for the management. Two interesting patterns emerge. First, as a result of coordination at the family level, in most cases, active funds vote with the management as often as index funds do. Second, there is a considerable variation across fund families in how often they vote for the management. For instance, while the fraction of funds from Rydex Investments voting for the management is almost 99.5%, for Fidelity, it is 88.9%.

[14] In the subsample used in Table 3.2, State Street Global Advisors is an outlier in its frequency of disagreement. This is driven by a single actively managed fund, SSGA IaM Shares Fund, that often deviated from the rest of the family and voted against the management 46% of the time in the earlier time period.

One possible interpretation of patterns in Table 3.2 is that a high degree of coordination is a result of actions taken by fund families. The alternative is that infrequent disagreement at the family arises mechanically because most proposals are procedural, and there is little disagreement among shareholders. To examine this alternative in more detail, I estimate what the level of coordination would have been if funds were randomly reassigned to different families. I randomly reassign fund family membership across funds while keeping the number of funds in each family the same. I then compute the disagreement measure for this bootstrapped sample and repeat the same procedure 1,000 times. I plot the resulting average disagreement in the last bar for each group.[15]

I show the results in Figure 3.1. I compute the average of the disagreement dummy for four groups of proposals: uncontested director elections, votes on executive pay, votes related to board structure, and votes related to the proxy voting rules at a firm. For each group, I report the average *Disagreement* among a family's index funds, active funds, and all funds. In the last bar, I show the bootstrapped benchmark described above.

Comparing the first two bars within each group, index funds in the same fund family vote in lockstep more often than active funds do. For example, when voting on proposals related to board composition, index funds in the same family only disagree 2.9% of the time, while active funds do so 9.3% of the time. For proposals that aim to amend proxy voting procedures at a firm and redistribute voting rights from the management to the shareholders, the numbers are 5.9% and 7.2%, respectively. Index funds vote as a block more frequently than active funds.

Bootstrapped benchmark disagreement frequency is much higher than the one observed in the data. In particular, on the proposals related to the allocation of voting rights, funds in the same family disagree 10% of the time, while if the funds were assigned to families randomly, the number would be 33.8%. Such finding confirms that fund families induce a significant degree of coordination in voting among its funds. The disagreement between fund families is more significant than disagreements between funds in the same family.

[15] I do not report the bootstrapped confidence intervals on the Figure because they are very tight and do not change the discussion above.

In the paper, I build on this observation and further explore whether the size of index fund holdings causes such high levels of coordination. Furthermore, I demonstrate how, given the coordination among funds, fund families use the blocks of shares held by its index funds to challenge the management on important proposals and affect the likelihood that a proposal passes or fails.

3.4 Empirical Framework

A. Empirical Strategy

To provide an empirical test for Hypotheses 1 and 2, I estimate the effect of a fund family's index and active fund holdings in a firm on its voting decisions. I compare how two fund families with identical active fund holdings, but different levels of index fund holdings vote on a firm's proposal. Families whose active funds have significant holdings in a firm are likely to exert more effort to make an informed vote on a proposal because it benefits those active funds. For this reason, it is important to take active holdings into account in this comparison. Importantly, I measure the holdings in dollar terms because it captures the size of gains that families receive from improving the firm's value (Edmans and Holderness, 2017). My baseline specification is

$$\begin{aligned} y_{ip} &= \beta_I \log \text{Index Holdings}_{ip} + \beta_A \log \text{Active Holdings}_{ip} + \\ &+ \gamma \log \text{MCap}_{j(p)t(p)} + \delta X_{ip} + \mu_{f(j)} + \lambda_{t(j)} + \epsilon_{ip}, \end{aligned} \tag{3.1}$$

where i is a fund family, p is a proposal that shareholders of firm $j(p)$ vote on at time $t(p)$. y_{ip} is the fund family's voting decision of interest. Index Holdings_{ip} (Active Holdings_{ip}) are the total dollar value of shares held by index (active) funds from fund family i at firm $j(p)$ before the meeting with proposal p. $\text{MCap}_{j(p)t(p)}$ is the market capitalization of firm $j(p)$ at time $t(p)$. X is the vector of additional control variables, such as the ISS recommendation and total size of fund family. Note that this specification is flexible and includes other common functional forms used in the literature.[16]

[16]For example, the specification in which I measure holdings as the fraction of a firm's outstanding equity

Identifying the causal effect of a fund family's holdings on voting decisions is empirically challenging. Fund holdings are not randomly assigned and are likely correlated with firm-level characteristics that also affect votes cast by fund families. In particular, well-governed firms, all else equal, perform better (Gompers, Ishii, and Metrick, 2003; Bebchuk, Cohen, and Ferrell, 2009) and receive higher weight in widely-used value-weighted benchmark indexes. Similarly, the governance structure of the firm may be an important factor for the active fund manager when selecting the firm's weight in the fund's portfolio. At the same time, the quality of the firm's governance has an effect on the votes cast by its shareholders. As a result, the OLS estimates are biased[17] and would suggest that fund families are more management-friendly than they actually are.[18]

To circumvent these issues and estimate the causal effect of index fund holdings on family-level voting decisions, I propose an instrumental variable approach. I use family-level flows of its index and active funds as instruments for the holdings of the family's funds in a given firm. I start by describing the construction of the instruments and then discuss their validity.

For each fund in the sample, I compute its flows in the last month prior to the vote's record date. I aggregate the flows across index and active funds from the same fund family and compute the percentage flows for the family's index and active funds using the following expression:

$$
\begin{aligned}
IndexFlow_{i,t} &= \frac{\sum_{k \in IndexFunds(i)} Flow_{k,t} TNA_{k,t-1}}{\sum_{k \in IndexFunds(f)} TNA_{k,t-1}} \\
ActiveFlow_{i,t} &= \frac{\sum_{k \in ActiveFunds(i)} Flow_{k,t} TNA_{k,t-1}}{\sum_{k \in ActiveFunds(i)} TNA_{k,t-1}},
\end{aligned}
\tag{3.2}
$$

where $IndexFunds(i)$ and $ActiveFunds(i)$ are the sets of index and active funds offered by a fund family i.

I use the flows defined in (3.2) as instruments for $\log IndexHoldings_{i,p}$ and $\log ActiveHoldings_{i,p}$.

can be obtained by setting $\gamma = -\beta_I - \beta_A$.

[17]Other unobservable firm-level variables can also be correlated with fund's votes and holdings. In particular, the firm's stock liquidity can affect both holdings of index funds and voting decisions (Edmans, Fang, and Zur, 2013).

[18]Governance is positively correlated with both fund holdings and the fraction of votes that go in favor of the management.

In my estimation, I use the GMM-IV estimator with the optimal weighting matrix. Standard errors are clustered at the firm levels to account for correlation in votes cast for proposals of the same firm.

B. Instrument Relevance

To be valid instruments, flows must have a strong correlation with the fund family's holdings. Holdings of fund i in firm j at time t can be expressed as

$$\begin{aligned} \log \text{Holdings}_{ijt} &= \log TNA_{it} + \log w_{ijt} \\ &= \log TNA_{i,t-1} + \log(1 + r_{i,t} + f_{i,t}) + \log w_{ijt} \\ &\approx \log TNA_{i,t-1} + r_{i,t} + f_{i,t} + \log w_{ijt}, \end{aligned} \tag{3.3}$$

where the TNA_{it} is the fund's size at time t and w_{ijt} is the weight of firm j in the fund's portfolio. $r_{i,t}$ and $f_{i,t}$ are, respectively, the fund's return and flow from time $t-1$ to time t, and the approximation assumes $r_{i,t} + f_{i,t}$ is small. The endogeneity described above arises because the portfolio weight w_{ijt} is not randomly assigned to funds and can depend on a firm's governance.

Equation (3.3) predicts a positive relation between holdings and past flows.[19] To test the strength of the instruments, I estimate the following first-stage regressions for $h_{ip} = \log ActiveHoldings_{ip}$ and $h_{ip} = \log IndexHoldings_{ip}$:

$$\begin{aligned} h_{ip} = \beta_I^{FS} \log \text{Index Flow}_{i,t(p)} + \beta_A^{FS} \log \text{Active Flow}_{i,t(p)} + \gamma^{FS} \log \text{MCap}_{j(p)t(p)} + \\ + \delta^{FS} X_{ip} + \tilde{\mu}_{f(j)} + \tilde{\lambda}_{t(j)} + \tilde{\epsilon}_{ip}. \end{aligned} \tag{3.4}$$

I report the estimates in Table 3.3. Estimates in Column (1) show that a 5% inflow to a family's index (active) funds increases index holdings by 4% (1.8%) of its standard deviation.[20] For holdings of active funds, a 1% inflows to the family's active funds increase active holdings by 4.4%. Kleibergen and Paap (2006) robust F-statistic for this specification is 18.39, well above the rule-of-thumb value of 10, indicating that the instruments are

[19] While Equation (3.3) is derived for a single fund, the analysis for a group of funds is similar.

[20] The log holdings are standardized to have a standard deviation of one. Computing the magnitude then reduced to $\Delta \log \text{IndexHoldings} = \hat{\beta} \Delta IndexFlow \cdot std(\log IndexHoldings)$.

strong. In Columns (3) and (4), I include additional controls and obtain qualitatively similar estimates.

In the subsequent results, I report both the robust first-stage F statistic and the Anderson-Rubin (AR) test statistic, which allows to conduct inference that is robust to weak instruments (Andrews, Stock, and Sun, 2019), which can be an issue when analyzing subsamples of the data.

C. Exclusion Restriction

Valid instruments should not only be relevant but must also satisfy the exclusion restriction. In my setting, it states that flows do not affect fund family voting decisions except through their impact on the size of holdings. The exclusion restriction is likely satisfied for several reasons.

First, the variation generated by family flows is at the family, not the proposal, level. Large fund families hold the majority of publicly listed firms. Families from my final sample, on average, vote on proposals from 1,376 firms in a given year. For such diversified portfolios, fund flows are most likely orthogonal to firm-level characteristics that drive voting decisions.

Second, the economic forces that determine fund flows do not directly affect voting decisions of a fund family. The vast literature on mutual fund flows provides two groups of explanations for the existence of fund flows. One strand of the literature argues that flows are a result of learning about a fund manager's skill on behalf of investors (Berk and Green, 2004). Fund flows arise as investors observe the performance of a fund and update their posterior belief of the manager's skill.[21] Alternatively, several studies argue that mutual fund investors can behave irrationally and make mistakes.[22] It is unlikely that either irrational behavior or learning has a direct effect on the votes cast by active fund managers or their fund families.[23] Neither of the two explanations tie the flows to voting at a particular firm

[21] Such explanation can apply to actively-managed funds because their performance deviates from the benchmark. For index funds, Clifford, Fulkerson, and Jordan (2014) document the sensitivity of ETF-flows to past performance. They conclude that rational explanation for this finding is unlikely.

[22] Choi, Laibson, and Madrian (2010) and Barber, Odean, and Zheng (2005) document that investors make suboptimal choices when selecting between mutual funds.

[23] The assumption I make is that fund's gains from voting are very small compared to their gains from stock picking. Such an assumption has some support based on the estimates of the value gains from proposal passage in Cuñat, Gine, and Guadalupe (2012).

because votes are likely driven by firm and proposal-specific characteristics.

The exclusion restriction might still be violated if fund managers try to attract flows by changing their voting behavior. For example, following outflows from the fund, the manager may decide to vote more actively to attract flow in the future. I argue that this is unlikely for two reasons. First, flows used to construct the instruments are typically two to four months before the votes are cast. It is unlikely that flows several months before the vote would change the voting behavior of fund families. Second, I argue that voting is unlikely to attract significant inflows. Estimates of a fund family's flow-performance sensitivity and value creating from voting in the literature suggest a very small impact of voting on flows.[24]

One drawback of my empirical strategy is that it does not allow to include fund family fixed effects. Because of the dynamic relationship between fund flows and holdings, including unit fixed effects results in a finite sample bias discussed in Stambaugh (1999).

3.5 Results

In this section, I present the main results. I start by examining how the size of index fund holdings affects the decision of funds within a family to coordinate their votes. Next, I consider the effect index funds have on the decision of their families to support the management. Finally, I evaluate the overall effect of a change in sizes of the index and active holdings of fund families on the votes they cast.

3.5 Family-level Coordination

I start by testing Hypothesis 1, which states that index funds increase voting coordination of their fund families. As holdings of a fund family's index funds in a firm increase, it increases the collective voting power of all funds from that fund family. Having more voting power can make those funds coordinate more frequently because they are more likely to be pivotal.

[24]To estimate the effect of voting on flows, assume the fund family is pivotal, and the proposal increases shareholder value by 2.8% if passed (Cuñat, Gine, and Guadalupe,2012). Following Lewellen and Lewellen (2021), assume a stock's weight of 1.56% in the fund family's portfolio. Additionally, assume that a 1% increase in quarterly returns leads to a 0.68 increase in the flows during the subsequent year. Voting for the proposal leads to a $2.8\% \cdot 1.56\% \cdot 0.68 = 3\text{bp}$ increase in the fund family's flows over the next year.

The outcome variable is the dummy $Disagreement_{ip}$, which equals one if the funds from fund family i do not vote identically on proposal p. I report the results in Table 3.4. The OLS coefficient on index holdings in Column (1) is positive but small in magnitude; a one standard deviation increase in the log of index holdings leads to a decrease in the coordination by 8% of its standard deviation. Interestingly, the IV estimate has the opposite sign and is larger in magnitude. The difference between the OLS and IV estimates is consistent with the sign of the omitted-variable bias discussed in the previous section. A one standard deviation increase in log index holdings reduces disagreement within a family by 0.58 of its standard deviations. If a fund family increases its index fund holdings by 20%, its funds will fully coordinate their voting in 0.95% more cases. Although seemingly small, this is a large effect as most proposals receive unanimous support from any given family. Funds within a family disagree only 4.1% of the time, meaning that a 20% increase in index fund holdings in a firm boosts the coordination by $0.95\%/4.1\% = 23.2\%$ relative to its sample mean.

The estimate $\hat{\beta}_A$ is also negative and statistically significant. Larger active fund holdings lead to more coordination within the family. Note, however, that the size of the estimate is much smaller than the estimate for index fund holdings. Having multiple active fund managers within a family likely makes coordination more difficult because they can have different objectives and information sets.

In Column (3), I include a control variable *ISSforMgt*, which is a dummy variable that equals one if the ISS issues a recommendation identical to the one released by the manager. Votes cast by mutual funds are affected by proxy advisors' recommendations (Malenko and Shen, 2016), affecting the family-level coordination. Additionally, I include a set of firm fixed effects to absorb a time-invariant level of coordination specific to proposals in a given firm. The estimates on the index and active holdings remain negative, statistically significant, and become slightly larger in magnitude. A 20% increase in the size of index fund holdings leads to 0.46% fewer disagreements among funds in the same family.

In Column (4), I control for the total size of a fund family. Larger fund families can collect more fees and allocate more funds to stewardship teams, leading to a higher quality of their votes. On the other hand, larger families may be more prone to conflicts of interest (Davis and Kim, 2007; Cvijanović, Dasgupta, and Zachariadis, 2016), which leads to more pro-management votes. The estimates remain negative and quantitatively similar.

Finally, I explore how the role of index funds differs depending on the importance of a proposal. I define a proposal as contentious if the ISS and the management issue conflicting recommendations. Because a large number of proposals in the sample are highly procedural, funds rarely oppose such proposals, and, as a result, families vote unanimously. As described in the hypothesis section, the coordination effect should be more pronounced when the benefits of voting are higher, which is more likely for contentious proposals.

I report the results for a sample of contentious proposals in Column (5) and non-contentious ones in Column (6). Column (5) estimate is almost four times larger than a full-sample estimate in Column (4). Following a 20% increase in the size of a family's index holdings, funds within a family disagree 2.5% less frequently. The disagreement occurs for 19% of contentious items, which results in a 5.6% relative increase in the coordination frequency. Similar to the previous results, the estimate for the effect of active fund holdings is smaller than the one for index fund holdings. The estimates for the effect of the index and active holdings in Column (6) are slightly smaller in magnitude but still negative and statistically significant.

To summarize, I find that index funds facilitate voting coordination within their fund families, as larger index holdings lead to more coordination within the family. In the Robustness section, I demonstrate that these results hold for alternative measure of family-level coordination in voting.

3.5 Voting for the Management

In the previous section, I show that index funds facilitate family-level coordination in voting. Whether such coordination benefits the shareholders is unclear. One possibility is that fund families use their increased voting power to monitor the management and vote against proposals that destroy shareholder value. Alternatively, fund families can have incentives to vote proxies of all its funds in favor of the management (Davis and Kim, 2007).

In this subsection, I study the effect of index fund holdings on a family's decision to vote with the management. Note that, if a family coordinates votes of its funds, votes of its active funds will depend not only on their own size but also on the size of index funds in the same family. Building on this observation, I use the fraction of active funds within a family that side with the management, *FractionActiveFundsForMgt*, as the main outcome variable.

I report the estimates in Table 3.5. The IV estimate on index holdings in Column (2) is negative; a one standard deviation increase in log index holdings decreases the fraction of active funds in a family that support the management by 0.17 of its standard deviation. Following a 20% increase in index holdings, 0.6% fewer active funds from the family vote with the management. The effect is marginally statistically significant but small in magnitude. The difference between the IV estimate and OLS estimate in Column (1) is consistent with the sign of omitted variable bias discussed above. The estimate on active fund holdings is negative, statistically significant, and close in magnitude to the estimated coefficient on index fund holdings.

In Column (3), I include a control for the ISS recommendation and firm fixed-effects. Further, in Column (4), I control for the total size of the fund family. The coefficients remain qualitatively similar and statistically significant. As most proposals receive significant support from shareholders, it is not surprising that the effect of index fund holdings on voting is small for a large sample of proposals.

Next, I repeat the analysis for samples of contentious and non-contentious proposals. I expect to find larger effects for contentious proposals because they likely have a larger effect on firm value, and institutional investors are more likely to be pivotal on such proposals. In Column (5), the estimate on log index holdings is 2.97. A 20% increase in index holdings decreases the fraction of active funds voting for the management by 4.8%. For non-contentious proposals in Column (6), the coefficient estimate on log index holdings is small in magnitude and not statistically significant.

The estimate on active holdings in Column (5) is also negative and statistically significant.

Interestingly, it is larger in magnitude than the estimate on index holdings. Although equal active and index holdings have identical voting power, the incentives to vote against the management are very different. Active fund managers are usually compensated based on their performance, while index fund managers receive a fee equal to the small fraction of the assets they manage. As a result, active fund holdings have a larger effect on voting than index holdings of the similar size (Bebchuk and Hirst, 2019). In the next section, I provide additional analysis that highlights the difference between the governance role of active and index funds.

In sum, I find that index funds allow active funds in the same family to vote against the management more often. The effect is particularly strong for a subset of contentious proposals.

3.5 Index Holdings and Close Proposals

In the results so far, I show that funds within a family are more likely to coordinate and oppose the management when they have larger index fund holdings in a firm. Do these actions affect the probability that a proposal passes or fails? To study this question, I compare the effect of index funds on the family's voting for subsets of closely contested proposals. I compute the margin of a vote as the absolute value of the difference between the fraction of votes in favor and the threshold required for the proposal to pass. I estimate my main specification for subsamples of proposals with the margins under 10%, between 10% and 50%, and above 50%.

I report the estimates in Table 3.6. The effect of index fund holdings on voting is the strongest in Column (1) for proposals with the smallest vote margin. A one standard deviation increase in log index holdings leads to almost a two standard deviation decrease in the fraction of a family's funds voting with the management. The coefficient estimate on index holdings in Column (2) is also negative and statistically significant but more than three times smaller in magnitude. For proposals that pass or fail by a wide margin, above 50%, index fund holdings do not affect the family's voting, according to the estimate in

Column (3).

The magnitude of coefficients on log active holdings across Columns (1) – (3) shows a very similar pattern. The size of holdings of a family's active funds have the strongest effect on voting on close proposals, and the effect becomes weaker for proposals that are less close. There are, however, two interesting differences as compared to index funds. First, the estimate for active holdings is larger than for index holdings for close proposals in Column (1). Second, the size of active fund holdings has a statistically significant effect even for proposals that are far from being close. These two patterns can emerge if fund families primarily vote in the interests of its active funds, but also take into account the size of their index holdings when a vote is close.

3.5 Did the Growth in Indexing Improve Governance?

The results so far show that index funds help fund families oppose the management on proposals, particularly the close ones. The big question unanswered so far is whether the growth in index investing has been beneficial for governance over time. As pointed out by Corum, Malenko, and Malenko (2021), there are at least two possibilities. On the one hand, index funds can grow over time because investors reallocate funds away from direct ownership or private savings. In that case, the size of active funds is unaffected, and families can use the increased size of their index funds and vote against value-destroying proposals, hence improving governance. On the other hand, growth in index funds can come at the expense of active funds. For example, index funds can attract flows that would otherwise go into actively managed funds. In the latter case, the growth in index investing can be harmful to governance because it adversely affects the size of active funds in a family and, in turn, their incentive to monitor.

Answering the question of whether the growth in indexing has been beneficial for governance is empirically challenging. The main difficulty is that a researcher does not observe the level of active holdings that would arise in equilibrium if index funds were constrained in their ability to attract new flows from investors. In this section, I present an attempt to estimate the effect of the growth in indexing over time on family-level monitoring efforts, measured by voting.

Using the estimates of the sensitivities of voting to holdings, $\hat{\beta}_I$ and $\hat{\beta}_A$, I can perform a statistical test to answer the following question: how would a fund family's propensity to vote against the management change following a Δ_{Index} (Δ_{Active}) percentage increase in index funds (active funds) holdings?

What values of Δ_{Index} and Δ_{Active} should one use to estimate the effect of the growth in indexing on governance? Ideally, one would need to observe Δ_{Active} for the counterfactual world in which growth in indexing was limited. Because such data is not available, I take an alternative approach. I estimate how often fund families would oppose the management if all the flows in index funds over my sample period went into actively managed funds instead. If growth in index funds indeed adversely affected the size of active funds, such exercise would put an upper bound on the absolute size of the negative effect of indexing on governance.

In my dataset, I find that, between 2007 and 2016, the average index fund holding increased by 58.91%, while the average active fund holding increased by 16.87%.[25] The average index holdings at the family increased by \$12.4 million over the time period. In the exercise, I decrease index holdings and increase active holdings by this amount, which corresponds to $\Delta_{Index} = -37.07\%$ and $\Delta_{Active} = 27.81\%$.[26] Then I test the following hypothesis

$$\mathcal{H}_0 : \beta_I \log(1 + \Delta_{Index}) + \beta_A \log(1 + \Delta_{Active}) = 0 \tag{3.5}$$

against a two-sided alternative $\mathcal{H}_1$.

I report the results in Table 3.7. In Column (1), I estimate the effect of capital reallocation on the propensity to vote for the management. I find that the fraction of funds voting with the management decreases by only 0.4%; the estimate is not statistically significant at the

[25] These numbers are based on a sample of families that have positive active and index holdings in a firm. In the unrestricted sample, average index fund holdings increased by 53.35%, while average active fund holdings increased by 12.84%.

[26] The average total holdings of a fund family's index and active funds in 2016, the last year of my sample, are \$33.53 million and \$44.71 million, respectively. Reallocating \$12.4 million from index funds to active funds corresponds to a $12.43/55.1 = 27.8\%$ increase in active holdings and $-12.43/33.53 = -37.07\%$ reduction in index holdings.

10% level. Next, I perform the test on subsamples of close proposals. For the subsample of proposals with a vote margin below 10%, the estimate in Column (2) is larger, a 6.7% reduction in the fraction of funds voting with the management, although still not statistically significant. In Columns (3) and (4), I consider proposals that have a vote margin between 10% and 50%, and above 50%, respectively. The only statistically significant estimate is in Column (4), for proposals that pass or fail by a very wide margin. In a counterfactual for these proposals 4.2% fewer funds from a family vote for management.

To summarize, I show that it is unlikely that the growth in index investing over time hurt governance on the margin. Index and active funds within a family act as close substitutes in the context of a fund family voting decisions. Several caveats are in order. First, the analysis in this section puts an upper bar on the negative impact of indexing because it is unlikely that all flows into index funds would go to active funds otherwise. Second, I use my main estimates $\hat{\beta}_A$ and $\hat{\beta}_I$, which capture marginal effects of holdings on voting. With a large reallocation, the effect of holdings on voting can be different. Finally, I assume that in the counterfactual outcome, active holdings increase by the same amount as index holdings did over the last ten years.

3.6 Robustness and Extensions

3.6 Proposal Sponsor

In my main analysis, I do not differentiate between proposals sponsored by the shareholders and the management. Although shareholders sponsor only 3.9% of proposals in my sample, they are important for my analysis. First, shareholder-sponsored proposals receive significantly lower support than the ones sponsored by the management. In my sample, the average percentage of votes in favor of shareholder-sponsored proposals is 34%, while for the proposals sponsored by the management, it is 93%. Hence, institutional investors are more likely to be pivotal while voting on such proposals. Second, a firm's management can negotiate which management-sponsored proposals are put to the vote with its institutional investors. In this case, the votes of large shareholders may not reflect the monitoring activities of fund families. Studying shareholder proposals alleviates this concern because they are less likely to be pre-negotiated (Gillan and Starks, 2000).

I estimate the main specification for subsamples of proposals sponsored by management and by the shareholders and report the results in Table 3.8. In Columns (1)–(3), I report the estimates for management-sponsored proposals. A 20% increase in the size of index holdings leads to 0.4% fewer active funds voting with the management. The estimated effect for contentious proposals is much larger, at 4.4%. Interestingly, the estimates for active holdings are larger than for index holdings.

In Columns (4)–(6), I focus on shareholder-sponsored proposals. There are several interesting findings. First, larger index fund holdings lead to lower support for the management by a family's active funds. A 20% increase in the size of index fund holdings leads to 7% fewer active funds voting with management across all shareholder-sponsored proposals and 9.1% for contentious ones. Second, the estimates $\hat{\beta}_I$ are larger for shareholder-sponsored proposals in Columns (4)–(6) than for management-sponsored ones in Columns (1)–(3). As shareholder proposals likely have a larger impact on firm value, coordination and voting against management are more important for this subset of proposals.

It is important to note that my instruments become weak for a subsample that only includes shareholder-sponsored proposals. For each specification, I report the Anderson-Rubin test-statistic and the corresponding p-value. This test allows testing a hypothesis of the joint significance of coefficients on the endogenous variables (index and active holdings, in my setting). Such a test is robust to weak instruments, which allows me to conduct inference in my setting. Across all specifications, I reject the null that coefficients on active and index holdings are jointly zero.

3.6 Proposal Types

In the next series of tests, I examine how my results vary across several common types of proposals. I focus on four broad groups of proposals: uncontested director elections, votes to approve CEO's compensation, and proposals related to the board composition and voting rights allocation. Studying the effects of index fund holdings on a family's voting for these groups of proposals can help get a deeper insight into families' preferences and the types of issues they engage with.

I start by examining votes on uncontested director elections. Although such proposals are rarely contested (Bebchuk, 2007), the evidence suggests that these votes have a real impact on the tenures of CEOs and directors (Cai, Garner, and Walkling, 2009; Fos, Li, and Tsoutsoura, 2018) and the firm's policies (Ertimur, Ferri, and Oesch, 2018). I report the estimate in Column (1) of Table 3.9. I find that, while the size of active holdings on votes of active funds is negative and statistically significant, it is small in magnitude: a 20% increase in the size of active fund holdings leads to only 0.4% fewer active funds in the family supporting the management. The estimate for index funds is not statistically significant. The null effect of index funds for uncontested director elections is not surprising as the benefits of opposing such candidates are virtually nonexistent for the vast majority of such proposals.

Next, I examine votes to approve the executive's compensation. Section 951 of the Dodd-Frank Wall Street and Consumer Protection Act requires public companies to put an advisory vote on the executives' compensation (often referred to as 'say on pay') on the proxy ballot. The SEC adopted these rules in 2011, and firms must hold such votes at least once every three years. Although non-binding, such votes have important implications for firms' governance and the level of CEO pay (Ertimur, Ferri, and Oesch, 2013; Correa and Lel, 2016; Cai, Garner, and Walkling, 2009).

I report the results in Column (2). Unlike director elections, a family's votes on compensation are significantly affected by the size of its index fund holdings. A 20% increase in the size of index fund holdings leads to 2.5% fewer active funds from the family supporting compensation proposals. The effect is large in relative terms as, on average, only 10.6% of active funds within a family oppose compensation proposals.

In Column (3), I consider proposals related to the board structure, such as declassification of the boards and separating the CEO and chairman roles. I fail to find any evidence that index funds influence the votes of active funds in the same family for this group of proposals. Interestingly, the size of active funds negatively affects voting with management on board-related proposals.

Finally, in Column (4), I examine a subset of proposals related to the allocation of voting rights between the shareholders and the managers. They include proposals to reduce supermajority vote requirements and allow proxy access. Such proposals are important for the firm's shareholders as they allow them to vote against the management much easier in the future.[27] Cuñat, Gine, and Guadalupe (2012) documents that passage of such proposals leads to positive stock market reactions, increasing the shareholder value. I find that index funds have the largest effect on votes of active funds in their family for these proposals. A 20% increase in index fund holdings leads to 10.9% fewer active funds supporting the management on such proposals. The effect of active fund holdings is similar in magnitude.

Because the F-statistic for this subsample is low, I perform inference robust to weak instruments. In particular, the Anderson-Rubin statistic is 55.4, and I reject the null hypothesis that coefficients on endogenous variables (index and active holdings) are jointly zero.

3.6 Alternative Measures of Coordination and Voting

In my main analysis, I use the dummy *Disagreement* as the measure of coordination at the fund family level. A concern with the measure is that it can overestimate disagreement if a small fund in the family votes differently from the rest.

To address this concern, I use an alternative measure of family-level coordination, *FractionSharesMinority*, the fraction of shares held by the family that are vote in the minority. To illustrate, assume a family has three funds holding 100, 350, and 50 shares in a firm, respectively. If the first two funds supported the proposal and the third one did not, then the $FractionSharesMinority = 0.1$[28]

[27] For example, proxy access allows shareholders to nominate their directors, change in voting rules makes it easier to vote down directors.

[28] The fund family holds a total of 500 shares, 450 are voted in favor and 50 against. The minority is voting against, which yields $FractionSharesMinority = 50/(100 + 350 + 50) = 0.1$. Importantly, the minority

I report the results in Columns (1)–(3) of Table 3.10. The coefficients on index fund holdings are negative and statistically significant; a one standard deviation increase in log index holdings decreases the fraction of shares voted in the minority by 0.45-0.51 of its standard deviation. The effect is almost twice as large for a subset of contentious proposals.

Next, I consider an alternative measure of voting for the management, the fraction of family-held shares voted for the management. Although I use the fraction of funds in my main analysis, it generally overweights the votes of small funds.

I report the results in Columns (4)–(6) of Table 3.10. A 20% increase in index holdings decreases *FractionSharesForMgt* by 0.9%–1% for the whole sample and by 6.2% for contentious proposals. Overall, the results are qualitatively similar to those in Table 3.5.

3.6 Votes of Fund Family's Active, Index Funds

In the main analysis, I demonstrate that the votes of a family's active funds depend on the size of holdings of index funds from the same family. I argue that the finding is caused by the coordination in voting at the fund family level. If this is the case, one should expect votes of index funds and both index and active funds combined in the same family to follow the same pattern. I test this prediction and report the results in Table 3.11.

In Columns (1)–(3) I study how *FractFundsForMgt* depend on index and active holdings. Higher active and index holdings lead to fewer funds voting for the management, particularly on contentious proposals. A 20% increase in index fund holdings decreases the fraction of the family's fund voting for management by 0.7%–0.9% and by 5.6% for contentious proposals.

In Columns (4)–(6), I examine how *FractIndexFundsForMgt* depends on size of active and index holds. Consistent with family-level coordination, I find that votes cast by index funds depend on the size of active fund holdings in the same family. A 20% increase in index fund holdings decreases the fraction of the family's fund voting for management by 0.75%-0.92% and by 4.5% for contentious proposals. These magnitudes are almost identical to those in the main results.

in this context is defined by the way most funds in the family vote, not by whether the proposal passes or fails.

3.6 Do Families Monitor When there are no Active Holdings?

My main analysis focuses on cases when both index holdings and active holdings of a family in a firm are positive. It is interesting to examine whether family-level coordination improves monitoring when either active or index holdings of the family's funds in a firm are zero. For example, it is possible that, if a family has large combined index holdings in a firm, the stewardship team has incentives to monitor the firm more closely. To investigate this possibility, I estimate the effect of index (active) fund holdings on family-level voting when active (index) funds holdings are zero.[29] I focus on a subset of contentious proposals sponsored by either the shareholders or the management.

I report the results in Table 3.12. In Column (1), I reestimate the main specification for reference. In Column (2), I focus on a subsample with zero active holdings. Interestingly, an increase in the size of index fund holdings still reduces the fraction of a family's funds voting with the management. It is worth noting that the magnitude of the estimate is slightly lower than in Column (1). It is likely the case that when active funds do not have a stake in the firm, the family's incentives to monitor are lower. Similarly, in Column (3), when index holdings are zero, larger active holdings lead to lower support for the management. The magnitude of the estimate, however, is much smaller than in Column (1). A possible explanation for this difference is that without large index fund holdings, their family is unlikely to change the vote outcome and, in turn, invest in costly monitoring.

3.7 Conclusion

In this paper, I demonstrate that index funds can benefit corporate governance by increasing their fund families' voting power. Even when index funds have no incentives to engage with a firm's management, their fund families can use large blocks of shares held by their index funds to monitor.

[29] These results are reported for a subsample of contentious proposals. I follow the same empirical strategy as in the main analysis. When analyzing subsample with zero active holdings, I instrument index holdings with index flows, and similarly for observations with zero index holdings. For consistency, I focus on the same set of fund families as in the main analysis.

I find that increasing holdings of index funds has two effects. First, it facilitates the fund's voting coordination within the same fund family, increasing the frequency of full voting coordination at the family level. Second, their families' voting power coming from index funds is utilized to cast fewer pro-management votes, particularly for proposals that significantly affect the firm's value.

There has been some critique of the proxy voting conducted by index funds. For instance, Lund (2018) makes a case against the right of index funds to vote. This paper provides evidence that restricting their ability to do so is likely to be detrimental to corporate governance. Such regulation will most likely decrease the voting power at the disposal of fund families and decrease their investment in costly monitoring.

3.8 Tables and Figures

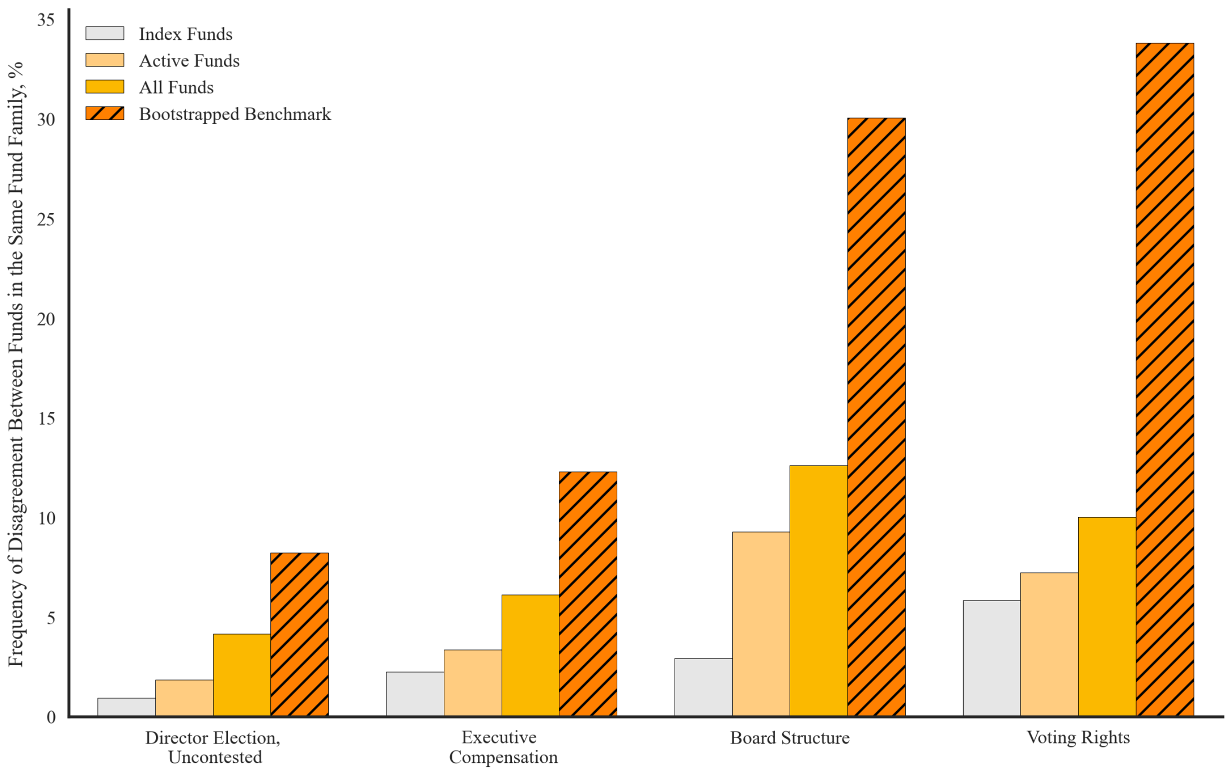

The figure shows the frequency of disagreement among funds from the same family voting on a proposal. The sample is restricted to fund family-proposal observations for which there are at least two index and two active funds voting on the proposal. Each bar shows the average *Disagreement* across observations from a particular group. I compute average *Disagreement* among index funds, active funds, and active and index funds from the same fund family. The above measure is computed for four types of proposals: uncontested director elections, advisory votes on the executive's compensation, and votes related to board composition and voting rules. For each type of proposal, I also compute the bootstrapped average frequency of *Disagreement*. It shows how often the disagreement would occur if funds were randomly assigned to fund families.

Figure 3.1. Fund Family-Level Voting Coordination

TABLE 3.1. Summary Statistics

This table reports summary statistics. The sample contains votes and holdings of 5,769 funds from 61 fund families from 2007 to 2016. **Panel A** reports descriptive statistics for the fund-level data. *Holding Value ($M)* is the value of a fund's holdings in a firm before the vote, and *MCap* is the firm's market value at the report date. *Fund TNA* is the size of the fund, and *Fund Flow* is the latest monthly flow to the fund before the report date. *Fund For Mgt* is a dummy that equals one if the fund voted according to the management's recommendation. **Panel B** reports summary statistics of the data aggregated at the fund family level. *Fraction (Active, Index) Funds for Mgt* is the fraction of (active, index) funds that vote according to the management's recommendation. *Fraction (Active, Index) Shares for Mgt* is the fraction of shares (shares held by active funds, index funds) that vote according to the management's recommendations. *Disagreement* is a dummy variable that equals one if funds from the same family did not cast identical votes on a proposal. *N Active (Index) Funds* is the number of active (index) funds voting on a proposal. *Index (Active) Value* is the total value of all index (active) fund holdings in the firm. *Index (Active) Flow* is the aggregated percentage monthly asset flow in all index (active) funds in the fund family over the month before the record date. *Fund Family TNA* is the fund family's total net asset value.

Panel A: Fund-Level Data ($N = 15,008,813$)					
	mean	std	10%	50%	90%
Holding Value ($M)	14.294	87.566	0.030	0.931	22.792
MCap ($B)	29.677	62.108	0.486	5.824	88.097
Fund TNA ($B)	2.217	7.898	0.008	0.242	4.785
Fund Flow	0.020	0.844	-0.038	-0.000	0.058
Fund for Mgt	0.921	0.270	1.000	1.000	1.000
Panel B: Fund Family-Level Data ($N = 546,322$)					
	mean	std	10%	50%	90%
Fraction Funds for Mgt	0.925	0.248	0.952	1.000	1.000
Fraction Active Funds for Mgt	0.920	0.267	1.000	1.000	1.000
Fraction Index Funds for Mgt	0.927	0.259	1.000	1.000	1.000
Fraction Shares for Mgt	0.929	0.250	0.999	1.000	1.000
Disagreement	0.041	0.199	0.000	0.000	0.000
Number of Funds	5.546	6.090	2.000	3.000	12.000
Number of Index Funds	2.664	3.868	1.000	1.000	5.000
Number of Active Funds	2.882	4.399	1.000	2.000	6.000
Index Value ($M)	55.075	424.919	0.135	1.856	56.748
Active Value ($M)	73.361	359.143	0.134	4.689	121.261
Index Flow	0.001	0.032	-0.028	0.003	0.031
Active Flow	-0.001	0.032	-0.024	0.000	0.022
Fund Family TNA ($B)	125.149	238.276	1.559	17.915	560.829

TABLE 3.2. Coordination and Voting in Large Fund Families

This table reports average coordination and support for the management in proxy voting by the largest fund families. The sample consists of proxy votes cast by families with a positive index fund and active fund holdings in a firm. Details on the sample construction are provided in Section 3. *Disagreement* is the dummy variable that equals one if funds in the same fund family did not cast identical votes on a proposal. *Fraction Funds for Mgt* is the fraction of funds within a family that votes according to the management's recommendation. The last two columns report the average fraction of funds that vote with management for a family's active and index funds, respectively.

Fund Family	Disagreement	Fraction Funds for Mgt	
		Active Funds	Index Funds
Vanguard	0.12%	94.35%	94.35%
BlackRock	4.21%	95.58%	95.63%
State Street Global Advisors	22.08%	74.23%	94.34%
Fidelity	12.08%	90.63%	86.61%
PowerShares Capital Management	3.53%	94.3%	96.61%
Rydex Investments	0.09%	99.47%	99.55%
T. Rowe Price	1.47%	92.46%	92.6%
TIAA-CREF Asset Management	0.0%	95.44%	95.43%

TABLE 3.3. First-Stage Estimates

This table reports the first-stage estimates of the instrumental variable analysis. The dependent variables are log*Index Holdings* and log*Active Holdings*. The instruments are *Index Flow* (*Active Flow*), defined as the total dollar inflow in the index (active) funds of a family scaled by the previous month's total family TNA. *log MCap* is the log of the market capitalization of a firm at the holdings report date. The observations are at the fund family-firm-proposal level. All variables, except for flows, are standardized. The standard errors clustered at the firm level are in parenthesis. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	(1) log Index Holdings	(2) log Active Holdings	(3) log Index Holdings	(4) log Active Holdings
IndexFlow	0.807***	-2.021***	1.341***	-1.846***
	(0.145)	(0.136)	(0.106)	(0.116)
ActiveFlow	0.358***	1.782***	0.048	1.642***
	(0.136)	(0.128)	(0.094)	(0.100)
log MCap	0.431***	0.385***	0.616***	0.643***
	(0.006)	(0.008)	(0.018)	(0.023)
log Fund Family TNA			0.638***	0.450***
			(0.005)	(0.005)
Constant	0.010*	0.013*	0.016***	0.016***
	(0.006)	(0.007)	(0.000)	(0.000)
Observations	519,896	521,502	519,896	521,502
R-squared	0.199	0.157	0.589	0.397
Firm FE	No	No	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

TABLE 3.4. Fund Family-Level Voting Coordination

This table reports the effect of fund family index and active holdings on family-level coordination in voting. The dependent variable in all columns is *Disagreement*, a dummy variable that equals one if the funds from the same family voted differently on the same proposal. log *Index Value* (log *Active Value*) is the log of total dollar holdings of index (active) funds from a family prior to the record day for a given proposal. log *MCap* is the log of the firm's market capitalization. Column (1) reports the OLS estimates. Columns (2)–(4) report IV estimates, where log *Active Holdings* and log *Index Holdings* are instrumented with $IndexFlow$ and $ActiveFlow$. *ISSforMgt* is the dummy that is equal to one if the voting recommendations of the ISS and the management coincide. log *FundFamilyTNA* is the log of the total size of the fund family's assets at the holdings report date. Column (5) contains estimates for a subset of contentious proposals for which the ISS and Management-issued recommendations differ. Column (6) focuses on a subsample of non-contentious proposals. The observations are at the fund family-proposal level. All variables are standardized. Standard errors clustered at the firm level are in the parenthesis. For each IV specification, I report the robust first stage F-statistic and the Anderson-Rubin test statistic. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	Dependent Variable = *Disagreement*					
	(1) OLS All Proposals	(2) IV All Proposals	(3) IV All Proposals	(4) IV All Proposals	(5) IV Contentious Proposals	(6) IV Non-Contentious Proposals
log IndexValue	0.080***	-0.582***	-0.642***	-0.441***	-1.506***	-0.296***
	(0.005)	(0.139)	(0.188)	(0.101)	(0.383)	(0.089)
log ActiveValue	-0.030***	-0.112*	-0.140**	-0.243***	-1.155***	-0.209***
	(0.004)	(0.059)	(0.067)	(0.063)	(0.436)	(0.051)
log MCap	0.022***	0.328***	0.355***	0.368***	1.701***	0.261***
	(0.005)	(0.070)	(0.128)	(0.098)	(0.533)	(0.083)
ISS for Mgt			-0.781***	-0.781***		
			(0.019)	(0.018)		
log Fund Family TNA				0.457***	1.759***	0.334***
				(0.089)	(0.425)	(0.077)
Observations	527,528	512,513	512,513	512,513	46,019	466,386
Firm FE	No	No	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
First Stage F-Stat		18.39	11.86	40.35	16.60	38.24
AR Test Statistic		43.36	33.55	32.47	31.93	24.09
AR Test p-value < 0.01		Yes	Yes	Yes	Yes	Yes

TABLE 3.5. Voting with the Management

This table reports the effect of fund family holdings on the family's decision to vote with the management. The dependent variable in all columns is *FractionActiveFunds-ForMgt*, the fraction of active funds from the fund family that voted according to the management-issued recommendation. log *Index Value* (log *Active Value*) is the log of total dollar holdings of index (active) funds from a family prior to the record day for a given proposal. log *MCap* is the log of the firm's market capitalization. The table focuses on management-sponsored proposals. Column (1) contains the OLS estimates. Columns (2)–(4) contains instrumental variable estimates, where log *Active Holdings* and log *Index Holdings* are instrumented with *Index Flow* and *Active Flow*. *ISSforMgt* is a dummy that equals one if the voting recommendations of the ISS and the management coincide. log *FundFamilyTNA* is the log of the total size of the fund family's assets at the holdings report date. Column (5) contains estimates for a subset of contentious proposals for which the ISS and Management-issued recommendations differed. Column (6) focuses on a subsample of non-contentious proposals. The observations are at the fund family-firm-proposal level. All variables are standardized. Standard errors clustered at the firm level are in the parenthesis. For each IV specification, I report the robust first stage F-statistic and the Anderson-Rubin test statistic. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	Dependent Variable = *FractionActiveFundsForMgt*					
	(1) OLS All Proposals	(2) IV All Proposals	(3) IV All Proposals	(4) IV All Proposals	(5) IV Contentious Proposals	(6) IV Non-Contentious Proposals
log IndexValue	-0.012***	-0.285**	-0.516***	-0.343***	-2.973***	0.064
	(0.005)	(0.119)	(0.165)	(0.094)	(0.747)	(0.066)
log ActiveValue	0.044***	-0.223***	-0.236***	-0.325***	-3.932***	-0.106***
	(0.004)	(0.035)	(0.044)	(0.050)	(0.762)	(0.034)
log MCap	-0.061***	0.150***	0.448***	0.459***	4.665***	0.074
	(0.008)	(0.057)	(0.104)	(0.085)	(0.964)	(0.059)
ISS for Mgt			1.881***	1.881***		
			(0.020)	(0.020)		
log Fund Family TNA				0.393***	4.175***	0.006
				(0.079)	(0.800)	(0.056)
Observations	527,528	512,513	512,513	512,513	46,019	466,386
Firm FE	No	No	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
First Stage F-Stat		18.39	11.86	40.35	16.60	38.24
AR Test Statistic		54.07	55.97	56.40	99.91	33.24
AR Test p-value < 0.01		Yes	Yes	Yes	Yes	Yes

TABLE 3.6. Fund Family Voting, Close Votes

This table reports the effect of the fund family index and active holdings on its voting for close proposals. The dependent variable in all columns is *Fraction Funds For Mgt*, the fraction of funds from the fund family that voted according to the management-issued recommendation. log *Index Value* (log *Active Value*) is the log of total dollar holdings of index (active) funds from a family prior to the record day for a given proposal. log *MCap* is the log of the firm's market capitalization. *ISSforMgt* is the dummy that equals one if the voting recommendations of the ISS and the management coincide. log *FundFamilyTNA* is the log of the total size of the fund family's assets at the holdings report date. All columns provide the IV estimates, where I instrument the holdings with fund family-level flows. *Vote Margin* is the absolute value of the difference between the percentage of votes received in favor of a proposal and the percentage of votes required for the proposal to pass. Column (1) reports estimates for a subsample of proposals that passed or failed with a margin lower than 10%. Column (2) focuses on the subsample with margins between 10% and 50%, Column (3) reports the estimates for proposals with a margin higher than 50%. The observations are at the fund family-proposal level. All variables are standardized. Standard errors clustered at the firm level are in the parenthesis. I also report the robust first stage F-statistic and the Anderson-Rubin test statistic. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	Dependent Variable = *Fraction Funds For Mgt*		
	(1) IV $\lvert$Vote Margin$\rvert$ < 10%	(2) IV 10% < $\lvert$Vote Margin$\rvert$ < 50%	(3) IV 50% < $\lvert$Vote Margin$\rvert$
log IndexValue	-1.989***	-0.624***	0.123
	(0.744)	(0.145)	(0.100)
log ActiveValue	-3.614***	-0.436***	-0.366***
	(0.790)	(0.068)	(0.081)
log MCap	3.577***	0.673***	0.254**
	(0.940)	(0.126)	(0.100)
ISS for Mgt	2.429***	2.002***	1.793***
	(0.141)	(0.026)	(0.043)
log Fund Family TNA	3.469***	0.639***	0.084
	(0.836)	(0.122)	(0.078)
Observations	9,713	359,605	142,598
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
First Stage F-Stat	11.28	22.97	36.91
AR Test Statistic	66.46	88.28	28.34
AR Test p-value < 0.01	Yes	Yes	Yes

TABLE 3.7. Counterfactual Analysis

This table reports estimates of the counterfactual exercise: how funds from a family would vote for management under a different allocation of funds between active and index funds. Formally, the table reports the results of a statistical test of the following null hypothesis:

$$\mathcal{H}_0 : \beta_I \log(1 + \Delta_{Index}) + \beta_A \log(1 + \Delta_{Active}) = 0 \tag{3.6}$$

against a two-sided alternative. The estimates of β_I and β_A are from Table 3.5. In all columns, I use $\Delta_{Index} = -37.07\%$ and $\Delta_{Active} = 27.81\%$. The Wald statistic is in the parenthesis. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

Outcome Variable = Δ *FractionFunds for Mgt*			
(1)	(2)	(3)	(4)
All Proposals	|Vote Margin| < 10%	10% < |Vote Margin| < 50%	|Vote Margin | > 50%
-0.4%	-6.7%	1.4%	-4.2%***
(0.28)	(1.92)	(2.44)	(12.47)

TABLE 3.8. Fund Family Voting, Proposal Sponsor

This table reports the effect of fund family holdings on the family's support for shareholder-sponsored and manager-sponsored proposals. The dependent variable in all columns is *FractionActiveFundsForMgt*, the fraction of active funds from the fund family that vote according to the management-issued recommendation. log *Index Value (*log *Active Value)* is the log of total dollar holdings of index (active) funds from a family before the record day for a given proposal. log *MCap* is the log of the firm's market capitalization. *ISSforMgt* is the dummy that is equal to one if the voting recommendations of the ISS and the management coincide. log *FundFamilyTNA* is the log of the total size of the fund family's assets at the holdings report date. All Columns report the IV estimates log *Active Holdings* and log *Index Holdings* are instrumented with *Index Flow* and *Active Flow*. Columns (1)–(2) and (4)–(5) report the estimates for samples of management-sponsored and shareholder-sponsored proposals, respectively. Columns (3) and (6) further restrict the samples to contentious proposals, for which the recommendations of the ISS and the management conflict with each other. The observations are at the fund family-proposal level. All variables are standardized. Standard errors clustered at the firm level are in the parenthesis. For each IV specification, I report the robust first stage F-statistic and the Anderson-Rubin test statistic. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	Dependent Variable = *FractionActiveFundsForMgt*					
	Management-Sponsored			Shareholder-Sponsored		
	(1) All Proposals	(2) All Proposals	(3) Contentious Proposals	(4) All Proposals	(5) All Proposals	(6) Contentious Proposals
log IndexValue	-0.187*	-0.174**	-1.981**	-2.112***	-3.226***	-4.152***
	(0.106)	(0.082)	(0.829)	(0.588)	(1.021)	(1.066)
log ActiveValue	-0.170***	-0.221***	-2.844***	-1.574***	-3.908***	-5.422***
	(0.029)	(0.040)	(0.670)	(0.479)	(1.227)	(1.434)
log MCap	0.154***	0.304***	3.538***	1.896***	4.667***	6.149***
	(0.049)	(0.072)	(0.968)	(0.429)	(1.494)	(1.663)
ISS for Mgt		1.704***			1.718***	
		(0.026)			(0.055)	
log Fund Family TNA		0.228***	2.896***		4.378***	5.915***
		(0.068)	(0.773)		(1.290)	(1.399)
Observations	482,648	482,648	25,174	29,865	29,808	20,790
Firm FE	No	Yes	Yes	No	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
First Stage F-Stat	16.83	40.33	12.84	6.950	5.220	7.692
AR Test Statistic	40.47	40.42	47.29	54.33	58.26	88.75
AR Test p-value < 0.01	Yes	Yes	Yes	Yes	Yes	Yes

TABLE 3.9. Fund Family Voting, Proposal Types

This table reports the effect of a fund family's index and active holdings on voting across several types of proposals. The dependent variable in all columns is *FractionActiveFunds-ForMgmt*, the fraction of active funds from the fund family that voted according to the management-issued recommendation. log *Index Value* (log *Active Value*) is the log of total dollar holdings of index (active) funds from a family prior to the record day for a given proposal. log *MCap* is the log of the market capitalization of a firm whose proposal is being voted on. *ISSforMgt* is the dummy that is equal to one if the voting recommendations of the ISS and the management coincide. log *FundFamilyTNA* is the log of the total size of the fund family's assets at the holdings report date. All columns provide the IV estimates, where I instrument the holdings with fund family-level flows. Column (1) provides the estimates for uncontested director elections, Column (2), for advisory votes on the executive's compensation. Column (3) provides the results for proposals to declassify the board of directors and separate the CEO and chairman roles. Column (4) provides the estimates for proposals that redistribute voting rights from the managers to shareholders, such as proposals to repeal the supermajority rule and proposals to allow proxy access. The observations are at the fund family-proposal level. All variables are standardized. Standard errors clustered at the firm level are in the parenthesis. I also report the robust first stage F-statistic and the Anderson-Rubin test statistic. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	Dependent Variable = *FractionActiveFundsForMgt*			
	(1)	(2)	(3)	(4)
	IV	IV	IV	IV
	Director Elections	Say on Pay	Board Related	Voting Rights
log IndexValue	-0.140	-1.033***	-0.484	-4.513**
	(0.087)	(0.273)	(0.360)	(2.191)
log ActiveValue	-0.208***	-0.220***	-0.780**	-4.557**
	(0.046)	(0.062)	(0.324)	(2.151)
log MCap	0.272***	0.850***	0.024	5.598**
	(0.079)	(0.192)	(0.466)	(2.568)
ISS for Mgt	1.655***	1.680***	2.372***	2.521***
	(0.034)	(0.041)	(0.096)	(0.248)
log Fund Family TNA	0.204***	0.856***	0.885**	5.571**
	(0.074)	(0.190)	(0.375)	(2.393)
Observations	362,305	36,349	5,469	5,184
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
First Stage F-Stat	37.15	20.93	13.35	2.918
AR Test Statistic	27.15	31.09	5.752	55.40
AR Test p-value < 0.01	Yes	Yes	No	Yes

TABLE 3.10. Family-Level Coordination and Voting, Alternative Measures

This table reexamines the effect of index fund holdings on fund family coordination in voting, using an alternative measure of coordination. The dependent variable in Columns (1)–(3) is *MinorityFractionShares*, the fraction of shares that are not voted as a majority by funds from the same family. For example, if the family has two funds that voted on their 100 and 300 shares differently, the minority vote is 25%. The dependent variable in Columns (4)–(6) is *FractionSharesForMgt*, the fraction of the family's shares voted for the management. log *Index Value* (log *Active Value*) is the log of total dollar holdings of index (active) funds from a family prior to the record day for a given proposal. log *MCap* is the log of the firm's market capitalization. All columns contains instrumental variable estimates, where log *Active Holdings* and log *Index Holdings* are instrumented with *Index Flow* and *Active Flow. ISSforMgt* is the dummy that equals one if the voting recommendations of the ISS and the management coincide. log *FundFamilyTNA* is the log of the total size of the fund family's assets at the holdings report date. Columns (3) and (6) report the estimates for a subsample of contentious proposals. The observations are at the fund family-firm-proposal level. All variables are standardized. Standard errors clustered at the firm level are in the parenthesis. I also report the robust first stage F-statistic and the Anderson-Rubin test statistic. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	MinorityFractionShares			*FractionSharesForMgt*		
	(1)	(2)	(3)	(4)	(5)	(6)
	IV	IV	IV	IV	IV	IV
	All Proposals	All Proposals	Contentious Proposals	All Proposals	All Proposals	Contentious Proposals
log IndexValue	-0.459***	-0.379***	-1.151***	-0.458***	-0.506***	-3.113***
	(0.121)	(0.095)	(0.401)	(0.143)	(0.109)	(0.865)
log ActiveValue	-0.021	-0.114**	-0.647*	-0.340***	-0.499***	-4.784***
	(0.044)	(0.051)	(0.371)	(0.045)	(0.057)	(0.877)
log MCap	0.212***	0.249***	1.060**	0.278***	0.670***	5.399***
	(0.057)	(0.085)	(0.495)	(0.067)	(0.099)	(1.110)
ISS for Mgt		-0.674***				
		(0.022)				
log Fund Family TNA		0.304***	1.098***		0.580***	4.563***
		(0.078)	(0.408)		(0.091)	(0.921)
Observations	512,513	512,513	46,019	512,513	512,513	46,019
Firm FE	No	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
First Stage F-Stat	18.39	40.35	16.60	18.39	40.35	16.60
AR Test Statistic	23.87	20.08	10.77	145.1	157.9	141.6
AR Test p-value < 0.01	Yes	Yes	Yes	Yes	Yes	Yes

TABLE 3.11. Robustness, Voting of Index Funds, and All Funds in the Family

This table reports the effect of index fund holdings on the decision of index funds in the family to vote with the management. In Columns (1)–(3), the dependent variable is *FractionFundsForMgt*, a fraction of funds within the family (index or active) that vote with the management. In Columns (4)–(6), the dependent variable is *FractionIndexFundsForMgt*, the fraction of index funds within a family that vote with the management. log *Index Value* (log *Active Value*) is the log of total dollar holdings of index (active) funds from a family prior to the record day for a given proposal. log *MCap* is the log of the firm's market capitalization. All columns report the IV estimates where the holdings are instrumented with *Index Flow* and *Active Flow*. *ISSforMgt* is the dummy that is equal to one if the voting recommendations of the ISS and the management coincide. log *FundFamilyTNA* is the log of the total size of the fund family's assets at the holdings report date. Columns (1)–(2) and (4)–(5) report the results for all proposals from my sample. Columns (3) and (6) contain the estimates for a subsample of contentious proposals for which the ISS and the management disagree in their voting recommendations. The observations are at the fund family-proposal level. All variables are standardized. Standard errors clustered at the firm level are in the parenthesis. I also report the robust first stage F-statistic and the Anderson-Rubin test statistic. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	FractionFundsForMgt			FractionIndexFundsForMgt		
	(1) All Proposals	(2) All Proposals	(3) Contentious Proposals	(4) All Proposals	(5) All Proposals	(6) Contentious Proposals
log IndexValue	-0.357***	-0.425***	-2.740***	-0.352***	-0.432***	-2.094***
	(0.132)	(0.102)	(0.796)	(0.134)	(0.105)	(0.760)
log ActiveValue	-0.307***	-0.446***	-4.387***	-0.378***	-0.527***	-4.317***
	(0.040)	(0.053)	(0.815)	(0.042)	(0.055)	(0.777)
log MCap	0.217***	0.588***	4.868***	0.246***	0.646***	4.424***
	(0.062)	(0.092)	(1.026)	(0.063)	(0.095)	(0.980)
ISS for Mgt		2.040***			1.986***	
		(0.020)			(0.020)	
log Fund Family TNA		0.501***	4.110***		0.524***	3.445***
		(0.085)	(0.855)		(0.087)	(0.819)
Observations	512,513	512,513	46,019	512,513	512,513	46,019
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
First Stage F-Stat	18.39	40.35	16.60	18.39	40.35	16.60
AR Test Statistic	109.5	122.1	116.3	176.5	191.6	128.5
AR Test p-value < 0.01	Yes	Yes	Yes	Yes	Yes	Yes

TABLE 3.12. Robustness: Fund Family Voting with Zero Index or Active Holdings

This table reports the effect of fund family holdings on voting when either index or active holdings of the family in a firm are zero. The outcome variable is *FractionFundsForMgt*, the fraction of funds within a family that voted with the management. I restrict the sample to contentious proposals. log *Index Value* (log *Active Value*) is the log of total dollar holdings of index (active) funds from a family before the record day for a given proposal. log *MCap* is the log of the firm's market capitalization. *ISSforMgt* is the dummy that equals one if the voting recommendations of the ISS and the management coincide. log *FundFamilyTNA* is the log of the total size of the fund family's assets at the holdings report date. In Column (1), I report my baseline specification for a subset of proposals for which both active and index funds from a family have positive holdings in a firm. I obtain the estimates by instrumenting active and index holdings with family-level flows, *Index Flow*, and *Active Flow*. In Column (2), I report the result for proposals for which only the family's index holdings are positive. I use *IndexFlow* as an instrument for index fund holdings. Similarly, In Column (3), I focus on a subset of proposals for which the family's active holdings are positive. I use *ActiveFLow* as the instrument for active holdings. The observations are at the fund family-proposal level. All variables are standardized. Standard errors clustered at the firm level are in the parenthesis. I also report the robust first stage F-statistic and the Anderson-Rubin test statistic. *, **, and *** indicate that the coefficient is statistically significant at the 10%, 5%, and 1% level, respectively.

	Dependent Variable = *FractionFundsForMgt*		
	(1) Index and Active	(2) Only Index	(3) Only Active
log IndexValue	-3.064***	-2.653***	
	(0.900)	(0.814)	
log ActiveValue	-4.522***		-1.881***
	(0.991)		(0.259)
log MCap	5.142***	1.590***	1.193***
	(1.231)	(0.490)	(0.164)
log Fund Family TNA	4.121***	1.469***	1.066***
	(0.997)	(0.473)	(0.089)
Observations	57,984	119,082	94,294
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
First Stage F-Stat	10.58	21.91	112.6
AR Test Statistic	95.97	34.14	97.66
AR Test p-value < 0.01	Yes	Yes	Yes

Bibliography

[1] Akey, P., and Appel, I. (2019). Environmental Externalities of Activism. Available at SSRN 3508808.

[2] Akey, P., and Appel, I. (2021). The Limits of Limited Liability: Evidence from Industrial Pollution. *Journal of Finance*, 76(1), 5-55.

[3] Alti, A. (2005). IPO Market Timing. *Review of Financial Studies*, 18(3), 1105-1138.

[4] Andrews, I., Stock, J. H., and Sun, L. (2019). Weak Instruments in Instrumental Variables Regression: Theory and Practice. *Annual Review of Economics*, 11, 727-753.

[5] Appel, I. R., Gormley, T. A., and Keim, D. B. (2016). Passive investors, not passive owners. *Journal of Financial Economics*, 121(1), 111-141.

[6] Appel, I. R., Gormley, T. A., and Keim, D. B. (2019). Standing on the Shoulders of Giants: The Effect of Passive Investors on Activism. *Review of Financial Studies*, 32(7), 2720-2774.

[7] Aran, Y. (2018). Beyond Covenants Not to Compete: Equilibrium in High-Tech Startup Labor Markets. *Stanford Law Review*, 70, 1235.

[8] Azar, J., Duro, M., Kadach, I., and Ormazabal, G. (2021). The Big Three and Corporate Carbon Emissions Around the World. *Journal of Financial Economics*.

[9] Babina, T., Ouimet, P., and Zarutskie, R. (2020). IPOs, Human Capital, and Labor Reallocation. Working Paper.

[10] Balasubramanian, N., Chang, J. W., Sakakibara, M., Sivadasan, J., and Starr, E. (2022). Locked in? The Enforceability of Covenants not to Compete and the Careers of High-Tech Workers. *Journal of Human Resources*, 57(S), S349-S396.

[11] Barber, B. M., Odean, T., and Zheng, L. (2005). Out of Sight, Out of Mind: The Effects of Expenses on Mutual Fund Flows. *Journal of Business*, 78(6), 2095-2120.

[12] Barber, B. M., Morse, A., and Yasuda, A. (2021). Impact Investing. *Journal of Financial Economics*, 139(1), 162-185.

[13] Barko, T., Cremers, M. and Renneboog, L. (2021). Shareholder Engagement on Environmental, Social, and Governance Performance. *Journal of Business Ethics*.

[14] Bebchuk, L. A., Cohen, A., and Ferrell, A. (2009). What Matters in Corporate Governance? *Review of Financial Studies*, 22(2), 783-827.

[15] Bebchuk, L. A., and Hirst, S. (2019). Index Funds and the Future of Corporate Governance: Theory, Evidence, and Policy. *Columbia Law Review*, 119(8), 2029-2146.

[16] Bebchuk, L. A. (2007). The Myth of the Shareholder Franchise. *Virginia Law Review*, 675-732.

[17] Benninga, S., Helmantel, M., and Sarig, O. (2005). The Timing of Initial Public Offerings. *Journal of Financial Economics*, 75(1), 115-132.

[18] Benveniste, L. M., Busaba, W. Y., and Wilhelm Jr, W. J. (2002). Information Externalities and the Role of Underwriters in Primary Equity Markets. *Journal of Financial Intermediation*, 11(1), 61-86.

[19] Berg, F., Koelbel, J. F., and Rigobon, R. (2019). Aggregate Confusion: the Divergence of ESG Ratings. Working Paper, MIT Sloan School of Management.

[20] Berk, J. B., and Green, R. C. (2004). Mutual Fund Flows and Performance in Rational Markets. *Journal of Political Economy*, 112(6), 1269-1295.

[21] Berk, J., and van Binsbergen, J. H. (2021). The Impact of Impact Investing. Available at SSRN 3909166.

[22] Bernstein, S. (2015). Does Going Public Affect Innovation? *Journal Finance*, 70(4), 1365-1403.

[23] Bernstein, S., Giroud, X., and Townsend, R. R. (2016). The Impact of Venture Capital Monitoring. *Journal of Finance*, 71(4), 1591-1622.

[24] Bialkowski, J., and Starks, L. T. (2016). SRI Funds: Investor Demand, Exogenous Shocks and ESG Profiles. Working Paper.

[25] Bias, D. (2022). Illiquid Equity, Labor Mobility, and Talent Allocation. Working Paper.

[26] Bishara, N. D. (2006). Covenants Not to Compete in a Knowledge Economy: Balancing Innovation from Employee Mobility Against Legal Protection for Human Capital Investment. *Berkeley Journal of Employment and Labor Law*, 27, 287.

[27] Bishara, N. D. (2011). Fifty Ways to Leave Your Employer: Relative Enforcement of Covenants not to Compete, Trends, and Implications for Employee Mobility Policy. *University of Pennsylvania Journal of Business Law*, 13, 751.

[28] Bolton, P., Li, T., Ravina, E., and Rosenthal, H. (2020). Investor Ideology. *Journal of Financial Economics*, 137(2), 320-352.

[29] Bodnaruk, A., Kandel, E., Massa, M., and Simonov, A. (2008). Shareholder Diversification and the Decision to Go Public. *Review of Financial Studies*, 21(6), 2779-2824.

[30] Borusyak, K., Jaravel, X., and Spiess, J. (2021). Revisiting Event Study Designs: Robust and Efficient Estimation. arXiv preprint arXiv:2108.12419.

[31] Brau, J. C., and Fawcett, S. E. (2006). Initial Public Offerings: an Analysis of Theory and Practice. *Journal of Finance*, 61(1), 399-436.

[32] Brav, A., Jiang, W., Li, T., and Pinnington, J. (2021). Picking Friends Before Picking (Proxy) Fights: How Mutual Fund Voting Shapes Proxy Contests. Working paper, Duke University.

[33] Broccardo, E., Hart, O. D., and Zingales, L. (2020). Exit vs. Voice. Working paper, National Bureau of Economic Research.

[34] Bubb, R., and Catan, E. (2020). The Party Structure of Mutual Funds. Working paper, New York University.

[35] Cai, J., Garner, J. L., and Walkling, R. A. (2009). Electing Directors. *Journal of Finance*, 64(5), 2389-2421.

[36] Callaway, B., and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200-230.

[37] Chatterji, A. K., Durand, R., Levine, D. I., and Touboul, S. (2016). Do Ratings of Firms Converge? Implications for Managers, Investors and Strategy researchers. *Strategic Management Journal*, 37(8), 1597-1614.

[38] Chemmanur, T. J., He, S., and Nandy, D. K. (2010). The Going-Public Decision and the Product Market. *Review of Financial Studies*, 23(5), 1855-1908.

[39] Chod, J., and Lyandres, E. (2011). Strategic IPOs and Product Market Competition. *Journal of Financial Economics*, 100(1), 45-67.

[40] Choi, J. J., Laibson, D., and Madrian, B. C. (2010). Why Does the Law of One Price Fail? An Experiment on Index Mutual Funds. *Review of Financial Studies*, 23(4), 1405-1432.

[41] Choi, S. J., Fisch, J. E., and Kahan, M. (2013). Who Calls the Shots: How Mutual Funds Vote on Director Elections. *Harvard Business Law Review*, 3, 35.

[42] Christoffersen, S. E., Musto, D. K., and Wermers, R. (2014). Investor Flows to Asset Managers: Causes and Consequences. *Annual Review of Financial Economics*, 6(1), 289-310.

[43] Clifford, C. P., Fulkerson, J. A., and Jordan, B. D. (2014). What Drives ETF Flows? *Financial Review*, 49(3), 619-642.

[44] Coase, R. H. (1960). The Problem of Social Cost. *Journal of Law and Economics*, 3, 1-44.

[45] Correa, R., and Lel, U. (2016). Say on pay laws, executive compensation, pay slice, and firm valuation around the world. *Journal of Financial Economics*, 122(3), 500-520.

[46] Corum, A. A., Malenko, A., and Malenko, N. (2021). Corporate Governance in the Presence of Active and Passive Delegated Investment. Working Paper, Cornell University.

[47] Cuñat, V., Gine, M., and Guadalupe, M. (2012). The Vote is Cast: The Effect of Corporate Governance on Shareholder Value. *Journal of Finance*, 67(5), 1943-1977.

[48] Cvijanović, D., Dasgupta, A., and Zachariadis, K. E. (2016). Ties That Bind: How Business Connections Affect Mutual Fund Activism. *Journal of Finance*, 71(6), 2933-2966.

[49] Dambra, M., Field, L. C., and Gustafson, M. T. (2015). The JOBS Act and IPO Volume: Evidence that Disclosure Costs Affect the IPO Decision. *Journal of Financial Economics*, 116(1), 121-143.

[50] Davis, G. F., and Kim, E. H. (2007). Business Ties and Proxy Voting by Mutual Funds. *Journal of Financial Economics*, 85(2), 552-570.

[51] De Chaisemartin, C., and d'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9), 2964-96.

[52] Dikolli, S. S., Frank, M. M., Guo, M. Z., and Lynch, L. J. (2021). Walk the Talk: ESG Mutual Fund Voting on Shareholder Proposals. Available at SSRN 3849762.

[53] Dixit, A. K., Pindyck, R. S. (1994). Investment Under Uncertainty. Princeton University Press.

[54] Doidge, C., Karolyi, G. A., and Stulz, R. M. (2013). The US Left Behind? Financial Globalization and the Rise of IPOs Outside the US. *Journal of Financial Economics*, 110(3), 546-573.

[55] Doidge, C., Karolyi, G. A., and Stulz, R. M. (2017). The US Listing Gap. *Journal of Financial Economics*, 123(3), 464-487.

[56] Draho, J. (2000). The Timing of Initial Public Offerings: A Real Option Approach. Available at SSRN 271351.

[57] Edmans, A., Fang, V. W., and Zur, E. (2013). The Effect of Liquidity on Governance. *Review of Financial Studies*, 26(6), 1443-1482.

[58] Edmans, A., and Holderness, C. G. (2017). Blockholders: a Survey of Theory and Evidence. *Handbook of the Economics of Corporate Governance*, 1, 541-636.

[59] Ertimur, Y., Ferri, F., and Oesch, D. (2013), Shareholder Votes and Proxy Advisors: Evidence from Say on Pay. *Journal of Accounting Research*, 51: 951-996.

[60] Ertimur, Y., Ferri, F., and Oesch, D. (2018). Understanding Uncontested Director Elections. *Management Science*, 64(7), 3400-3420.

[61] Evans, R. B. (2010). Mutual Fund Incubation. *Journal of Finance*, 65(4), 1581-1611.

[62] Ewens, M., and Farre-Mensa, J. (2020). The Deregulation of the Private Equity Markets and the Decline in IPOs. *Review of Financial Studies*, 33(12), 5463-5509.

[63] Ewens, M., and Marx, M. (2018). Founder Replacement and Startup Performance. *Review of Financial Studies*, 31(4), 1532-1565.

[64] Field, L. C., and Hanka, G. (2001). The Expiration of IPO Share Lockups. *Journal of Finance*, 56(2), 471-500.

[65] Field, L. C., and Karpoff, J. M. (2002). Takeover Defenses of IPO Firms. *Journal of Finance*, 57(5), 1857-1889.

[66] Filali Adib, F. Z. (2019). Passive Aggressive: How Index Funds Vote on Corporate Governance Proposals. Working paper, Copenhagen Business School.

[67] Friedman, H. L., and Heinle, M. S. (2021). Interested Investors and Intermediaries: When Do ESG Concerns Lead to ESG Performance? Jacobs Levy Equity Management Center for Quantitative Financial Research Paper.

[68] Fos, V., Li, K., and Tsoutsoura, M. (2018). Do Director Elections Matter? *Review of Financial Studies*, 31(4), 1499-1531.

[69] Gao, X., Ritter, J. R., and Zhu, Z. (2013). Where Have All the IPOs Gone? *Journal of Financial and Quantitative Analysis*, 48(6), 1663-1692.

[70] Garmaise, M. J. (2011). Ties That Truly Bind: Noncompetition Agreements, Executive Compensation, and Firm Investment. *Journal of Law, Economics, and Organization*, 27(2), 376-425.

[71] Gaspar, J. M., Massa, M., and Matos, P. (2006). Favoritism in Mutual Fund Families? Evidence on Strategic Cross-Fund Subsidization. *Journal of Finance*, 61(1), 73-104.

[72] Gibson Brandon, R., Krueger, P., and Schmidt, P. S. (2021). ESG Rating Disagreement and Stock Returns. Financial Analysts Journal, 77(4), 104-127.

[73] Gillan, S. L., and Starks, L. T. (2000). Corporate Governance Proposals and Shareholder Activism: The Role of Institutional Investors. *Journal of Financial Economics*, 57(2), 275-305.

[74] Gillan, S. L., Koch, A., and Starks, L. T. (2021). Firms and Social Responsibility: a Review of ESG and CSR Research in Corporate Finance. *Journal of Corporate Finance*, 101889.

[75] Gompers, P., Ishii, J., and Metrick, A. (2003). Corporate Governance and Equity Prices. *Quarterly Journal of Economics*, 118(1), 107-156.

[76] Goodman-Bacon, A. (2021). Difference-in-differences with Variation in Treatment Timing. *Journal of Econometrics.*

[77] Gormley, T. A., Gupta, V. K., Matsa, D. A., Mortal, S., and Yang, L. (2021). The Big Three and Board Gender Diversity: The Effectiveness of Shareholder Voice. European Corporate Governance Institute–Finance Working Paper, 714, 2020.

[78] Grewal, J., and Serafeim, G. (2020). Research on Corporate sustainability: Review and Directions for Future Research. *Foundations and Trends in Accounting*, 14(2), 73-127.

[79] Gupta, A. (2022). Labor Mobility, Firm Monopsony, and Entrepreneurship: Evidence from Immigration Wait-Lines. Working Paper, University of North Carolina.

[80] Gurun, U. G., Stoffman, N., and Yonker, S. E. (2021). Unlocking Clients: The Importance of Relationships in the Financial Advisory Industry. *Journal of Financial Economics*, 141(3), 1218-1243.

[81] Hart, O., and Moore, J. (1994). A Theory of Debt Based on the Inalienability of Human Capital. *Quarterly Journal of Economics*, 109(4), 841-879.

[82] Hart, O., and Zingales, L. (2017). Companies Should Maximize Shareholder Welfare not Market Value. ECGI-Finance Working Paper, (521).

[83] Hartzmark, S. M., and Sussman, A. B. (2019). Do Investors Value Sustainability? A Natural Experiment Examining Ranking and Fund Flows. *Journal of Finance*, 74(6), 2789-2837.

[84] Heath, D., Macciocchi, D., Michaely, R., and Ringgenberg, M. (2021). Do Index Funds Monitor? *Review of Financial Studies.*

[85] Heath, D., Macciocchi, D., Michaely, R., and Ringgenberg, M. C. (2021). Does Socially Responsible Investing Change Firm Behavior? Available at SSRN 3837706.

[86] Heinkel, R., Kraus, A., and Zechner, J. (2001). The Fffect of Green Investment on Corporate Behavior. *Journal of Financial and Quantitative Analysis*, 36(4), 431-449.

[87] Holmström, B., and Tirole, J. (1993). Market Liquidity and Performance Monitoring. *Journal of Political Economy*, 101(4), 678-709.

[88] Hong, H., and Kacperczyk, M. (2009). The Price of Sin: The Effects of Social Norms on Markets. *Journal of Financial Economics*, 93(1), 15-36.

[89] Ibbotson, R. G., and Jaffe, J. F. (1975). 'Hot Issue' Markets. *Journal of Finance*, 30(4), 1027-1042.

[90] Jeffers, J. (2019). The Impact of Restricting Labor Mobility on Corporate Investment and Entrepreneurship. Working Paper.

[91] Iliev, P., and Lowry, M. (2015). Are Mutual Funds Active Voters? *Review of Financial Studies*, 28(2), 446-485.

[92] Investment Company Institute (2021). Factbook: a Review of Trends and Activities in the Investment Company Industry. https://www.ici.org/system/files/2021-05/2021_factbook.pdf

[93] Kleibergen, F., and Paap, R. (2006). Generalized Reduced Rank Tests Using the Singular Value Decomposition. *Journal of Econometrics*, 133(1), 97-126.

[94] Lakkis, E. (2021). Fund Family Matters: How Index Funds Improve Corporate Governance. Available at SSRN 3557588.

[95] Landier, A., and Lovo, S. (2020). ESG Investing: How to Optimize Impact? HEC Paris Research Paper No. FIN-2020-1363.

[96] Larcker, D. F., Tayan, B., and Watts, E. M. (2018). Cashing it in: Private-company exchanges and employee stock sales prior to ipo. Rock Center for Corporate Governance at Stanford University Closer Look Series: Topics, Issues and Controversies in Corporate

Governance No. CGRP-73, Stanford University Graduate School of Business Research Paper, (18-45).

[97] Levit, D., and Malenko, N. (2011). Nonbinding Voting for Shareholder Proposals. *Journal of Finance*, 66(5), 1579-1614.

[98] Lewellen, J., and Lewellen, K. (2021). Institutional Investors and Corporate Governance: The Incentives to Be Engaged, *Journal of Finance*, forthcoming.

[99] Liang, H., and Renneboog, L. (2017). On the Foundations of Corporate Social Responsibility. *Journal of Finance*, 72(2), 853-910.

[100] Liang, H., and Renneboog, L. (2020). Corporate Social Responsibility and Sustainable Finance: a Review of the Literature. European Corporate Governance Institute–Finance Working Paper, (701).

[101] Loughran, T., and Ritter, J. (2004). Why Has IPO Underpricing Changed Over Time? *Financial management*, 5-37.

[102] Lowry, M., Michaely, R., and Volkova, E. (2017). Initial Public Offerings: A Synthesis of the Literature and Directions for Future Research. *Foundations and Trends in Finance.*

[103] Lowry, M., Michaely, R., and Volkova, E. (2020). Information Revealed through the Regulatory Process: Interactions between the SEC and Companies ahead of Their IPO. *Review of Financial Studies*, 33(12), 5510-5554.

[104] Lund, D. S. (2017). The Case Against Passive Shareholder Voting. *Journal of Corporate Law.*, 43, 493.

[105] Malenko, N., and Shen, Y. (2016). The Role of Proxy Advisory Firms: Evidence from a Regression-Discontinuity Design. *Review of Financial Studies*, 29(12), 3394-3427.

[106] Marx, M., Strumsky, D., and Fleming, L. (2009). Mobility, Skills, and the Michigan Non-Compete Experiment. *Management Science*, 55(6), 875-889.

[107] Marx, M., Singh, J., and Fleming, L. (2015). Regional Disadvantage? Employee Non-Compete Agreements and Brain Drain. *Research Policy*, 44(2), 394-404.

[108] Matos, P. (2020). ESG and Responsible Institutional Investing Around the World: a Critical Review.

[109] Matsa, D. A. (2018). Capital structure and a firm's workforce. *Annual Review of Financial Economics*, 10, 387-412.

[110] Matvos, G., and Ostrovsky, M. (2008). Cross-ownership, Returns, and Voting in Mergers. *Journal of Financial Economics*, 89(3), 391-403.

[111] Matvos, G., and Ostrovsky, M. (2010). Heterogeneity and Peer Effects in Mutual Fund Proxy Voting. *Journal of Financial Economics*, 98(1), 90-112.

[112] McAdams, J. M. (2019). Non-compete Agreements: A Review of the Literature. Available at SSRN 3513639.

[113] McCahery, J. A., Sautner, Z., and Starks, L. T. (2016). Behind the Scenes: The Corporate Governance Preferences of Institutional Investors. *Journal of Finance*, 71(6), 2905-2932.

[114] Metrick, A., and Yasuda, A. (2011). Venture Capital and the Finance of Innovation. John Wiley and Sons.

[115] Michaely, R., Ordonez-Calafi, G., and Rubio, S. (2021). ES Votes that Matter. Working Paper.

[116] Morgan, A., Poulsen, A., Wolf, J., and Yang, T. (2011). Mutual Funds as Monitors: Evidence from Mutual Fund Voting. *Journal of Corporate Finance*, 17(4), 914-928.

[117] Morgan, A., Poulsen, A., Wolf, J., and Yang, T. (2011). Mutual Funds as Monitors: Evidence from Mutual Fund Voting. *Journal of Corporate Finance*, 17(4), 914-928.

[118] Morningstar (2020). 2020 U.S. Fund Fee Study. https://www.morningstar.com/content/dam/marketing/shared/pdfs/Research/annual-us-fund-fee-study-updated.pdf

[119] Nanda, V., Wang, Z. J., and Zheng, L. (2004). Family Values and the Star Phenomenon: Strategies of Mutual Fund Families. *Review of Financial Studies*, 17(3), 667-698.

[120] Naaraayanan, S. L., Sachdeva, K., and Sharma, V. (2021). The Real Effects of Environmental Activist Investing. European Corporate Governance Institute–Finance Working Paper, (743).

[121] Nishesh, N., Ouimet, P., and Simintzi, E. (2022). Labor and Corporate Finance. Available at SSRN.

[122] Oyer, P., and Schaefer, S. (2005). Why Do Some Firms Give Stock Options to All Employees?: An Empirical Examination of Alternative Theories. *Journal of Financial Economics*, 76(1), 99-133.

[123] Pagano, M. (1993). The Flotation of Companies on the Stock Market: a Coordination Failure Model. *European Economic Review*, 37(5), 1101-1125.

[124] Pagano, M., Panetta, F., and Zingales, L. (1998). Why Do Companies go Public? An Empirical Analysis. *Journal of Finance*, 53(1), 27-64.

[125] Pagano, M., and Röell, A. (1998). The Choice of Stock Ownership Structure: Agency Costs, Monitoring, and the Decision to Go Public. *Quarterly Journal of Economics*, 113(1), 187-225.

[126] Pagano, M. (2020). Risk Sharing Within the Firm: A Primer. *Foundations and Trends® in Finance* 12(2):117–198.

[127] Pástor, Ľ., Stambaugh, R. F., and Taylor, L. A. (2021). Sustainable Investing in Equilibrium. *Journal of Financial Economics*, 142(2), 550-571.

[128] Pigou, A. C. (1924). The Economics of Welfare. Macmillan.

[129] Puri, M., and Zarutskie, R. (2012). On the Life Cycle Dynamics of Venture-Capital-and non-Venture-Capital-Financed Firms. *Journal of Finance*, 67(6), 2247-2293.

[130] Renneboog, L., Ter Horst, J., and Zhang, C. (2008). Socially Responsible Investments: Institutional Aspects, Performance, and Investor Behavior. *Journal of Banking and Finance*, 32(9), 1723-1742.

[131] Ritter, J. R. (1984). The 'Hot Issue' Market of 1980. *Journal of Business*, 215-240.

[132] Rothberg, B., and Lilien, S. (2006). Mutual Funds and Proxy Voting: New Evidence on Corporate Governance. *Journal of Business and Technology Law*, 1, 157.

[133] Sanati, A. (2018). How does Labor Mobility Affect Corporate Leverage and Investment? Available at SSRN 2986582.

[134] Shleifer, A., and Vishny, R. W. (1986). Large Shareholders and Corporate Control. *Journal of Political Economy*, 94(3, Part 1), 461-488.

[135] Stambaugh, R. F. (1999). Predictive Regressions. *Journal of Financial Economics*, 54(3), 375-421.

[136] Starr, E. (2019). The Use, Abuse, and Enforceability of Non-Compete and No-Poach Agreements. Washington, DC: Economic Innovation Group.

[137] Starr, E., Prescott, J. J., and Bishara, N. (2021). Noncompetes in the US Labor Force. *Journal of Law and Economics.*

[138] Stein, J. C. (1997). Internal Capital Markets and the Competition for Corporate Resources. *Journal of Finance*, 52(1), 111-133.

[139] Stuart, T., and Sorenson, O. (2003). The Geography of Opportunity: Spatial Heterogeneity in Founding Rates and the Performance of Biotechnology Firms. *Research policy*, 32(2), 229-253.

[140] Sun, L., and Abraham, S. (2021). Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects. *Journal of Econometrics*, 225(2), 175-199.

[141] Xiong, X., and Png, I. P. (2019). Location of US Manufacturing, 1987-2014: A New Dataset. Available at SSRN 3401582.

[142] Yermack, D. (2010). Shareholder Voting and Corporate Governance. *Annual Review of Financial Economics*, 2(1), 103-125.

[143] Ysmailov, B. (2020). Non-Compete Agreements and Capital Structure Decisions. Working Paper, Nazarbayev University.

[144] Younge, K. A., and Marx, M. (2016). The Value of Employee Retention: Evidence from a Natural Experiment. *Journal of Economics and Management Strategy*, 25(3), 652-677.