Expanding Shareholder Voice: The Impact of SEC Guidance on Environmental and Social Proposals*

Kenneth Khoo[†]
kenneth.khoo@nus.edu.sg

Roberto Tallarita[‡]
rtallarita@law.harvard.edu

Abstract

After a dramatic increase over the past decade, shareholder support for environmental and social (E&S) proposals seems to have waned. In this Article, we examine whether this recent decline is linked to a 2021 shift in the SEC's policy, which expanded the ability of shareholders to influence E&S corporate decisions. We suggest that this regulatory shift has led to an increase in "prescriptive" E&S proposals, which typically call for more aggressive but costlier E&S policies by companies. Using a combination of supervised and unsupervised machine learning techniques to identify prescriptive proposals, we find that these proposals generally receive less shareholder support and seem to be driving a substantial part of the decline in support for E&S proposals. This decline is observed among the vast majority of institutional investors, including many ESG funds. However, there is considerable heterogeneity in the magnitude of this decrease across different investor groups. By classifying investors according to their ideological preferences over E&S issues, we find that investors with more intense preferences for E&S issues are more likely to support prescriptive proposals, while those at the opposite end of this spectrum are more likely to oppose them. Our results suggest that while investors continue to vote along ideological lines on E&S issues, the financial costs of prescriptive proposals may outweigh the intensity of E&S preferences for most of them.

*For helpful comments and discussions, we are grateful to Sumit Agarwal, Lucian Bebchuk, Emiliano Catan, John Coates, Alma Cohen, Jared Ellias, Henry Hansmann, Oliver Hart, Scott Hirst, Edwin Hu, Louis Kaplow, Kobi Kastiel, Alvin Klevorick, Mariana Pargendler, Mark Roe, Alan Schwartz, Holger Spamann, Matthew Stephenson, Johan Sulaeman, Hans Tjio, Rory Van Loo, Alan Miller, Virginia Harper Ho, and participants at the UC Berkeley Forum on Corporate Governance, the 2024 Stigler Center-CEPR Political Economy of Finance Conference (Chicago Booth), the 2024 Conference on Empirical Legal Studies at Emory School of Law, the Corporate and Financial Law Hub Workshop at LSE, the Empirical Law & Finance Workshop at Harvard Law School, the Law, Economics, and Organizations Workshop at Harvard Law School, the Finance Brownbag Seminar at NUS Business School, the SIG Applied Microeconomics Workshop at NUS Economics Department, the 2025 Annual Meeting of the American Law and Economics Association at NYU Law School, the 2025 Annual Meeting of the Asian Law and Economics Association at the City University of Hong Kong, and the Harvard Law School Faculty Workshop. We are also grateful to the EW Barker Centre for Law & Business (Khoo) and the Harvard Law School Summer Research Fund (Tallarita) for generous funding, and to Will Magliocco, Megan Tan, Joshua Toh, and Bai Xueqi for outstanding research assistance.

[†]National University of Singapore, Faculty of Law.

[‡]Harvard Law School.

1 Introduction

Shareholders of major publicly traded corporations frequently use the Securities Exchange Commission (SEC)'s Rule 14a-8 to introduce proposals addressing environmental, political, ethical, or social issues. Such proposals have included calls for ExxonMobil to reduce its greenhouse gas emissions, for Meta to remedy its gender pay gaps, and for Lululemon to discontinue the use of down feathers in its merchandise. Historically, these environmental and social ("E&S") proposals have generally garnered low shareholder support (He et al. (2023)). However, this landscape has changed significantly in recent years. As Tallarita (2022) documents, the average shareholder vote in favor of E&S proposals was 18% in 2010, yet it nearly doubled to over 35% by 2021. While from 2010 to 2019, only 1% of E&S proposals attained majority support at annual meetings, this figure rose to 16% in 2020 and 2021.

The rising support for E&S proposals among shareholders is consistent with a number of plausible theories. For example, investors may have become increasingly aware of the growing risks of climate change and other E&S-related risks to their investment portfolios (Krueger et al. (2020); Ilhan et al. (2023); Bolton and Kacperczyk (2021); Ilhan et al. (2021); Bolton and Kacperczyk (2023)). Furthermore, investors may derive non-pecuniary utility from acting in a pro-social fashion (Andreoni (1990)), or may otherwise have altruistic preferences (Hart and Zingales (2017)) and are therefore willing to sacrifice pecuniary returns in order to pursue social goals (Barber et al. (2021); Hirst et al. (2023); Hart et al. (2024)). While voting has historically been viewed as a costlier alternative to simply exiting the firm (Admati and Pfleiderer (2009)), recent studies by Li et al. (2022) and Brav et al. (2022) have shown that shareholders often prefer to exercise their voting rights over divestment. Indeed, Broccardo et al. (2022) illustrate how voting may be more effective than exit in achieving "socially responsible" outcomes.

Despite this dramatic rise over the last decade, the trend has notably shifted in the opposite direction in 2022 and 2023. In Panel A of Figure 1, we plot shareholder support for E&S proposals over time, measured as the mean percentage of "votes for as a percentage of votes cast." In 2021, shareholder support for environmental proposals stood at 40.24%, exceeding that for governance proposals (35.52%). However, in 2022, support for environmental proposals dropped to 34.00%, followed by a further decline to 19.01% in 2023. Social proposals also experienced a decrease in support, from 35.27% in 2021 to 24.51% in 2022, and then to 16.96% in 2023. Meanwhile, as Figure 1 illustrates, the decline in voting support for governance proposals has been comparatively muted, decreasing from 35.52% to 26.50% between 2021 and 2023.

In this article, we examine whether the declining support for E&S proposals is related to a recent policy change that expanded the ability of shareholders to submit such proposals.[1] Traditionally, under the so-called "ordinary business exclusion," the SEC allowed corporate management to exclude E&S proposals if they included specific goals, methods, or time-frames for implementing the proposals. In November 2021, however, the SEC's Division of Corporation Finance issued a new "staff legal bulletin" (the "2021 Guidance") (SEC (2021)) that permitted the submission of such "prescriptive" proposals.

The concept of "prescriptiveness" captures the degree to which a shareholder proposal seeks to influence matters typically reserved for the board, by requesting the company to adopt specific policies, timelines, or targets. This usage aligns both with the lay meaning of the term (WSJ (2022); Blackrock (2022)) and with the technical meaning adopted by the SEC (SEC (2017)). For example, a non-prescriptive proposal might request a description of the company's plans to reduce its carbon emissions, whereas a prescriptive proposal might ask the board to identify specific short-, medium-, and long-term targets to reduce the company's carbon emissions in line with the Paris Climate Agreement's goals.[2]

A plausible economic intuition for why prescriptiveness matters to investors is that prescriptive proposals signal a heightened commitment to E&S objectives. When voting on E&S proposals–which deal with the negative effects of the company's activities on workers, consumers, non-human animals, or the environment–shareholders weigh their commitment to E&S objectives against the potential pecuniary costs to the firm and, by extension, to their own portfolios.[3] By articulating stronger commitments to E&S objectives, prescriptive proposals change the relevant trade-off between those objectives and the pecuniary costs required to achieve them. Returning to the earlier example, the disclosure of management's climate-transition plans would, in most practical scenarios, entail lower implementation costs than the identification of quantified emissions targets to be achieved within a specified timeline. Voting in favor of the latter therefore confronts shareholders with a starker trade-off: achieving a greater pro-social benefit may come at the cost of diminished financial returns.

Note that, regardless of how "prescriptive" the language is, shareholder proposals are merely advisory in nature and do not bind the board of directors even if they are approved by the majority of shareholders (Tallarita (2022)). However, proposals that secure majority or significant minority support can still influence managerial decisions because directors value their reputations and

[1]For an illustration of how scholars leverage regulatory changes to identify shifts in behavior, see Dammann (2022).

[2]See Table 2 for some examples of these such proposals.

[3]We use the terms "financial" and "pecuniary" interchangeably throughout this article, as well as "pro-social" and "non-pecuniary."

their long-term relationships with shareholders (Del Guercio and Hawkins (1999); Ertimur et al. (2010); Matvos and Ostrovsky (2010)). Empirical work shows that directors do respond to the voting outcomes of advisory proposals, even in settings characterized by low-cost shareholder activism (Del Guercio et al. (2008)). All else equal, adopting a prescriptive proposal should make concrete and costlier E&S actions more likely.

According to this framework, the 2021 Guidance's impact on shareholder voting should be significantly affected by the intensity of E&S preferences among investors. For instance, if the majority of investors had relatively weak pro-social preferences, one would expect to see a decrease in voting support attributable to the 2021 Guidance and the consequent increase in prescriptive proposals (Bolton et al. (2020); Curtis et al. (2021); Bebchuk and Hirst (2019); Bebchuk and Hirst (2022); Griffin (2020); Zytnick (2022)).

Reactions to whether the 2021 Guidance had an effect on the support for E&S proposals were mixed. On one hand, journalists (WSJ (2022)), legal practitioners (Posner (2022); Gibson-Dunn (2022); Gibson-Dunn (2023)), and even major institutional investors (Blackrock (2022)) have suggested that the greater "prescriptiveness" of shareholder proposals associated with the 2021 Guidance may have contributed to the decline in support for E&S proposals. On the other hand, skeptics like Morgan (2024) have argued that the 2021 Guidance has had little to no effect on voting support. Indeed, Morgan (2024) contends that the observed decrease in voting for E&S proposals was due to "a dramatic increase in anti-ESG proposals, which have proven largely unpersuasive and acted as an anchor on average vote totals," and that "the SEC's no-action process remains remarkably lop-sided in favor of [management]."

To assess the effect of the SEC's policy shift on the shareholder support for E&S proposals, we employ a mix of supervised and unsupervised machine learning methods in Natural Language Processing (NLP) to ascertain the prescriptive nature of proposals. Our supervised machine learning algorithm exploits the SEC staff's own assessment of a contested proposal's prescriptiveness before 2021. In other words, to establish whether a given proposal is prescriptive, we do not use our subjective interpretation of the pre-2021 policy; instead, we train the algorithm to recognize prescriptive proposals based on the SEC's decisions on hundreds of contested proposals.

We use proposals contested under the "ordinary business exclusion" prior to 2021 as a training dataset for Google's BERT (Bidirectional Encoder Representations from Transformers) algorithm (Liu and Lapata (2019)).[4] Assuming that the SEC is more likely to exclude highly prescriptive

[4] As noted in Sections 4.2 and 9.2, a portion of these proposals is reserved as a validation dataset to test the model on out-of-sample data that it has not encountered before.

proposals under the "ordinary business exclusion," we then use this algorithm to classify all uncontested and withdrawn proposals from 2018 to 2021, as well as all proposals from 2022 and 2023. Furthermore, recognizing that this classification might not capture every prescriptive proposal, in a secondary step, we adopt an unsupervised "Topic Modeling" strategy (Grootendorst (2022)) to identify additional prescriptive proposals. This approach aims to uncover clusters of proposals potentially associated with prescriptive textual elements, such as the request for the adoption of specific policies.

After adjusting for a wide range of characteristics, our analysis reveals that prescriptive proposals tend to attract lower levels of shareholder support. More pertinently, by exploiting the regulatory shock created by the 2021 Guidance to induce quasi-exogenous changes in proposal prescriptiveness, we find a marked decline in support for prescriptive proposals post-2021, relative to their non-prescriptive counterparts.

To determine whether the decline in support for E&S proposals might be influenced by specific shareholder groups known for their ideological stances on E&S issues, we investigate the relationship between individual fund-level voting behavior and the prescriptiveness of proposals. We find that mutual funds are, on average, less likely to support prescriptive proposals across the entire period, with this tendency becoming even more pronounced after 2021.

More importantly, we find evidence of a decline in support for prescriptive proposals among institutional shareholders with varying ideological preferences on E&S issues. Following Bolton et al. (2020) and Michaely et al. (2021), we construct an ideological spectrum for funds on E&S issues, with "pro-social" funds at one end and "financially-oriented" funds at the other. While support for prescriptive proposals has generally decreased across the large majority of funds since 2021, funds with stronger preferences for E&S issues are more likely to support these proposals, whereas funds with stronger financial preferences are more likely to oppose them. For example, within the set of ESG funds, those belonging to E&S-focused families (which lean pro-social) are more likely to support prescriptive proposals compared to the average fund. In contrast, ESG funds in non-E&S families (which lean financially-oriented) show no significant difference from the average fund. Additionally, active mutual funds, which are typically more financially-oriented, are more likely to oppose prescriptive proposals. Conversely, the "Big Three" funds and other predominantly passive mutual funds, which align closer to the median voter on the E&S ideological spectrum, are indistinguishable from the average fund.

Our results are consistent with the hypothesis that many institutional investors do not "walk the talk" when E&S concerns conflict with pecuniary maximization objectives (Goshen and Hamdani

(2023); Michaely et al. (2021); Li et al. (2023b); Heath et al. (2021); Aggarwal et al. (2023)). If prescriptive proposals generally entail a higher level of commitment to E&S issues than their non-prescriptive counterparts—as suggested by the ideological voting patterns—the evidence implies that most investors are reluctant to embrace commitments that impose significant financial costs. Although investors still vote along ideological lines on E&S matters, and those with weaker E&S preferences are less likely to support prescriptive proposals, the financial costs of these proposals may outweigh the intensity of E&S preferences for most investors.

Although prescriptive proposals could, in principle, differ from non-prescriptive proposals in ways that make them less attractive to investors regardless of the intensity of their E&S preferences—e.g., poorer drafting quality or ill-advised recommendations—we regard our interpretation as more convincing. In practice, heightened prescriptiveness typically involves adding concrete goals and implementation timelines, often drawing on guidance from expert third parties or international bodies (Tallarita (2022)). A common shift, for example, replaces a general request to "reduce carbon emissions" with a specific request to align with Paris-Agreement trajectories or peer-company Net-Zero pledges. These refinements typically heighten the level of E&S commitments requested from the firm while leaving the proposals' basic focus unchanged.

Our article is organized as follows. In Section 2, we briefly review the existing literature and explain how our study contributes to it. In Section 3, we outline the regulatory context surrounding Rule 14a-8 and describe our data sources. In Section 4, we explain how we measure the precriptiveness of shareholder proposals and present our findings on the voting support for prescriptive and non-prescriptive E&S proposals among shareholders. In Section 5, we combine proposal-level data from Section 4 with individual fund-level voting information and present the relevant findings. Section 6 provides robustness tests concerning a key threat to our identification strategy—the presence of political backlash. Section 7 concludes. Finally, an Online Appendix (Section 9) details additional results secondary to our primary analysis, the machine learning techniques used to develop the prescriptiveness indicator referenced in Sections 4 and 5, the data-cleaning procedures employed for the findings presented in Section 5, and includes additional tables and figures that support our primary analysis.

2 Contributions to the Literature

Our article contributes to two distinct debates on E&S proposals. The first debate concerns the potential relationship between the prescriptive nature of shareholder proposals and shareholder support. Although many commentators—including journalists (WSJ (2022)), legal practitioners (Posner (2022); Gibson-Dunn (2022); Gibson-Dunn (2023); Morgan (2024)), legal academics (Tallarita (2022); Fisch and Robertson (2023)), and even major institutional investors (Blackrock (2022))—have suggested that the "prescriptiveness" of shareholder proposals may influence voting outcomes, we are not aware of any study that formally investigates whether (and to what extent) such a relationship exists.

Establishing how prescriptive proposals affect voting behavior or other corporate outcomes presents challenges in two main respects. One is that the textual composition of a proposal is inherently shaped by its proponents, who may strategically craft their wording to sway voting outcomes or avoid SEC preclusion (Gantchev and Giannetti (2021); Tallarita (2022)). Thus, any quasi-exogenous variation in "prescriptiveness" must arise from a significant shift (regulatory or otherwise) that alters proponents' incentives—a context exemplified by the 2021 change in the SEC's Rule 14a-8 policy, which we exploit. This shift could increase the prescriptiveness of proposals in at least two ways: it may encourage new proponents to submit more prescriptive proposals, and it could also allow existing proponents, who were previously constrained, to align their proposals more closely with their actual preferences.

Furthermore, the textual contents of proposals are intrinsically unstructured and high-dimensional compared to traditional quantitative measures used in causal inference (Egami et al. (2022)), making it difficult to devise a "prescriptiveness" metric free from subjective biases. To address these issues, we employ recent advances in corporate governance research that incorporate machine learning methods, such as embedding models and dimensionality reduction (Michaely et al. (2023); Rajan et al. (2023); Li et al. (2023a); Briscoe-Tran (2023); Andrikogiannopoulou et al. (2022)). A key contribution we make to the literature is using legal outcomes to label data for supervised machine learning, rather than relying on researcher-coded labeling (Badawi (2023); Gompers et al. (2003); Frankenreiter et al. (2021); Porta et al. (1998); Spamann (2010)). Indeed, our objective is not to construct an expert-driven, "objective" measure of prescriptiveness as advocated by Bainbridge (2016), but rather to replicate the SEC's own interpretation of prescriptiveness under Rule 14a-8(i)(7).

The second debate is about the determinants of shareholder voting support.[5] The factors affect-

[5]A separate strand of literature examines the economic impact of shareholder proposals that proceed to a vote. Some studies suggest that mutual fund support for E&S proposals represents informative signals about firms' potential E&S risks, highlighting the potential advantages of expanding shareholder voice on E&S issues (He et al. (2023)). However, other research presents evidence that investors respond favorably when the SEC allows for the exclusion of certain proposals, suggesting that, on average, these excluded proposals were perceived as detracting from firm value (Matsusaka et al. (2019); Matsusaka et al. (2021)). Our work is adjacent to this line of literature.

ing voting support in corporate decisions are numerous and varied, encompassing firm characteristics (Cuñat et al. (2012)), shareholder characteristics (Brav et al. (2024); Brav et al. (2022)), proponent characteristics (Gantchev and Giannetti (2021); Bebchuk et al. (2020)), proposal topics (Bolton et al. (2020); Bubb and Catan (2022); Curtis et al. (2021)), the strategic incentives of voting investors (Michaely et al. (2021); Li et al. (2023b)), proxy advisory recommendations (Iliev and Lowry (2015); Iliev and Vitanova (2022); Hu et al. (2024)), and the dynamics of management-shareholder relations (Matvos and Ostrovsky (2010)). Our key contribution to this literature underscores the interplay between investor ideology and the *regulatory landscape* in determining voting outcomes, an aspect that has not been fully explored by other authors.

As discussed in Section 1, we consider the 2021 Guidance to be a quasi-exogenous shock that increased the "prescriptiveness" of proposals.[6] We hypothesize that prescriptive proposals exhibit a higher level of commitment to E&S issues but are more costly to implement. Therefore, a quasi-exogenous increase in prescriptiveness should illuminate how investors balance pro-social and financial goals in their voting decisions.

Prior literature has shown that investor ideology is a primary determinant of voting behavior on E&S issues (Bolton et al. (2020); Michaely et al. (2021); Dikolli et al. (2022); Curtis et al. (2021)). These empirical findings align with a broader theoretical literature suggesting that investor behavior may involve balancing pro-social and pecuniary objectives (Hart and Zingales (2017); Hart and Zingales (2022); Broccardo et al. (2022); Barber et al. (2021); Hirst et al. (2023); Hart et al. (2024)). Collectively, this scholarship supports our hypothesis that, *ceteris paribus*, an increase in proposal prescriptiveness should induce greater voting support among pro-social investors and reduced support among financially-oriented investors, with the average fund influencing the overall outcome. Overall, our findings suggest that for most investors, the financial costs implicit in more prescriptive proposals may outweigh their pro-social appeal, even though individual voting responses continue to vary along the E&S ideological spectrum.

[6]We can also view this quasi-exogenous shock as an increase in the "intensity" of E&S issues to be voted on.

3 Legal Framework and Data

3.1 Rule 14a-8 and the 2021 Guidance

Rule 14a-8 requires public companies to include shareholder proposals that meet certain formal and substantive criteria in the proxy materials circulated to shareholders.[7] Since most shareholders do not attend the annual meeting in person and vote instead by proxy, inclusion in the proxy materials is effectively the only means by which these proposals can be presented to, and voted on by, other shareholders. Consequently, the submission of shareholder proposals under Rule 14a-8 is often described as a form of "low-cost activism." (Kastiel and Nili (2020); Gantchev and Giannetti (2021); Bainbridge (2016))

However, Rule 14a-8(i) allows companies to omit a shareholder proposal from the proxy statement if it fails to meet certain conditions. For instance, companies may exclude proposals that address the company's "ordinary business operations," are "materially false or misleading," or have already been implemented. Here, we focus on Rule 14a-8(i)(7), which permits the exclusion of proposals "deal[ing] with a matter relating to the company's ordinary business operations."

Management may seek to exclude a shareholder proposal by submitting a "no-action letter" request to the SEC, outlining its basis for exclusion. This adversarial process is akin to litigation, allowing the proponent to respond before the SEC staff issues a decision to side with either the company or the proponent. Nevertheless, even when the SEC allows the proposal to proceed to a vote, management nearly always recommends that shareholders vote against these proposals (Tallarita (2022)).

Over the years, the SEC and its Corporation Finance Division have issued several interpretive documents to clarify which proposals fall within the "ordinary business exclusion." Beginning with a sequence of Staff Bulletins in November 2017, the SEC staff indicated that companies could exclude "social policy" proposals that tried to "micromanage" the company by including "the imposition or assumption of specific time-frames or methods for implementing complex policies" (SEC (2017)).

According to this interpretation, proposals seeking annual reporting on "short-, medium- and long-term greenhouse gas targets aligned with the greenhouse gas reduction goals established by the Paris Climate Agreement to keep the increase in global average temperature to well below 2 degrees Celsius and to pursue efforts to limit the increase to 1.5 degrees Celsius" is "overly

[7] 17 C.F.R. § 240.14a-8.

prescriptive" as it tries to micromanage the company. By constrats, proposals that "defers to management's discretion to consider if and how the company plans to reduce its carbon footprint" are acceptable (SEC (2017)).

However, in Staff Legal Bulletin No. 14L, published in November 2021, the SEC rescinded these previous documents and reversed its prior position. Contrary to its earlier guidance, the SEC announced that "social policy proposals" "seeking detail or seeking to promote timeframes or methods do not per se constitute micromanagement." To illustrate this policy shift, the SEC noted that proposals requesting that "companies adopt timeframes or targets to address climate change" would henceforth be considered non-excludable (SEC (2021)).[8]

In the terminology used by the SEC and by industry and policy experts—and adopted in this Article—the SEC's 2021 Guidance opened the door to more "prescriptive" E&S proposals. Table 2 presents three illustrative pairs of prescriptive proposals that are nearly identical in substance. In each pair, the first was excluded under the old regime, whereas the second was allowed onto the ballot after the 2021 Guidance. For example, before the 2021 Guidance, the Commission granted no-action relief to omit a proposal asking the board to adopt policies that would prevent its underwriting practices from supporting new fossil-fuel projects inconsistent with the IEA's Net Zero Emissions by 2050 pathway. Once the 2021 Guidance was in effect, however, a nearly identical proposal, which asked the board to adopt a policy by the end of 2022 prohibiting the financing of new fossil-fuel supplies misaligned with the same pathway, was permitted to proceed.

Many observers of the 2021 Guidance highlighted this as a significant departure in the SEC's approach, describing it as a "clear move by the SEC to encourage sustainability efforts." (Era et al. (2021)). From a practical standpoint, legal practitioners also suggested that the guidance created a more difficult threshold for no-action relief, and would likely result in more E&S shareholder proposals either making it onto the agenda of shareholder meetings or ending in a settlement for the company (Era et al. (2021)). Our primary strategy in this Article is to exploit this policy change to study the effects of "prescriptiveness" on shareholder support for E&S proposals.

[8]The SEC's position on Rule 14a-8(i)(5), the economic relevance exception, was also revised in Staff Bulletin No. 14L. Henceforth, we will refer to this staff bulletin as the "2021 Guidance".

3.2 Data Sources

We procure our data from multiple sources. Our primary dataset is Factset,[9] which provides information on all environmental, social, and governance (ESG) shareholder proposals at Russell 3000 companies from 2018 to 2023.[10] Although data on shareholder proposals extends beyond 2018, we restrict our scope to post-2018 data for three reasons: (1) to focus on the causal impact of the 2021 Guidance, which rescinded the Staff Bulletins issued in November 2017; (2) to limit the potential influence of other confounding events, such as the 2015 *Trinity Wall Street v. Wal-Mart Stores, Inc.* case; [11] and (3) to keep the analysis tractable when merging proposal data with extensive fund-level voting data.[12]

The Factset dataset also captures several proposal characteristics, including whether a no-action letter was requested, whether it was granted, and the specific sub-topics of the proposal. In addition, it offers limited information on the proponent's characteristics, such as proponent type (e.g., a pension fund or an individual) and the proponent's name.[13] Since voting outcomes and proponent targets often hinge on firm-level attributes (Cuñat et al. (2012); Bebchuk et al. (2020)), we gather firm characteristics from the CRSP-Compustat-Merged (CCM) database and merge these data with the Factset dataset at the firm-year level.[14]

In Section 5 of our article, we investigate the relationship between individual fund-level voting behavior and the prescriptiveness of proposals. Because the Factset database does not include data on fund-level voting, we begin by merging the Factset dataset with the ISS Voting Analytics (Company Vote Results) database, which provides proxy voting outcomes at the firm-year level. In the absence of a direct common identifier, we match observations using a firm identifier, the relevant meeting date, and aggregate votes (for, against, and abstentions). We then integrate this combined dataset with the ISS Voting Analytics (Mutual Fund Vote Records) database using the unique identifiers assigned by ISS. This step allows us to obtain detailed voting results at the individual fund level for each firm.

Similar to firm characteristics, prior research has shown that voting outcomes are also influenced

[9] While most scholars like He et al. (2023) and Gantchev and Giannetti (2021) have traditionally used a similar dataset from Institutional Shareholder Services (ISS) as their primary source, we have chosen Factset because, unlike ISS, it includes data on the textual content of shareholder proposals.

[10] In Figures 1 and Table A12, we also utilize data on the same set of firms from 2013 to 2018. However, we exclude this data from subsequent analysis.

[11] *Trinity Wall Street v. Wal-Mart Stores, Inc.*, 792 F.3d 323 (3d Cir. 2015)

[12] For instance, although Zytnick (2022) focuses solely on E&S proposals from 2015 to 2017, incorporating individual fund-level voting data still yields a dataset with nearly five million observations.

[13] Further information about these variables can be found in Table A1.

[14] Further information about the variables collected from this dataset can be found in Table A1.

by fund-specific attributes (Brav et al. (2024); Brav et al. (2022)). To integrate these attributes for each fund in our dataset, we merge the previously mentioned data with the CRSP mutual fund database.[15] Because neither dataset contains common identifiers, we obtain the fund names linked to each N-PX identifier (as reported in the ISS dataset) from the SEC's EDGAR database.[16] We then employ fuzzy-matching techniques to align these EDGAR fund names with the corresponding entries in the CRSP mutual fund database. Finally, we merge the combined datasets using a fund identifier, a firm identifier, and relevant record dates.[17]

4 Prescriptive Proposals and Shareholder Support

4.1 The Decline in Shareholder Support for E&S Proposals

After a significant increase over the past decade, voting support for E&S proposals has seen a notable downturn from 2022 to 2023. In Figure 1, we show that average support for E&S proposals (across various metrics) steadily increased from 2018 through 2021 before dropping after 2021. Although governance proposals—focusing on takeover defenses, independent directors, and shareholder rights—also experienced a reduction in support post-2021 (Eldar and Wittry (2021)), the magnitude of this decline was considerably more muted.[18] Table 1 offers additional context and summary statistics for our key variables. On average, governance proposals garner 34.00% support, compared to 28.33% for environmental and 25.06% for social proposals. Overall, governance issues dominate, making up 53.97% of all shareholder proposals.

Our objective is to pinpoint the mechanism behind this marked shift. We hypothesize that the 2021 Guidance triggered the emergence of more prescriptive E&S proposals, which were ultimately disfavored by the majority of investors. While anecdotal evidence suggests this mechanism may be plausible, formally testing our hypothesis requires a clear quantitative measure of a proposal's "prescriptiveness" (Ilhan et al. (2023)).

[15]We also incorporate fund characteristics from the Thomson Reuters S12 Mutual Fund database. For further details on our data-cleaning procedures, see Section 9.5.

[16]Mutual funds and other registered management investment companies must disclose proxy votes pursuant to Section 30 of the Investment Company Act of 1940 and Sections 13 and 15(d) of the Securities Exchange Act of 1934. These disclosures, referred to as "Form N-PX" disclosures, connect each fund name in the SEC's EDGAR database with a non-unique N-PX identifier.

[17]A more detailed description of these data-cleaning procedures is provided in Section 9.5. Further information about the variables collected from this dataset can be found in Table A1.

[18]In the Appendix, we formally examine this trend using the Synthetic Difference-in-Differences (SDID) methodology developed by Arkhangelsky et al. (2021). See Section 9.1.

4.2 Constructing a Measure for Proposal Prescriptiveness

4.2.1 Supervised Model

We construct our measure of proposal "prescriptiveness" using a combination of supervised and unsupervised machine learning methods in Natural Language Processing (NLP). The supervised approach replicates the SEC's own assessment of a proposal's prescriptiveness, drawing on hundreds of contested proposals where the SEC has rendered decisions on the applicability of the micromanagement exclusion under Rule 14a-8(i)(7). As explained in Section 3.1, when a company challenges an E&S proposal for "micromanaging" the company (i.e., for being overly prescriptive), the SEC staff adjudicates the dispute and sides with either the company or the proponent.[19]

We assume that, through 2021, consistent with the legal guidance then in effect, the SEC tended to exclude proposals displaying a higher degree of prescriptiveness. Accordingly, when examining all proposals from 2001 to 2021 (i.e., before the 2021 Guidance) that were contested under Rule 14a-8(i)(7), we assign a prescriptiveness indicator of 1 to proposals that were excluded and 0 to those that proceeded to a vote. This approach yields a "training" set of 927 proposals and a "validation" set of 231 proposals. We use the training set to train Google's BERT model (Bidirectional Encoder Representations from Transformers) to distinguish prescriptive proposals from non-prescriptive proposals, and the validation set to to evaluate the model's predictive accuracy on out-of-sample data. Our validation tests show that BERT is able to predict SEC's decision with about 74% accuracy, indicating that BERT is able to replicate the SEC's assessment of excessive "prescriptiveness" under the old policy in a substantial majority of cases.

Then, we use BERT to classify all other E&S proposals in our dataset—including those that were uncontested or withdrawn between 2018 and 2021, as well as all the proposals from 2022 to 2023—as prescriptive or non-prescriptive.[20]

[19]Consistent with the findings of Tallarita (2022) and Matsusaka et al. (2021), a substantial majority (61.7%) of all contested proposals in our dataset are disputed on the grounds that they would interfere with a company's ordinary business operations, thus qualifying for exclusion under Rule 14a-8(i)(7).

[20]Contested proposals from 2018 to 2021 are not classified because they form part of the training set. This approach aligns with methodologies in Michaely et al. (2023), Rajan et al. (2023), and Liu and Lapata (2019). The BERT model is pre-trained on approximately 3.2 billion words from Wikipedia and on 11,000 books from various genres, enabling it to generate context-specific embeddings (i.e., numerical weights assigned to words) (Liu and Lapata (2019)). For a detailed explanation of how we implement these machine learning algorithms, see Section 9.2.

4.2.2 Unsupervised Model

Although the supervised approach is quite accurate in identifying prescriptive proposals, it may not capture the full spectrum of such proposals. This limitation arises because the training dataset is drawn from proposals clustered near the SEC's decision boundary, which can misclassify proposals whose language or context lies far outside that training manifold (i.e., the surface in embedding-space traced out by the training points in machine learning).[21]

To address this concern, we also employ an unsupervised "Topic Modeling" strategy (Grootendorst (2022)) to identify groups of proposals sharing common themes linked to "prescriptive content," such as the implementation of specific policies. This approach provides a more nuanced perspective on the proposals' characteristics. Given the likely differences in content between environmental and social proposals, we run separate topic modeling analyses for each category.

Initially, we apply an embedding model to assign context-specific weights to individual words (or word combinations) in our dataset. Next, we use the UMAP (Uniform Manifold Approximation and Projection) algorithm to reduce the dimensionality of the textual data, retaining the most important features of each environmental or social proposal. We then employ a vectorization model to filter out common stop-words in these proposals.[22] Finally, we use the HDBSCAN (Hierarchical Density-Based Spatial Clustering of Applications with Noise) algorithm to group similar proposals, thereby identifying distinct clusters.[23]

The application of topic modeling algorithms to our dataset reveals distinct topic clusters aligning with characteristics of prescriptive proposals noted by the SEC, legal practitioners, and institutional investors (Era et al. (2021); Blackrock (2022); WSJ (2022); Gibson-Dunn (2022); Gibson-Dunn (2023); Posner (2022)). For instance, among environmental proposals, a clear cluster emerges that urges companies to set "time-bound" emissions targets. Another set of proposals calls for companies to "adopt a [specific] policy" (or similar phrases like "implementing," "adopting," or "committing to a policy"), such as phasing out fossil fuel exploration and development. We identify these clusters of proposals as ostensibly "prescriptive" in nature, and assign a prescriptiveness indicator of 1 (Li et al. (2023a)).

[21]Moreover, the relatively small size of the training dataset increases the risk of misclassification. We provide a formal discussion of these limitations in Section 9.3.

[22]Stop-words are frequently used words in a language (e.g., "the," "is," "and") that typically carry little analytical significance.

[23]A more detailed description of how we implement these machine learning algorithms is provided in Section 9.3.

4.3 Voting Support for Prescriptive Proposals

4.3.1 Growth in Prescriptive Proposals

Table 3 displays the number and percentage of prescriptive proposals that proceed to a vote (Panel A) compared to those that do not (Panel B) over the entire 2018–2023 sample period. As Table 3 shows, there is a clear uptick in both the number and proportion of prescriptive proposals that make it to the ballot after 2021, rising from 49 in 2021 to 113 in 2022, and then to 142 in 2023. More importantly, even as the total volume of E&S proposals (prescriptive and non-prescriptive alike) continues to grow, the share of prescriptive proposals increases from 31.61% in 2021 to 40.36% in 2022, and then to 46.71% in 2023. In contrast, there is a notable decline in the fraction of prescriptive proposals excluded by the SEC, and thus prevented from going to a vote. Specifically, 63.64% of these proposals were excluded in 2021, dropping to 56.00% in 2022, and again to 49.23% in 2023. Figure 2 graphically depicts these developments. Overall, these trends support the hypothesis that the 2021 Guidance contributed to an increase in the number of prescriptive proposals reaching a vote.[24]

4.3.2 Investor Support for Prescriptive Proposals after the 2021 Guidance

As an initial test of shareholder support for prescriptive proposals, we run panel regressions to estimate the relation between prescriptiveness and voting support, both before and after the 2021 Guidance, based on the following specification:

$$y_{ijktn} = \alpha + p_{ijktn}\beta + X_{ijktn}\xi + \theta_i + \eta_j + \psi_k + \nu_t + \varepsilon_{ijktn} \quad (1)$$

where i indexes firms, j indexes industries, k indexes proponent-types, n indexes proposals, X is a vector of firm-proposal controls[25], while θ_i, η_j, λ_k, and ν_t represent firm, industry, proponent-type, and year fixed effects, respectively. Meanwhile, y_{ijktn} relates to a measure of voting support–in this case, the percentage of affirmative votes out of the total votes cast, while the binary indicator p_{ijktn} denotes whether a proposal is prescriptive in line with Section 4.2.

Tables 4 and A2 report results from several variations of this specification, incorporating various

[24] However, as shown in Section 4.3.2, the rise in prescriptive proposals does not appear to be the key driver behind the decline in voting support for E&S proposals.

[25] Further information about these variables can be found in Table A1.

fixed effects, sub-samples, dependent variables, and selection bias models.[26] Our findings show that prescriptiveness is associated with much lower voting support among shareholders. On average, shareholder voting support is 28.33% for environmental proposals and 25.06% for social proposals (see 1). However, prescriptiveness is associated with a reduction in voting support of 5.22 percentage points—11.16 for environmental proposals and 4.74 for social proposals.

Next, to estimate the impact of the 2021 Guidance, we revise the baseline specification 1 by introducing the interaction term $p_{ijktn} \times Post_t$, where $Post_t$ is a binary indicator that denotes whether the proposal is post-treatment (i.e., in 2022 or 2023).[27]

$$y_{ijktn} = \alpha + p_{ijktn}\beta + (p_{ijktn} \times Post_t)\gamma + X_{ijktn}\xi + \theta_i + \eta_j + \psi_k + \nu_t + \varepsilon_{ijktn} \quad (2)$$

Table 5 presents our findings. We vary the specifications by altering the measure of voting support, by distinguishing whether proposals address environmental, social, or both categories, and by using different fixed effects. For instance, specifications (1), (3), (5), and (6) employ firm fixed effects, whereas specifications (2) and (4) use industry fixed effects.[28] In our main specification (1), which includes firm fixed effects, the interaction term $p_{ijkt} \times Post_t$ has a negative coefficient that is statistically significant at the 1% level, indicating a substantial decrease in shareholder support for prescriptive proposals following the 2021 Guidance. Specifically, prescriptive E&S proposals received 8.48% less support compared to their non-prescriptive counterparts after 2021. Our primary results remain robust across the various specifications. However, the coefficient on p_{ijkt} alone remains statistically indistinguishable from zero in all models. This pattern suggests that the decline in voting support is driven primarily by changes in the degree of prescriptiveness of proposals after 2021, rather than any differences in prescriptiveness before the 2021 Guidance.

To examine how support for prescriptive proposals has evolved over time in light of the 2021 Guidance, we estimate coefficients for interaction terms that combine year-specific indicators with a binary variable distinguishing prescriptive from non-prescriptive proposals.[29] These coef-

[26] Sections 4.3.3 and 9.4 provide a detailed discussion of how potential selection biases are addressed.

[27] Because year-fixed effects are included, the term $Post_t$ is perfectly collinear with these fixed effects and is therefore excluded from the regression. Furthermore, as our design does not incorporate staggered treatments, we do not apply the recent methodological innovations detailed in Baker et al. (2022).

[28] Industry fixed effects amplify the treatment coefficients, whereas firm fixed effects—by stripping out unobserved, time-invariant firm traits—attenuate them. We report the more conservative firm-fixed-effect estimates to demonstrate robustness; see Section 9.7.

[29] In estimating these coefficients, we include firm-proposal controls, along with fixed effects for firm, year, and proponent type. We use 2021 as the baseline year (when the treatment occurred), and the dependent variable is the percentage of affirmative votes out of the total votes cast.

ficients are displayed over time in Figure 3.[30] The figure indicates a clear decline in the estimated coefficients for all E&S proposals after 2021, with environmental proposals showing a particularly sharp reduction compared to social proposals.

A plausible explanation for these findings is that prescriptive proposals became considerably *more* prescriptive after 2021, both in terms of their overall volume (see Table 3) and substantive content.[31] Prior to 2021, proponents likely moderated these proposals to stay within the boundaries set by the SEC, presumably to advance E&S objectives without risking exclusion under Rule 14a-8. Once the 2021 Guidance took effect, however, proponents leveraged this new flexibility by introducing new proposals with significantly more prescriptive elements or revising existing proposals to be more prescriptive. As a result, while investors were largely indifferent between prescriptive and non-prescriptive proposals before 2021 (as evidenced by the coefficient on p_{ijkt} being statistically indistinguishable from zero), they have a strong negative response to the increasingly prescriptive proposals under the revised SEC policy.[32]

In Table A3, we further corroborate this interpretation by using a continuous measure of prescriptiveness derived from the raw probability values generated by our supervised algorithms.[33] Although the coefficients in Table A3 are smaller than those in Table 5, they still indicate a negative relationship between the post-2021 interaction term ($p_{ijkt} \times Post_t$) and voting support, alongside no significant difference in investor response prior to 2021.

4.3.3 Addressing Potential Selection Effects

To establish the causal impact of the 2021 Guidance, the effect of this regulatory shock on voting outcomes must arise solely through its influence on proposal prescriptiveness. In other words, the shock should neither directly alter voting outcomes nor do so indirectly through mechanisms unrelated to prescriptiveness. While ruling out direct effects is relatively straightforward, the possibility remains that indirect pathways—such as selection effects—could play a role.

One potential source of selection bias arises from endogenous or non-random inclusion in the sample. For instance, the 2021 Guidance could prompt corporate management, especially at larger firms facing heightened reputational risks, to refrain from contesting proposals (Bebchuk et al. (2020)). If managers shift their behavior for reasons unrelated to prescriptiveness, any observed

[30] Panel A of Figure 3 illustrates trends for all E&S proposals, while Panels B and C depict trends specifically for environmental and social proposals, respectively.

[31] In other words, the rise in prescriptiveness reflects an "extensive margin" and an "intensive margin" effect.

[32] We explore the "extensive" and "intensive" effects in greater detail in Section 4.3.4.

[33] We apply a log transformation to these probability values to address skewness in the distribution.

changes in voting outcomes may not be attributable solely to the 2021 Guidance's effect on proposal content. To address this potential bias, we employ a Heckman selection model (Heckman (1979)), which corrects for the selective exclusion or withdrawal of proposals (Zytnick (2022); Brav et al. (2024)).[34]

Another form of potential selection bias involves how proposals are "selected" for treatment, which, in this context, relates to a proposal's prescriptiveness. For example, the 2021 Guidance might encourage proponents to direct more prescriptive proposals towards larger firms, believing these E&S proposals will have a higher likelihood of proceeding to a vote post-2021 (Era et al. (2021); Bebchuk et al. (2020)). To ameliorate these concerns, we calculate propensity scores for prescriptive (treatment) and non-prescriptive (control) proposals, representing each proposal's likelihood of receiving the "treatment" based on an array of observable characteristics. Incorporating these scores into the analysis helps ensure that the two groups differ only in their level of "prescriptiveness," minimizing systematic differences apart from the treatment.[35]

In Table A4, we revisit the specifications from columns (1), (3), and (4) of Table 5, applying the previously described corrections for potential selection bias. The results indicate that the key coefficients (specifically, on $p_{ijkt} \times Post_t$) remain consistent with those in Table 5, suggesting that selection bias based on observable characteristics is unlikely to explain the observed treatment effects.

4.3.4 Anti-ESG Proposals, New Proponents, and New Targets

Other potential confounders in our analysis might arise from changes in the ideological nature of E&S proposals (e.g., a rise in so-called "anti-ESG" proposals), changes in the identity of proponents, or changes in the identity of target companies, following the 2021 Guidance. To address the first concern, Table A5 replicates the analyses in Table 5 while excluding the (small) subset of anti-ESG proposals, which comprise roughly 7.37% of our sample. The results in Table A5 confirm that our main findings remain robust despite the removal of these proposals.[36]

To address the second concern, we investigate two distinct hypotheses. One posits that the

[34]Note that shareholder proposals must be contested by firm management before being excluded by the SEC. Bebchuk et al. (2020) describe numerous firm and proponent characteristics that may influence whether a proposal is contested, including the activist's stake, insider ownership, share class structure, performance, historical success rates, and board composition.

[35]Further details about these models are provided in Section 9.4.

[36]We discuss anti-ESG proposals in greater detail in Section 6.3. Although they are included in our primary specifications—given that our main variable of interest, prescriptiveness, is correlated with them—our results still hold when these proposals are excluded, as demonstrated in Table A5, Table A11, and throughout Section 5.

post-2021 decline in voting support largely reflects existing proponents altering the prescriptiveness of their proposals. Another posits that the decline stems primarily from new proponents—previously deterred by the old policy—now submitting prescriptive proposals. This latter scenario may also explain a decrease in E&S proposal quality, potentially due to insufficient expertise or sophistication among newer proponents (Gantchev and Giannetti (2021)).

We investigate these hypotheses by modifying specification (2) to include proponent fixed effects (as opposed to proponent-type fixed effects), so that all variation is limited to within-proponent variation over time. In a related specification, we incorporate $pshare_{kt}$, the proportion of prescriptive proposals submitted by each proponent in a given year. Finally, we introduce an additional binary variable, $FirstAppearance_n$, into the key interaction term, $p_{ijkt} \times Post_t \times FirstAppearance_n$. We define $FirstAppearance_n$ to take on the value 1 when a new proponent name is first observed for a given proposal, and 0 otherwise.

Table 6 presents these results. Column (1) reproduces the baseline specification from column (1) of Table 5. In columns (2) and (3), we replace proponent-type fixed effects with proponent fixed effects, thereby removing any variation between different proponents.[37] Although the coefficient on $p_{ijkt} \times Post_t$ is smaller than in the baseline, it remains negative and statistically significant at the 1% level, implying that prescriptive E&S proposals received 6.59% less support post-2021 compared to non-prescriptive proposals. In column (4), we include we include $pshare_{kt}$, but while the coefficient on $p_{ijkt} \times Post_t$ remains similar to column (2), the coefficient on $pshare_{kt}$ is not statistically significant. In columns (5) and (6), we add the binary variable $FirstAppearance_n$, effectively conducting a triple difference-in-difference analysis under firm and industry fixed effects, respectively. However, the coefficient on $p_{ijkt} \times Post_t \times FirstAppearance_n$ is not statistically different from zero.

Collectively, these findings suggest that the post-2021 decline in voting support aligns more closely with the hypothesis that existing proponents are making their proposals more prescriptive in response to the 2021 Guidance, even after accounting for the overall rise in the proportion of prescriptive proposals (see Figure 2).[38] This finding is consistent with Tallarita (2022), who observes that the shareholder proposal market is dominated by a relatively small number of specialized actors who connect shareholders with pro-social motives with corporate stakeholders, citizens, and social and policy activists.

Finally, to address the third concern (i.e., a potential shift in the identity of target companies),

[37] Column (2) is our main specification here, while column (3) replaces firm fixed effects with industry fixed effects.

[38] In other words, our evidence indicates that the "intensive margin" largely explains the changes in voting support. In fact, the average proportion of prescriptive proposals only rose by about 7.5% after the 2021 Guidance.

we focus on the subset of "stable firms" that appear in our sample both before and after the 2021 Guidance. Figure 4 indicates that these stable firms are the primary targets of shareholder proponents, accounting for an average of 78.34% of all proposals in our dataset. This observation mirrors Tallarita (2022), who notes that specialized actors in the E&S proposals market tend to concentrate on large firms perceived to have a substantial social impact. The figure displays the yearly distribution of E&S proposals submitted to stable firms—70.98% in 2019 and 88.68% in 2022—underscoring their central role throughout the sample period. Given the dominance of stable firms in the dataset, firm-level sample selection is unlikely to pose a material threat to the validity of our main regression estimates.[39] Moreover, Table A6 shows that our main findings remain robust even when non-stable firms are excluded from the analysis.

5 Prescriptive Proposals and Investor Characteristics

5.1 Mutual Fund Voting on Prescriptive Proposals

As discussed in Section 1, numerous studies have shown that different mutual funds often vote differently, especially on E&S issues (Bolton et al. (2020); Curtis et al. (2021); Bebchuk and Hirst (2019); Bebchuk and Hirst (2022); Griffin (2020); Zytnick (2022)). In light of this, we examine whether the decrease in support for prescriptive proposals described in Section 4.3 is driven by particular shareholder groups. To that end, we merge proposal-level data with mutual fund-level voting information, capturing over 900,000 individual fund votes on E&S proposals.[40]

Table 7 presents summary statistics on our fund-level data.[41] Our primary dependent variable of interest, "Binary Fund Vote," is coded as 1 if a specific fund votes in favor of a proposal and 0 otherwise. Consistent with Brav et al. (2024), another dependent variable, "Ordered Fund Vote," takes a value of 1 for a "yes" vote, 0.5 for an abstention, and 0 for any other outcome.

To examine how fund voting support varies with the prescriptiveness of proposals, we estimate

[39]Figure 4 shows that most firms received multiple shareholder proposals during the sample period, justifying the use of firm fixed effects.

[40]Integrating proposal-level and fund-level voting data is challenging due to inconsistencies and gaps in the databases, including missing entries, which necessitates excluding many unmatched records. Moreover, voting information is only available for mutual funds subject to N-PX filing requirements under Section 30 of the Investment Company Act of 1940; thus, pension funds, banks, and retail investors are not obliged to disclose their votes, resulting in the exclusion of substantial information from the merged dataset.

[41]Further information about these variables can be found in Table A1.

the following panel regressions:[42]

$$y_{ijktnm} = \alpha + p_{ijktnm}\beta + X_{ijktn}\xi + V_{ijktnm}\delta + \theta_i + \eta_j + \psi_k + \kappa_m + \nu_t + \varepsilon_{ijktnm} \quad (3)$$

where i indexes firms, j indexes industries, k indexes proponent-types[43], m indexes funds, n indexes proposals, X is a vector of firm-proposal controls[44], V is a vector of fund-level controls[45], while θ_i, η_j, ψ_k, κ_m, and ν_t represent firm, industry, proponent-type, fund, and year fixed effects, respectively. Meanwhile, y_{ijktnm} relates to a measure of voting support (e.g., the "Binary Fund Vote" measure described above), while p_{ijktnm} is a measure of prescriptiveness that denotes whether a given proposal is prescriptive or not, in line with Section 4.2.

Table 8 presents our results from specification (3). In our main specification (1) with firm fixed effects, fund-level voting support for prescriptive proposals is about 9.2% lower than for non-prescriptive proposals, reflecting the pattern we observed at the firm-proposal level. These findings remain robust across multiple variants of specification (3), including different sets of control variables, fixed effects, the addition of a binary "index-fund" variable,[46] and alternative measures of voting support. We also observe a negative and statistically significant coefficient on firm ownership (the percentage of the security held by a given fund) across all specifications. Under the assumption that concentrated owners have more "skin in the game" and thus emphasize pecuniary outcomes, our results suggest a tension between the financial and non-financial aspects of

[42]We follow Brav et al. (2024) and Brav et al. (2022) in using an Ordinary Least Squares (OLS) approach rather than a probit model with fixed effects. In particular, OLS coefficients directly capture the average change in the dependent variable resulting from a one-unit shift in the independent variable, unlike probit coefficients that alter outcome probabilities through the standard normal cumulative distribution function. Additionally, OLS does not rely on the normality of errors under the Gauss-Markov conditions and is less sensitive to distributional assumptions than probit. Its computational simplicity also facilitates easier implementation.

[43]We do not include a separate index for individual proponents in this specification.

[44]Further information about these variables can be found in Table A1.

[45]Further information about these variables can be found in Table A1. These variables are also explicitly enumerated in Table 8.

[46]In column (8), we replicate the specification from column (7) but introduce an "index-fund" variable as the only fund-level control in V. To identify "Index Funds," we begin with the CRSP mutual fund database classification of funds as an index fund or ETF. We then include funds whose names contain any of the terms "Index, Idx, Indx, INDEX, Ind, ETF, Russell, S&P (and its variants such as S&P, SandP, S and P, and SP), DOW (and its variants such as Dow and DJ), MSCI, Bloomberg, KBW, NASDAQ, NYSE, FTSE, Wilshire, Morningstar, 100, 400, 500, 600, 900, 1000, 1500, 2000, 3000, or 5000." Unlike columns (1) through (7), the index-fund coefficient in column (8) is negative and statistically significant at the 1% level, suggesting that a substantial portion of the variability in index-fund voting may be attributed to fund-level factors—such as a fund's expense ratio or assets under management. This observation echoes prior research noting that most index funds belong to large institutional investors characterized by very low expense ratios and high asset levels (Bebchuk and Hirst (2019); Fisch et al. (2019)). Although index funds are not our central focus, we present this specification mainly to show that our findings do not contradict Brav et al. (2024) and Zytnick (2022), who document that index funds vote against E&S proposals more frequently than other fund types.

E&S proposals (Choi (2018)).

5.2 Fund Voting Support after the 2021 Guidance

To analyze the effect of the 2021 Guidance on fund-level voting behavior, we estimate the specification:

$$y_{ijktnm} = \alpha + p_{ijktnm}\beta + (p_{ijktnm} \times Post_t)\gamma + X_{ijktn}\xi + V_{ijktnm}\delta + \theta_i + \eta_j + \psi_k + \kappa_m + \nu_t + \varepsilon_{ijktnm} \quad (4)$$

where specification (3) is modified so that an additional interaction term, $p_{ijktn} \times Post_t$ is included.[47]

Table 9 presents our results, varying the measure of voting support, employing different fixed effects, applying alternate sets of control variables, and introducing selection bias corrections outlined in Section 4.3.3. In our main specification with firm fixed effects (column (1)), the interaction term $p_{ijktnm} \times Post_t$ is negative and statistically significant at the 1% level, indicating a 10.8% decline in support for prescriptive E&S proposals post-2021 relative to their non-prescriptive counterparts. Similar to our baseline findings for all shareholders, the coefficient on p_{ijktnm} is negative but does not reach significance at the 10% level, implying that the variation in prescriptiveness is largely tied to proposals after 2021. Our results concerning the interaction term remain stable across the different specifications. This fund-level analysis parallels the firm-level results presented in Section 4.3.2, reinforcing the notion that the 2021 Guidance has led to a discernible decrease in voting support for these proposals.

5.3 Fund Categories and Support for Prescriptive Proposals

To explore the possibility that certain shareholder groups may be driving the observed decrease in support for prescriptive proposals (see Section 4.3), we modify specification (4) to include the additional binary variable, $FundCat_m$, where $FundCat_m$ denotes a certain class of funds in our dataset (e.g., the "Big Three" funds, or ESG-related funds).[48] Accordingly, we estimate the triple DID specification:

[47] As detailed earlier in Section 9.1, $Post_t$ is a binary indicator that denotes whether the proposal occurs post-treatment (i.e., in 2022 or 2023). Furthermore, as year-fixed effects are incorporated, the term $Post_t$ becomes perfectly collinear with these fixed effects and is therefore excluded from the specification.

[48] In this specification, y_{ijktnm} relates to the binary indicator, "Binary Fund Vote".

$$
\begin{aligned}
y_{ijktnm} = \alpha &+ p_{ijktnm}\beta + FundCat_m\delta \\
&+ (p_{ijktnm} \times Post_t)\epsilon + (p_{ijktnm} \times FundCat_m)\zeta \\
&+ (Post_t \times FundCat_m)\eta + (p_{ijktnm} \times Post_t \times FundCat_m)\theta \\
&+ X_{ijktn}\xi + V_{ijktnm}\iota + \theta_i + \eta_j + \psi_k + \kappa_m + \nu_t + \varepsilon_{ijktnm}
\end{aligned}
\tag{5}
$$

which uses the same indices as specifications (3) and (4). Our triple DID framework aims to capture both the *main* effect of the treatment across the average fund and the *marginal* effect for the specific fund category in question. Notably, a positive and statistically significant marginal effect does not imply that this category of funds has increased its absolute support for prescriptive E&S proposals post-2021; instead, it indicates that these funds support prescriptive proposals *more* than the average fund does.

Before we present the results of our triple DID specification, Table 10 reports the outcomes from estimating specification (4) (with industry fixed effects) across ten different fund categories, including the Big Three funds, Blackrock, active funds, funds sorted by assets under management (AUM), funds sorted by ownership concentration, ESG funds, and ESG funds associated with E&S families.[49] We find a negative and statistically significant coefficient on the interaction term of interest, $p_{ijktnm} \times Post_t$, for 8 out of the 10 categories at the 1% level.[50]

5.3.1 ESG Funds

We begin the triple DID analysis by implementing specification (5), where $FundCat_m$ serves as a binary indicator identifying ESG-related funds and variants, with fixed effects for industry, proponent type, and year.[51] ESG funds are of particular interest because they may exhibit pro-social preferences which may dominate purely pecuniary considerations (Bolton et al. (2020); Michaely et al. (2021)), prompting the question of how they respond to prescriptive E&S proposals. We identify ESG funds in a manner similar to Zytnick (2022), beginning with Morningstar's 2022 list of "sustainable funds," which either integrate ESG factors into investment processes or declare sustainability-related objectives in their prospectuses. We also include funds whose names contain "Sustainable," "ESG," "Social," or "Clean Energy," as well as funds belonging to five es-

[49]The rationale for using industry-fixed effects is explained in Section 9.7.

[50]For the remaining two categories, we find a negative and statistically significant coefficient on the interaction term $p_{ijktnm} \times Post_t$ for ESG funds at the 10% level, while the same interaction term for ESG funds belonging to ES families is negative but not statistically significant.

[51]Following our discussion in Section 9.7, we do not use firm fixed effects in all specifications throughout Section 5.3.

tablished ESG fund families: Calvert, Pax, Parnassus, Trillium, and Praxis. In line with Michaely et al. (2021), who find that E&S funds in non-ESG families may oppose E&S proposals when pivotal votes are at stake, we further construct a measure of family E&S preferences based on each family's average support for E&S proposals in year $t-1$. Specifically, for each family-year, we compute the mean fraction of votes cast in favor of E&S proposals and classify families with below-median support as "non-ES" and those with above-median support as "ES."[52]

We report our findings in Table 11. In column (1), we estimate specification (5) with $FundCat_m$ indicating all ESG funds. Although the coefficient on $p_{ijktnm} \times Post_t$ remains negative and significant at the 1% level, the key interaction term, $p_{ijktnm} \times Post_t \times FundCat_m$, is positive and statistically significant at the 5% level. These results suggest that, relative to the average fund in our sample, ESG funds increased their support for prescriptive E&S proposals by about 5.4% after the 2021 Guidance. Moreover, the coefficient on $FundCat_m$ itself is positive and significant at the 1% level, indicating that ESG funds already provided about 24.5% more voting support for E&S proposals than non-ESG funds, indicating that E&S proposals received approximately 24.5% more voting support from ESG funds compared to non-ESG funds, even when the effects of the 2021 Guidance are not taken into account. This finding aligns with a substantial body of research indicating that ESG funds generally offer stronger support for E&S proposals than non-ESG funds (Dikolli et al. (2022); Curtis et al. (2021); Bolton et al. (2020); Zytnick (2022)).

The results in column (1) of Table 11 point to an ideological dimension in shareholder voting behavior on E&S issues. Socially-oriented funds may be more inclined to back prescriptive E&S proposals even when other funds do not. Indeed, Michaely et al. (2021) suggest that ESG funds in non-E&S families are less ideological in their voting behavior, as they must balance incorporating the pro-social stakeholders interests' they advertise while maximizing shareholder value favored by their families. To test this hypothesis, we construct binary indicators classifying ESG funds into those belonging to "non-ES" families, which face this tradeoff, and those in "ES families", which may prioritize pro-social goals over shareholder value maximization.

In column (2), we present findings where $FundCat_m$ denotes ESG funds in ES families. Consistent with Michaely et al. (2021), these "ideological" funds appear more likely to vote for prescriptive E&S proposals, evidenced by an even larger coefficient on $p_{ijktnm} \times Post_t \times FundCat_m$ compared to column (1). Column (3) focuses on ESG funds in non-ES families. In contrast to columns (1) and (2), we find no evidence that these funds are likelier to support prescriptive E&S proposals

[52]For each family in each year, we calculate the proportion of E&S votes in favor out of all E&S votes cast. We designate those with a below-median level of support as "non-ES," while those above the median are deemed "ES" families (Michaely et al. (2021)).

post-2021. Column (4) presents results for ESG funds belonging to five well-known ESG fund families (Calvert, Pax, Parnassus, Trillium, and Praxis); these findings are consistent with those in column (2), reinforcing our earlier hypothesis about pro-social fund ideologies. Finally, in the remaining columns, we employ an alternative measure of voting support for specifications (1) and (2). The results indicate that our core findings hold under these variations in the measure of voting support.

In Section 2, we posited that preferences over governance issues play a secondary role in this analysis. In Table A7, we provide empirical support for this claim. In column (1), we replicate the baseline specification from column (1) of Table 11. Next, we construct a measure of fund-family governance ("G") preferences, following the same approach used for E&S fund families (Michaely et al. (2021)), but based on the previous year's average support for governance proposals at the family level. In columns (2) and (3), we apply the same specification to ESG funds in G families ("anti-management") and non-G families ("pro-management"). Our findings show that ESG funds maintain consistent voting behavior across both types of families, evidenced by similar coefficients on all key interaction terms (e.g., $p_{ijktnm} \times Post_t \times FundCat_m$).[53] Taken together, these results point to significant homogeneity in how ESG funds approach governance issues, indicating that governance preferences are not the primary driver of voting behavior in our dataset.

5.3.2 Big Three and Active Funds

Our findings in Section 5.3.1 suggest that ESG funds do not play a primary role in reducing support for more prescriptive E&S proposals following the 2021 Guidance. Consequently, we first turn our attention to the "Big Three" fund families (Blackrock, Vanguard, and State Street), who are the largest shareholders in many publicly listed firms where E&S proposals are advanced (Dasgupta et al. (2021)).

In column (1) of Table 12, we report results from specification (5), where $FundCat_m$ is a binary indicator for membership in the "Big Three" (Blackrock, Vanguard, or State Street). The coefficient on $p_{ijktnm} \times Post_t$ remains negative and significant at the 1% level, while the interaction term of interest, $p_{ijktnm} \times Post_t \times FundCat_m$, is positive but not statistically significant at the 10% level. These results imply insufficient evidence to suggest that the Big Three funds differ from the average fund in their support of prescriptive proposals. However, unlike ESG funds, the coefficient on $FundCat_m$ is negative and significant at the 1% level, indicating that E&S proposals receive about

[53] In columns (4) and (5), we repeat the analysis from columns (1) and (2) using the Ordered Fund Vote as the dependent variable, yielding comparable results.

26.3% less support from the "Big Three" than from other funds. This corresponds with prior research showing a tendency for the "Big Three" to oppose E&S proposals (Bolton et al. (2020); Bubb and Catan (2022); Griffin (2020); Pinnington (2023); Heath et al. (2022)). Indeed, Bebchuk and Hirst (2019) and Lund (2018) propose that the Big Three fund families frequently follow management recommendations, which typically oppose E&S proposals. This behavior is attributed to the low-fee index structures characteristic of much of their portfolios, reducing incentives to acquire firm-specific information.

In column (2) of Table 12, we replicate the specification from column (1), but define $FundCat_m$ as a binary indicator for funds belonging to Blackrock, given its public statements alluding to a retreat from backing prescriptive proposals (Blackrock (2022)). Again, we find no indication that Blackrock's voting support diverges from that of the average fund. Although Blackrock may have decreased its backing of prescriptive proposals in absolute terms after the 2021 Guidance, our examination of cross-fund voting behavior suggests that Blackrock's behavior closely resembles the broader mutual fund landscape. This outcome is unsurprising, as Bolton et al. (2020) note that Blackrock and Vanguard generally occupy ideological positions near the average voter on E&S issues.

Given the positive and significant coefficient on $p_{ijktnm} \times Post_t \times FundCat_m$ for ESG funds in ES families, we hypothesize that "financially-oriented" funds, ideologically opposed to ESG funds, may be more inclined to reject prescriptive E&S proposals relative to the average fund (Bolton et al. (2020)). Moreover, a substantial literature finds that "active" mutual funds, which exercise more deliberate voting decisions, adopt markedly different voting stances than "passive" funds, which constitute most of the "Big Three" families (Iliev and Lowry (2015); Brav et al. (2024)).

To probe this hypothesis further, we generate two fund characteristics that serve as "active" mutual fund measures. The first measure, termed "Active (Measure 1)," follows Riley (2021) and Brav et al. (2024). Specifically, we exclude all funds identified by CRSP as index funds,[54] exchange-traded funds, variable annuity funds, funds with Lipper codes indicating a traditional long-only U.S. equity strategy, and funds holding less than 70% of their assets in common equities. We subsequently exclude all such funds in ES families, leaving only "active" funds in non-ES families. For the second "active" mutual fund measure, we draw on evidence suggesting that active funds tend to earn higher alphas (Iliev and Lowry (2015)). We capture this by labeling funds in the top quintile of expense ratios as "active" and again excluding all such funds in ES families, referring

[54]This includes funds whose names contain any of the following terms: "Index, Idx, Indx, INDEX, Ind, ETF, Russell, S&P (and its variants: S & P, S and P, SandP, SP), DOW (and its variants: Dow, DJ), MSCI, Bloomberg, KBW, NASDAQ, NYSE, FTSE, Wilshire, Morningstar, 100, 400, 500, 600, 900, 1000, 1500, 2000, 3000, or 5000."

to this metric as "Active (Measure 2)."

In column (3) of Table 12, we reapply the specification from column (1), defining $FundCat_m$ as a binary indicator for "Active (Measure 1)." Similar to columns (1) and (2), the coefficient on $p_{ijktnm} \times Post_t$ remains negative and significant at the 1% level. However, the key interaction term, $p_{ijktnm} \times Post_t \times FundCat_m$, is negative and statistically significant at the 1% level, implying that—relative to the average fund in our sample—active funds reduced their support for prescriptive E&S proposals by about 5.8% after the 2021 Guidance. This result contrasts sharply with the findings in Table 11, where ESG funds (in ES families) demonstrated a relative increase in support for these proposals. In column (4), we use a similar specification with $FundCat_m$ reflecting "Active (Measure 2)" and obtain findings comparable to column (3). Finally, columns (5) through (8) replicate the analyses from columns (1) to (4) using the "Ordered Fund Vote" variable described in Section 5.1, confirming that our results are robust to alternative measures of voting support.

6 Political Backlash

To establish the causal impact of the 2021 Guidance, the regulatory shock's effect on voting outcomes must occur exclusively through its influence on the prescriptiveness of shareholder proposals. In other words, the shock should not directly affect voting outcomes or do so via channels unrelated to proposal prescriptiveness (see Section 4.3.2). However, recent work by several scholars proposes an alternative explanation for the decrease in voting outcomes—"political backlash" (Garrett and Ivanov (2024); Zhang (2024); Tang et al. (2024); Padfield (2022)). As Curtis (2024) notes, "more than twenty states have adopted at least some type of anti-ESG measure, flows into ESG funds are [ostensibly] declining, and the performance of many ESG funds lagged the broader market in 2022." While the precise source of this backlash remains unclear, we acknowledge that political mechanisms could account for the observed drop in voting outcomes for E&S proposals, potentially violating our identification strategies in earlier sections.

6.1 Big Three Fund-Families

To address concerns about political backlash, we first note that our results in Table 12 may not fully support this narrative. As Bebchuk and Hirst (2022) emphasize, the Big Three fund families have a vested interest in minimizing the risk of public and political backlash, given historical examples where comparable concentrations of financial power provoked such responses. Bebchuk and Hirst (2022) further observe that these funds are inclined to curb these risks by taking a deferential stance toward corporate managers.

Under a counterfactual scenario where political considerations overwhelmingly depress support for E&S proposals, one would expect a strongly negative association between E&S proposal support and the Big Three after 2021, regardless of the proposals' prescriptiveness. However, columns (1) and (5) of Table 12 indicate that the coefficients on the interaction terms $Post_t \times FundCat_m$ and $p_{ijktnm} \times Post_t \times FundCat_m$ are statistically insignificant at the 10% level. In other words, we find no evidence suggesting that the Big Three have altered their voting behavior in response to any hypothesized political backlash.

6.2 ESG Fund Flows

To further address the possibility that our results might be influenced by political backlash, we follow Curtis (2024) in evaluating whether ESG funds have experienced lower fund flows than non-ESG funds since 2021. If political mechanisms were indeed driving the decline in voting support for E&S proposals, we would expect to see a corresponding decrease in flows to ESG funds relative to non-ESG funds post-2021.

We measure fund flows according to the standard definitions in the literature (Sirri and Tufano (1998); Ferreira et al. (2012)), where the fund flow for fund i in month t is calculated as:

$$Flow_{i,t} = \frac{(Assets_{i,t} - Assets_{i,t-1})(1 + Return_{i,t})}{Assets_{i,t-1}}$$

The monthly net assets and returns of funds are obtained from CRSP. Subsequently, we estimate the DID specification:

$$Flow_{it} = \alpha + ESG_i + (ESG_i \times Post_t)\beta + X_{it}\xi + \theta_i + \nu_t + \varepsilon_{it} \tag{6}$$

where i indexes funds, X is a vector of fund-level controls,[55] and θ_i and ν_t represent fund and month fixed effects, respectively. We define ESG_i as a binary indicator marking whether a fund is identified as an ESG fund (pursuant to Section 5.3.1), and $Post_t$ as a binary indicator equal to 1 for periods after November 2021 and 0 otherwise.

[55]Further information about these variables can be found in Table A1.

We report the results of specification (6) in Table A8. All interaction terms ($ESG_i \times Post_t$) in the table are positive but lack statistical significance at the 10% level. This outcome counters a "political backlash" hypothesis, which would predict a drop in ESG fund flows post-2021—evidenced by a negative and significant coefficient on these terms.

To reinforce the findings in Table A8, we employ the triple DID framework from Section 5.3 on the subset of ESG fund votes, letting $FundCat_m$ serve as the variable of interest. In this context, $FundCat_m$ identifies ESG funds in the lowest decile, quintile, or quartile of fund flows.[56] Specifically, we assess how the 2021 Guidance influenced voting on prescriptive proposals among ESG funds with the most pronounced outflows. The coefficient of interest, $p_{ijktnm} \times Post_t \times FundCat_m$, captures how these funds voted on prescriptive proposals relative to the average ESG fund in the dataset.

Our results, presented in Table A9, show that funds with significant negative outflows indeed exhibit weaker support for prescriptive E&S proposals,[57] yet we do not detect statistically significant effects for $p_{ijktnm} \times Post_t \times FundCat_m$ in any specification.[58] These findings imply that even ESG funds with the largest negative outflows did not deviate from the average ESG fund's voting patterns post-2021, further challenging the "political backlash" narrative, which posits that such funds would cut back on support for prescriptive E&S proposals compared to funds with smaller outflows.

6.3 Anti-ESG Proposals

Finally, to address concerns that political mechanisms may be driving the observed decline in voting outcomes for E&S proposals, we emphasize the strong connection between political backlash and the rise of "anti-ESG" proposals. As Welsh (2023) points out, anti-ESG proponents often share political ideologies with politicians who have attempted to pass state laws rejecting ESG considerations in the investment process.

Anti-ESG proposals, led by advocates who urge companies to "stop doing things," strive to "roll back the clock to a mid-20th century world where businesses operated with little consideration of their social and environmental impacts," despite the fact that anti-ESG ideas have gained little

[56]Because fund flows can be negative, the lowest decile, quintile, or quartile corresponds to funds with the largest negative outflows.

[57]This finding supports the notion that negative flows may push fund managers to emphasize financial objectives more strongly (Li et al. (2022)).

[58]Similarly, we observe no statistically significant effect for $p_{ijktnm} \times Post_t$, consistent with Table 11 and its indication that some ESG funds maintain support for prescriptive proposals.

recent traction with investors at large (Welsh (2023)). Nevertheless, the potential influence of these proposals on the decline in voting outcomes for E&S proposals cannot be dismissed, given that their number has more than doubled over the past three years—from 30 in 2021 to 79 in 2023.[59]

We classify anti-ESG proposals according to Welsh (2023), incorporating all proposals by the "National Center for Public Policy Research," the "National Legal and Policy Center," "Inspire Investing LLC," the "Bahnsen Family Trust," the "American Conservative Values ETF," and "Steve J. Milloy." In Table A10, we summarize the subset of anti-ESG proposals that proceeded to a vote. Although these proposals represent just 7.37% of our total sample, a slightly higher share of prescriptive proposals are anti-ESG (9.70%) compared to non-prescriptive proposals (5.77%). Moreover, 53.6% of all anti-ESG proposals are prescriptive. These observations motivate a closer look at whether such proposals have contributed to the reduced voting support for E&S proposals since the 2021 Guidance.

To investigate whether anti-ESG proposals have shaped voting support for E&S proposals post-2021, we adapt specification (2), replacing p_{ijktn} with a binary indicator, $AntiESG_{ijktn}$, to denote an anti-ESG proposal. Table A11 reports our findings under various specifications that include different fixed effects, IPTW weights, and voting support measures. Across all models, the coefficients on $AntiESG_{ijktn}$ are negative and significant at the 1% level, reflecting a general lack of support for such proposals. However, the interaction terms $AntiESG_{ijktn} \times Post_t$ are negative but not statistically significant at the 10% level, suggesting that anti-ESG proposals do not appear to drive the observed post-2021 reduction in voting support. This outcome diverges from a "political backlash" hypothesis, which would anticipate negative and statistically significant coefficients for these terms.

7 Conclusion

In this Article, we explore the significant reduction in shareholder support for E&S proposals post-2021, a trend that reverses the dramatic surge in shareholder support for E&S proposals from 2016 to 2021. Our research unveils novel evidence linking this decline to a change in the SEC's interpretation of Rule 14a-8 in 2021, which effectively allowed shareholders to submit more "prescriptive" E&S proposals. Due to the challenge of directly quantifying a proposal's prescrip-

[59]Unlike the results in Table A10, this figure includes proposals excluded by the SEC that do not proceed to a vote.

tiveness, we employ a combination of supervised and unsupervised machine learning techniques within Natural Language Processing (NLP) to determine the prescriptive nature of these proposals.

Our findings reveal that prescriptive proposals are less favored by investors, receiving approximately 3.75% to 5.38% less voting support compared to their non-prescriptive counterparts. This gap grew substantially after the 2021 Guidance, with support dropping by approximately 6.60% to 8.50%. Although the decline in support is evident across different institutional shareholder categories, there is marked heterogeneity in the degree to which they endorse these proposals. Specifically, funds with stronger E&S preferences are more inclined to back prescriptive proposals, while funds with more financially-oriented objectives are more likely to oppose them. Our results remain robust under multiple tests considering the "political backlash" hypothesis, in which political forces could explain the drop in voting outcomes for E&S proposals.

More broadly, our findings reinforce the viewpoint that many institutional investors might not "walk the talk" when E&S issues clash with pecuniary maximization goals (Goshen and Hamdani (2023); Michaely et al. (2021); Heath et al. (2021)). Although scholars have emphasized pro-social preferences in combating social and environmental externalities (Hart and Zingales (2017); Hart and Zingales (2022); Broccardo et al. (2022); Barber et al. (2021); Hirst et al. (2023); Hart et al. (2024)), we demonstrate that for most funds, the financial costs of prescriptive proposals may outweigh the strength of their E&S commitments.

Finally, while our analysis focuses on voting outcomes, it does not address how the SEC's screening mechanisms shape firm and market behavior beyond the ballot box. Future research could examine their effects on outcomes such as valuation, profitability, exposure to E&S risks (including greenhouse gas emissions), proxy-advisory guidance, and negotiated settlements (He et al. (2023)). Such evidence would enable a fuller welfare assessment of screening rules beyond pass/fail rates and voting tallies.

8 Bibliography

Admati, A. R. and Pfleiderer, P. (2009). The "wall street walk" and shareholder activism: Exit as a form of voice. *The Review of Financial Studies*, 22(7):2645–2685. 2

Aggarwal, D., Litov, L. P., and Rajgopal, S. (2023). Big three (dis) engagements. *Northwestern Law & Econ Research Paper*, (23-17):23–53. 6

Aigner, D. J. et al. (1973). Regression with a binary independent variable subject to errors of observation. *Journal of Econometrics*, 1(1):49–59. 62

Andreoni, J. (1990). Impure altruism and donations to public goods: A theory of warm-glow giving. *Economic Journal*, 100(41):464–477. 2

Andrikogiannopoulou, A., Krueger, P., Mitali, S. F., and Papakonstantinou, F. (2022). Discretionary information in esg investing: A text analysis of mutual fund prospectuses. *Available at SSRN 4082263*. 7

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., and Wager, S. (2021). Synthetic difference-in-differences. *American Economic Review*, 111(12):4088–4118. 12, 57, 58, 89

Badawi, A. B. (2023). How informative is the text of securities complaints? *The Journal of Law, Economics, and Organization*, 39(3):801–827. 7

Bai, J. (2009). Panel data models with interactive fixed effects. *Econometrica*, 77(4):1229–1279. 75

Bainbridge, S. M. (2016). Revitalizing sec rule 14a-8's ordinary business exclusion: Preventing shareholder micromanagement by proposal. *Fordham L. Rev.*, 85:705. 7, 9, 60

Baker, A. C., Larcker, D. F., and Wang, C. C. (2022). How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics*, 144(2):370–395. 16

Balogh, A. and Yonker, S. E. (2024). Is there a business case for racial diversity on corporate boards? *Available at SSRN*. 58

Barber, B. M., Morse, A., and Yasuda, A. (2021). Impact investing. *Journal of Financial Economics*, 139:162–185. 2, 8, 31

Bebchuk, L. and Hirst, S. (2019). Index funds and the future of corporate governance: Theory, evidence, and policy. 119:2029. 4, 20, 21, 26

Bebchuk, L. and Hirst, S. (2022). Big three power, and why it matters. *BUL Rev.*, 102:1547. 4, 20, 27

Bebchuk, L. A., Brav, A., Jiang, W., and Keusch, T. (2020). Dancing with activists. *Journal of Financial Economics*, 137(1):1–41. 8, 11, 17, 18, 70, 71

Belkin, M., Niyogi, P., and Sindhwani, V. (2006). Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of machine learning research*, 7(11). 68

Blackrock, I. S. T. (2022). 2022 climate-related shareholder proposals more prescriptive than 2021. https://www.wsj.com/articles/investors-put-forward-more-proposals-dialing-up-pressure-on-companies-11659432601. Accessed: 2024-03-01. 3, 4, 7, 14, 26

Bolton, P. and Kacperczyk, M. (2021). Do investors care about carbon risk? *Journal of financial economics*, 142(2):517–549. 2

Bolton, P. and Kacperczyk, M. (2023). Global pricing of carbon-transition risk. *The Journal of Finance*, 78(6):3677–3754. 2

Bolton, P., Li, T., Ravina, E., and Rosenthal, H. (2020). Investor ideology. *Journal of Financial Economics*, 137(2):320–352. 4, 5, 8, 20, 23, 24, 26

Brav, A., Cain, M., and Zytnick, J. (2022). Retail shareholder participation in the proxy process: Monitoring, engagement, and voting. *Journal of Financial Economics*, 144(2):492–522. 2, 8, 12, 21

Brav, A., Jiang, W., Li, T., and Pinnington, J. (2024). Shareholder monitoring through voting: New evidence from proxy contests. *The Review of Financial Studies*, 37(2):591–638. 8, 12, 18, 20, 21, 26, 70, 73

Briscoe-Tran, H. (2023). Do employees have useful information about firms' esg practices? *Fisher College of Business Working Paper*, (2021-03):21. 7

Broccardo, E., Hart, O., and Zingales, L. (2022). Exit versus voice. *Journal of Political Economy*, 130(12):3101–3145. 2, 8, 31

Bubb, R. and Catan, E. M. (2022). The party structure of mutual funds. *The Review of Financial Studies*, 35(6):2839–2878. 8, 26

Chen, S., Wong, S., Chen, L., and Tian, Y. (2023). Extending context window of large language models via positional interpolation. *arXiv preprint arXiv:2306.15595*. 65

Choi, A. H. (2018). Concentrated ownership and long-term shareholder value. *Harv. Bus. L. Rev.*, 8:53. 22

Cuñat, V., Gine, M., and Guadalupe, M. (2012). The vote is cast: The effect of corporate governance on shareholder value. *The journal of finance*, 67(5):1943–1977. 8, 11

Curtis, Q. (2024). The esg backlash and the demand for esg mutual funds. Working paper. 27, 28

Curtis, Q., Fisch, J., and Robertson, A. Z. (2021). Do esg mutual funds deliver on their promises? *Mich. L. Rev.*, 120:393. 4, 8, 20, 24

Dammann, J. (2022). Fee-shifting bylaws: An empirical analysis. *The Journal of Law and Economics*, 65(1):1–37. 3

Dasgupta, A., Fos, V., and Sautner, Z. (2021). Institutional investors and corporate governance. *Foundations and Trends in Finance, forthcoming, European Corporate Governance Institute–Finance Working Paper*, 700:2020. 25

Del Guercio, D. and Hawkins, J. (1999). The motivation and impact of pension fund activism. *Journal of Financial Economics*, 52(3):293–340. 4

Del Guercio, D., Seery, L., and Woidtke, T. (2008). Do boards pay attention when institutional investor activists "just vote no"? *Journal of Financial Economics*, 90(1):84–103. 4

Dikolli, S. S., Frank, M. M., Guo, Z. M., and Lynch, L. J. (2022). Walk the talk: Esg mutual fund voting on shareholder proposals. *Review of Accounting Studies*, 27(3):864–896. 8, 24

Egami, N., Fong, C. J., Grimmer, J., Roberts, M. E., and Stewart, B. M. (2022). How to make causal inferences using texts. *Science Advances*, 8(42):eabg2652. 7

Eldar, O. and Wittry, M. D. (2021). Crisis poison pills. *The Review of Corporate Finance Studies*, 10(1):204–251. 12

Era, A., Maia, G., and Scott, L. (2021). Sec's new approach to no-action requests for shareholder esg proposals. https://corpgov.law.harvard.edu/2021/12/04/secs-new-approach-to-no-action-requests-for-shareholder-esg-proposals/. Accessed: 2024-03-01. 10, 14, 18, 71

Ertimur, Y., Ferri, F., and Stubben, S. R. (2010). Board of directors' responsiveness to shareholders: Evidence from shareholder proposals. *Journal of Corporate Finance*, 16(1):53–72. 4

Ferreira, M. A., Keswani, A., Miguel, A. F., and Ramos, S. B. (2012). The flow-performance relationship around the world. *Journal of Banking & Finance*, 36(6):1759–1780. 28

Fisch, A., Jaakkola, T., and Barzilay, R. (2022). Calibrated selective classification. *arXiv preprint arXiv:2208.12084*. 65

Fisch, J., Hamdani, A., and Solomon, S. D. (2019). The new titans of wall street: A theoretical framework for passive investors. *University of Pennsylvania Law Review*, pages 17–72. 21

Fisch, J. E. and Robertson, A. (2023). Shareholder proposals and the debate over sustainability disclosure. *U of Penn, Inst for Law & Econ Research Paper*, (23-29). 7

Frankenreiter, J., Hwang, C., Nili, Y., and Talley, E. (2021). Cleaning corporate governance. *U. Pa. L. Rev.*, 170:1. 7

Gantchev, N. and Giannetti, M. (2021). The costs and benefits of shareholder democracy: Gadflies and low-cost activism. *The Review of Financial Studies*, 34(12):5629–5675. 7, 8, 9, 11, 19, 74

Garrett, D. and Ivanov, I. (2024). Gas, guns, and governments: Financial costs of anti-esg policies. *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper*. 27

Gibson-Dunn (2022). Client alert: Shareholder proposal developments during the 2022 proxy season. https://www.gibsondunn.com/wp-content/uploads/2022/07/shareholder-proposal-developments-during-the-2022-proxy-season.pdf. Accessed: 2024-03-01. 4, 7, 14, 57

Gibson-Dunn (2023). Client alert: Shareholder proposal developments during the 2023 proxy season. https://www.gibsondunn.com/wp-content/uploads/2023/07/shareholder-proposal-developments-during-the-2023-proxy-season.pdf. Accessed: 2024-03-01. 4, 7, 14, 57

Gompers, P., Ishii, J., and Metrick, A. (2003). Corporate governance and equity prices. *The quarterly journal of economics*, 118(1):107–156. 7

Goshen, Z. and Hamdani, A. (2023). Will systematic stewardship save the planet? *Available at SSRN 4605549*. 5, 31

Griffin, C. N. (2020). Environmental & social voting at index funds. *Del. J. Corp. L.*, 44:167. 4, 20, 26

Grootendorst, M. (2022). Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*. 5, 14, 68

Gunel, B., Du, J., Conneau, A., and Stoyanov, V. (2020). Supervised contrastive learning for pre-trained language model fine-tuning. *arXiv preprint arXiv:2011.01403*. 65

Hart, O., Thesmar, D., and Zingales, L. (2024). Private sanctions. *Economic Policy*, pages 203–268. 2, 8, 31

Hart, O. and Zingales, L. (2017). Companies should maximize shareholder welfare not market value. *Journal of Law, Finance, and Accounting*, 2(2):247–275. 2, 8, 31

Hart, O. and Zingales, L. (2022). The new corporate governance. *The University of Chicago Business Law Review*, 1(1):7. 8, 31

He, Y. E., Kahraman, B., and Lowry, M. (2023). Es risks and shareholder voice. *The Review of Financial Studies*, 36(12):4824–4863. 2, 7, 11, 31

Heath, D., Macciocchi, D., Michaely, R., and Ringgenberg, M. C. (2021). Does socially responsible investing change firm behavior? *Available at SSRN 3837706*. 6, 31

Heath, D., Macciocchi, D., Michaely, R., and Ringgenberg, M. C. (2022). Do index funds monitor? *The Review of Financial Studies*, 35(1):91–131. 26

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica: Journal of the econometric society*, pages 153–161. 18, 70

Hendrycks, D. and Gimpel, K. (2016). A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*. 68

Hirst, S., Kastiel, K., and Kricheli-Katz, T. (2023). How much do investors care about social responsibility. *Wisconsin Law Review*, 2023:977–1039. 2, 8, 31

Hu, E., Malenko, N., and Zytnick, J. (2024). Custom proxy voting advice. *Available at SSRN*. 8

Ilhan, E., Krueger, P., Sautner, Z., and Starks, L. T. (2023). Climate risk disclosure and institutional investors. *The Review of Financial Studies*, 36(7):2617–2650. 2, 12

Ilhan, E., Sautner, Z., and Vilkov, G. (2021). Carbon tail risk. *The Review of Financial Studies*, 34(3):1540–1571. 2

Iliev, P. and Lowry, M. (2015). Are mutual funds active voters? *The Review of Financial Studies*, 28(2):446–485. 8, 26

Iliev, P. and Vitanova, S. (2022). Investors' voting preferences. *Available at SSRN 4082127*. 8

Jiang, A. Q., Sablayrolles, A., Mensch, A., Bamford, C., Chaplot, D. S., Casas, D. d. l., Bressand, F., Lengyel, G., Lample, G., Saulnier, L., et al. (2023). Mistral 7b. *arXiv preprint arXiv:2310.06825*. 68

Joffe, M. M., Ten Have, T. R., Feldman, H. I., and Kimmel, S. E. (2004). Model selection, confounder control, and marginal structural models: review and new applications. *The American Statistician*, 58(4):272–279. 72

Joseph, V. R. (2022). Optimal ratio for data splitting. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 15(4):531–538. 60

Kastiel, K. and Nili, Y. (2020). The giant shadow of corporate gadflies. *S. Cal. L. Rev.*, 94:569. 9

Krueger, P., Sautner, Z., and Starks, L. T. (2020). The importance of climate risks for institutional investors. *The Review of Financial Studies*, 33(3):1067–1111. 2

Lewbel, A. (2007). Estimation of average treatment effects with misclassification. *Econometrica*, 75(2):537–551. 63

Li, K., Mai, F., Wong, G., Yang, C., and Zhang, T. (2023a). Female equity analysts and corporate environmental and social performance. *Available at SSRN 4154013*. 7, 14

Li, S. Z., Maug, E., and Schwartz-Ziv, M. (2022). When shareholders disagree: Trading after shareholder meetings. *The Review of Financial Studies*, 35(4):1813–1867. 2, 29

Li, T., Naaraayanan, S. L., and Sachdeva, K. (2023b). Conflicting objectives of esg funds: Evidence from proxy voting. *Available at SSRN 3760753*. 6, 8

Liu, Y. and Lapata, M. (2019). Text summarization with pretrained encoders. *arXiv preprint arXiv:1908.08345*. 4, 13, 60

Lund, D. S. (2018). The case against passive shareholder voting. *Journal of Corporation Law*, 43(3). 26

Matsusaka, J. G., Ozbas, O., and Yi, I. (2019). Opportunistic proposals by union shareholders. *The Review of Financial Studies*, 32(8):3215–3265. 7

Matsusaka, J. G., Ozbas, O., and Yi, I. (2021). Can shareholder proposals hurt shareholders? evidence from securities and exchange commission no-action-letter decisions. *The Journal of Law and Economics*, 64(1):107–152. 7, 13, 74

Matvos, G. and Ostrovsky, M. (2010). Heterogeneity and peer effects in mutual fund proxy voting. *Journal of Financial Economics*, 98(1):90–112. 4, 8, 74

Michaely, R., Ordonez-Calafi, G., and Rubio, S. (2021). Mutual funds' strategic voting on environmental and social issues. *European Corporate Governance Institute–Finance Working Paper*, (774). 5, 6, 8, 23, 24, 25, 31

Michaely, R., Rubio, S., and Yi, I. (2023). Voting rationales. *Available at SSRN 4521854*. 7, 13, 60

Morgan, L. (2024). Reality catches up with the anti-shareholder narrative. *Harvard Law School Forum on Corporate Governance*. Accessed: 2024-07-22. 4, 7

Mosbach, M., Andriushchenko, M., and Klakow, D. (2021). On the stability of fine-tuning bert: Misconceptions, explanations, and strong baselines. In *International Conference on Learning Representations*. 61

Moskalev, A. (2019). Funds of a feather: Influencing corporate elections by voting together. *Job Market Paper*. 72

Padfield, S. J. (2022). An introduction to anti-esg legislation. *Transactions: Tenn. J. Bus. L.*, 24:291. 27

Parnami, A. and Lee, M. (2022). Learning from few examples: A summary of approaches to few-shot learning. *arXiv preprint arXiv:2203.04291*. 65

Pinnington, J. (2023). Are passive investors biased voters? *Available at SSRN 4396879*. 26

Porta, R. L., Lopez-de Silanes, F., Shleifer, A., and Vishny, R. W. (1998). Law and finance. *Journal of political economy*, 106(6):1113–1155. 7

Posner, C. (2022). More prescriptive proposals, less support for 2022 proxy season. https://corpgov.law.harvard.edu/2022/08/22/more-prescriptive-proposals-less-support-for-2022-proxy-season/. Accessed: 2024-03-01. 4, 7, 14

Quiñonero-Candela, J., Sugiyama, M., Schwaighofer, A., and Lawrence, N. D. (2022). *Dataset shift in machine learning*. Mit Press. 68

Rajan, R., Ramella, P., and Zingales, L. (2023). What purpose do corporations purport? evidence from letters to shareholders. Technical report, National Bureau of Economic Research. 7, 13, 60

Rambachan, A. and Roth, J. (2023). A more credible approach to parallel trends. *Review of Economic Studies*, 90(5):2555–2591. 57

Riley, T. B. (2021). Portfolios of actively managed mutual funds. *Financial Review*, 56(2):205–230. 26

Rosenbaum, P. R. (1987). Model-based direct adjustment. *Journal of the American statistical Association*, 82(398):387–394. 72

SEC (2017). Shareholder proposals: Staff legal bulletin no. 14i (cf); no. 14j (cf), no. 14k (cf). https://www.sec.gov/corpfin/staff-legal-bulletin-14i-shareholder-proposals?; https://www.sec.gov/corpfin/staff-legal-bulletin-14j-shareholder-proposals.htm?; https://www.sec.gov/corpfin/staff-legal-bulletin-14k-shareholder-proposals? Accessed: 2024-03-01. 3, 9, 10

SEC (2021). Shareholder proposals: Staff legal bulletin no. 14l (cf). https://www.sec.gov/corpfin/staff-legal-bulletin-14l-shareholder-proposals? Accessed: 2024-03-01. 3, 10

Sirri, E. R. and Tufano, P. (1998). Costly search and mutual fund flows. *The journal of finance*, 53(5):1589–1622. 28

Spamann, H. (2010). The "antidirector rights index" revisited. *The Review of Financial Studies*, 23(2):467–486. 7

Tallarita, R. (2022). Stockholder politics. *Hastings LJ*, 73:1697. 2, 3, 6, 7, 9, 13, 19, 20, 57, 75

Tang, O., Shi, X., and Jiu, L. (2024). Value creation or political trick? an event study on anti-esg regulations. *Finance Research Letters*, 65:105530. 27

Wang, Y., Qu, W., and Ye, X. (2024). Selecting between bert and gpt for text classification in political science research. *arXiv preprint arXiv:2411.05050*. 65

Wang, Y., Yao, Q., Kwok, J. T., and Ni, L. M. (2020). Generalizing from a few examples: A survey on few-shot learning. *ACM computing surveys (csur)*, 53(3):1–34. 65

Welsh, H. (2023). Anti-esg shareholder proposals in 2023. https://corpgov.law.harvard.edu/2023/06/01/anti-esg-shareholder-proposals-in-2023/. Accessed: 2024-06-29. 29, 30

Wen, B., Yao, J., Feng, S., Xu, C., Tsvetkov, Y., Howe, B., and Wang, L. L. (2025). Know your limits: A survey of abstention in large language models. *Transactions of the Association for Computational Linguistics*, 13:529–556. 65

WSJ, J. W.-A. (2022). Investors put forward more proposals, dialing up pressure on companies. https://www.wsj.com/articles/investors-put-forward-more-proposals-dialing-up-pressure-on-companies-11659432601. Accessed: 2024-03-01. 3, 4, 7, 14

Zhang, H. (2021). How using machine learning classification as a variable in regression leads to attenuation bias and what to do about it. *SocArXiv*. 63

Zhang, J. D. (2024). How anti-esg pressure affects investment: Evidence from retirement savings. 27

Zytnick, J. (2022). Do mutual funds represent individual investors? *NYU Law and Economics Research Paper*, (21-04). 4, 11, 18, 20, 21, 23, 24, 70, 74

Figure 1: Average Voting Support over Time (E&S Proposals)

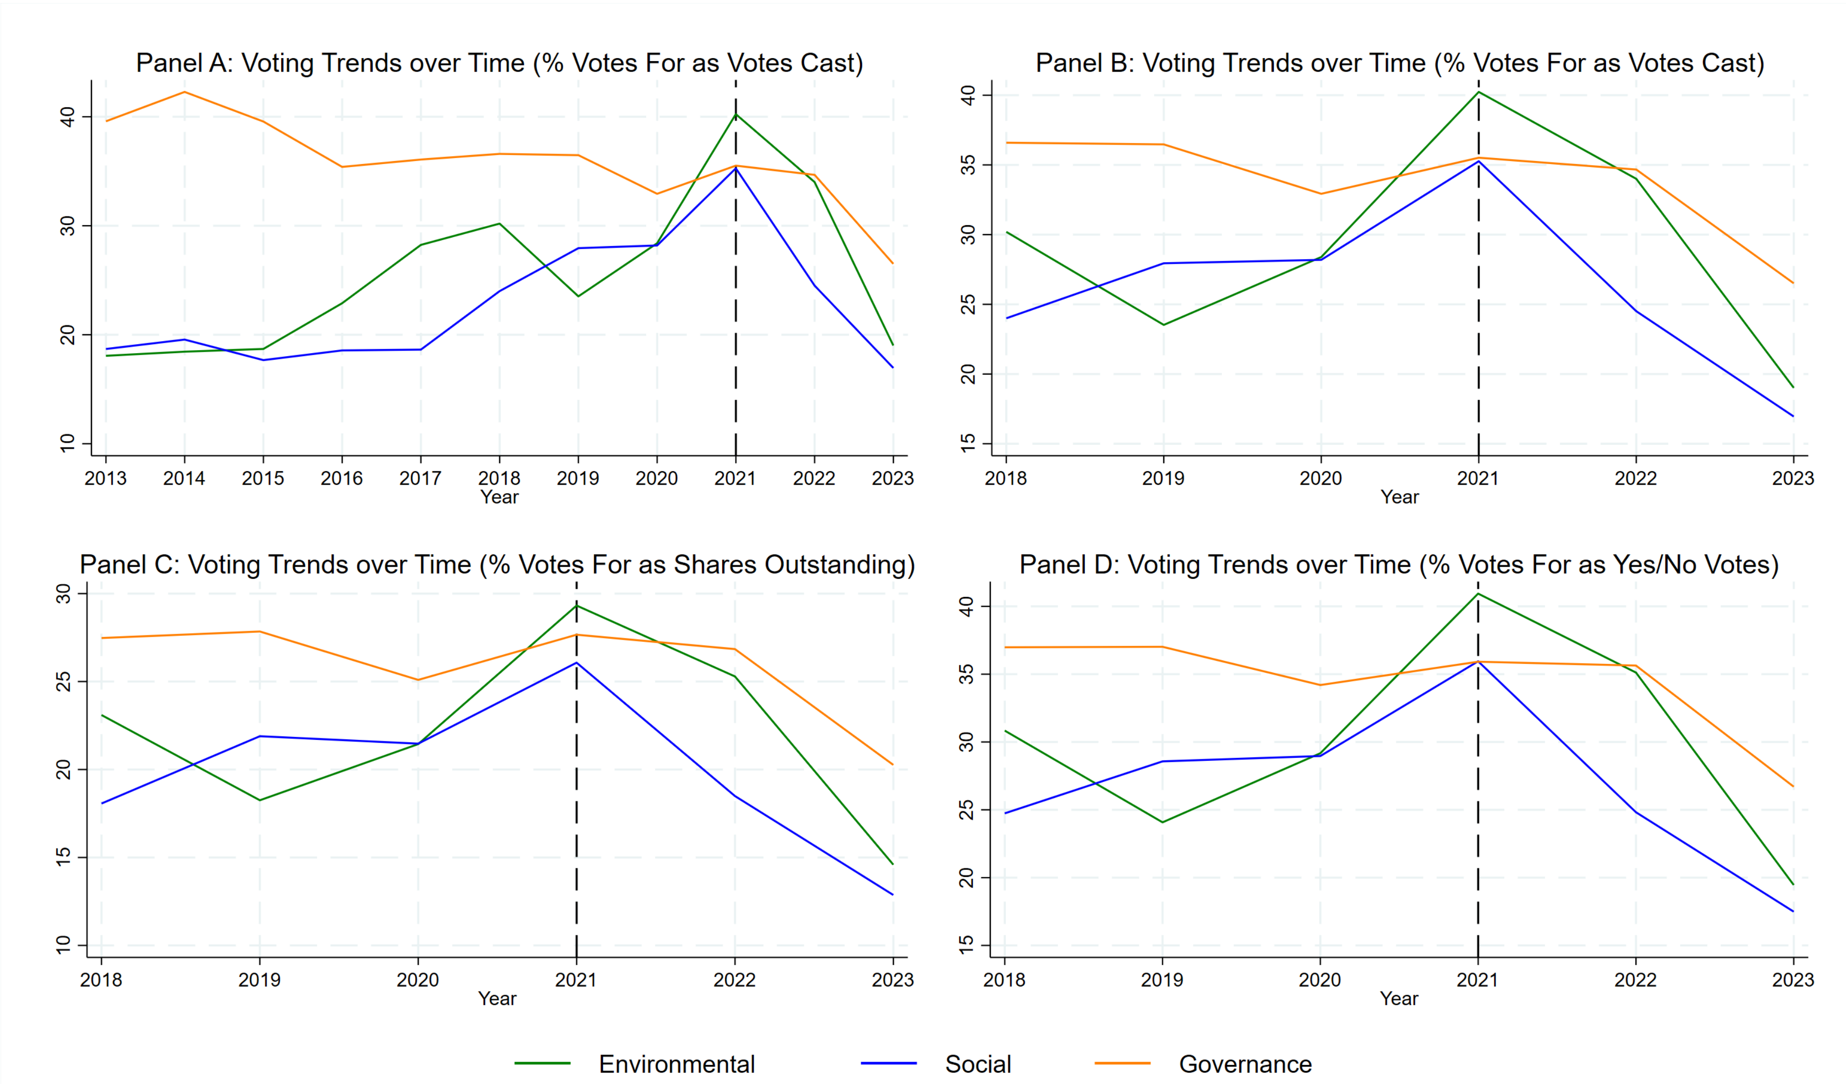

Note: This figure displays the trends over time in voting support for shareholder proposals on environmental (depicted in green) and social (in blue) issues. Panel A charts the support for these proposals from 2013 to 2023, quantified by the proportion of affirmative votes out of the total votes cast. Panel B mirrors Panel A but focuses on the period from 2018 to 2023. Panel C offers a comparative view for the years 2018 to 2023, but measures voting support differently, using the percentage of affirmative votes out of all outstanding shares. Lastly, Panel D, akin to Panels B and C, illustrates voting support for the same timeframe but determines it as the percentage of affirmative votes out of the total of affirmative and negative votes.

Table 1: Summary Statistics: Firm-Proposal Level Data

	Proxy Category			
	Environmental (N=463)	**Governance (N=2,182)**	**Social (N=1,398)**	**Total (N=4,043)**
Votes For As % Votes Cast	28.33 (21.26)	34.00 (21.48)	25.06 (17.41)	30.54 (20.64)
Votes For As % Shares Out	21.33 (16.00)	26.03 (16.89)	18.99 (13.41)	23.28 (16.08)
Votes For As % Yes & No	29.04 (21.92)	34.64 (21.97)	25.65 (17.77)	31.17 (21.10)
Log Mkvalt	10.91 (1.71)	10.34 (1.98)	11.36 (1.71)	10.76 (1.92)
Tobin's Q	2.11 (1.66)	2.50 (2.33)	2.95 (2.26)	2.61 (2.25)
RoA	0.12 (0.08)	0.12 (0.11)	0.14 (0.10)	0.13 (0.10)
Leverage Ratio	-2.75 (107.94)	1.63 (26.11)	0.32 (15.76)	0.68 (42.25)
Firm Size	11.06 (1.85)	10.38 (2.01)	11.11 (1.70)	10.71 (1.93)
HHI	0.26 (0.25)	0.29 (0.25)	0.32 (0.27)	0.30 (0.26)
Inst Own	0.66 (0.21)	0.71 (0.19)	0.66 (0.19)	0.69 (0.20)
Inst HHI	0.04 (0.03)	0.04 (0.03)	0.04 (0.02)	0.04 (0.03)
Proxy Subcategory				
Board Related	0 (0.0%)	508 (23.3%)	0 (0.0%)	508 (12.6%)
Capital Stock	0 (0.0%)	2 (0.1%)	0 (0.0%)	2 (0.0%)
Environmental Issues	463 (100.0%)	0 (0.0%)	0 (0.0%)	463 (11.5%)
Executive Compensation Related	0 (0.0%)	288 (13.2%)	0 (0.0%)	288 (7.1%)
Fund Related	0 (0.0%)	1 (0.0%)	0 (0.0%)	1 (0.0%)
Miscellaneous	0 (0.0%)	37 (1.7%)	0 (0.0%)	37 (0.9%)
Miscellaneous Corporate Governance	0 (0.0%)	196 (9.0%)	0 (0.0%)	196 (4.8%)
Proxy Fight Specific	0 (0.0%)	34 (1.6%)	0 (0.0%)	34 (0.8%)
Shareholder Rights/Takeover Defense	0 (0.0%)	1,083 (49.6%)	0 (0.0%)	1,083 (26.8%)
Social Issues Related	0 (0.0%)	0 (0.0%)	1,398 (100.0%)	1,398 (34.6%)
Value Maximization	0 (0.0%)	33 (1.5%)	0 (0.0%)	33 (0.8%)
Proponent Type Description				
Misc	52 (11.2%)	148 (6.8%)	114 (8.2%)	314 (7.8%)
Corporation	0 (0.0%)	2 (0.1%)	0 (0.0%)	2 (0.0%)
Hedge Fund Company	6 (1.3%)	29 (1.3%)	5 (0.4%)	40 (1.0%)
Individual	97 (21.0%)	1,472 (67.5%)	298 (21.3%)	1,867 (46.2%)
Investment Adviser	63 (13.6%)	36 (1.6%)	101 (7.2%)	200 (4.9%)
Labor Union	9 (1.9%)	103 (4.7%)	104 (7.4%)	216 (5.3%)
Mutual Fund Manager	3 (0.6%)	1 (0.0%)	11 (0.8%)	15 (0.4%)
Other Institutions	11 (2.4%)	32 (1.5%)	85 (6.1%)	128 (3.2%)
Other Stake Holders	141 (30.5%)	193 (8.8%)	364 (26.0%)	698 (17.3%)
Public Pension Fund	32 (6.9%)	130 (6.0%)	153 (10.9%)	315 (7.8%)
Religious Groups	49 (10.6%)	36 (1.6%)	163 (11.7%)	248 (6.1%)
Has No Action Letter Sought				
No	272 (58.7%)	1,524 (69.8%)	913 (65.3%)	2,709 (67.0%)
Yes	191 (41.3%)	658 (30.2%)	485 (34.7%)	1,334 (33.0%)

Note: This table provides summary statistics for our dataset at the firm-proposal-year level, omitting information related to the "prescriptiveness" metric and fund-level data. Mean values for continuous variables are presented without the use of parentheses, whereas their standard deviations are enclosed in parentheses. In the case of factor or binary variables, the frequencies of these variables are provided without parentheses, while the percentages of factor variables are indicated within parentheses.

Table 2: Examples of Prescriptive Proposals

Policy	Resolution
Proposal Pair #1 (Similar Proposals)	
Pre-2021 Guidance (Precluded: No Vote)	**Resolved:** Shareholders request that the Travelers' Board of Directors adopt and disclose new policies to help ensure that its underwriting practices do not support new fossil fuel supplies, in alignment with the IEA's Net Zero Emissions by 2050 Scenario.
Post-2021 Guidance (Permitted: Proceeds to Vote)	**Resolved:** Shareholders request that JPMorgan Chase (JPMC) adopt a policy by the end of 2022 in which the company takes available actions to help ensure that its financing does not contribute to new fossil fuel supplies that would be inconsistent with the IEA's Net Zero Emissions by 2050 Scenario.
Proposal Pair #2 (Similar Proposals)	
Pre-2021 Guidance (Precluded: No Vote)	**Resolved:** Shareholders request management review its policies related to human rights to assess areas where the Company needs to adopt and implement additional policies and to report its findings, omitting proprietary information and prepared at reasonable expense, by December 2018.
Post-2021 Guidance (Permitted: Proceeds to Vote)	**Resolved:** Shareholders direct the board of directors of Meta Platforms, Inc. (formerly Facebook, Inc.) to publish an independent third-party Human Rights Impact Assessment (HRIA) examining the actual and potential human-rights impacts of Facebook's targeted advertising policies and practices throughout its business operations. This HRIA should be conducted at reasonable cost; omit proprietary and confidential information, as well as information relevant to litigation or enforcement actions; and be published on the company's website by June 1, 2023.
Proposal Pair #3 (Similar Proposals)	
Pre-2021 Guidance (Precluded: No Vote)	**Resolved:** Shareholders request Wal-Mart prepare a report on the risks to the company associated with emerging public policies on the gender pay gap, including associated reputational, competitive, and operational risks, and risks related to recruiting and retaining female talent. The report should be prepared at reasonable cost, omitting proprietary information, litigation strategy, and legal compliance information.
Post-2021 Guidance (Permitted: Proceeds to Vote)	**Resolved:** Shareholders ask that the board commission and publish a report on (1) whether the Company participates in compensation and workforce practices that prioritize Company financial performance over the economic and social costs and risks created by inequality and racial and gender disparities and (2) the manner in which any such costs and risks threaten returns of diversified shareholders who rely on a stable and productive economy.

This table lists three illustrative pairs of proposals that are nearly identical in substance. In each pair, the first proposal was excluded under the pre-Guidance regime and never came to a vote, whereas the second was permitted to appear on the ballot—and therefore proceeded to a vote—after the 2021 Guidance.

Table 3: Summary Statistics: Frequencies and Percentages of Prescriptive Proposals

	Year						
	2018	2019	2020	2021	2022	2023	Total
Panel A: Prescriptiveness (Voted On)							
Prescriptiveness = 0							
Frequency	95	98	107	106	167	162	735
Percent (Within-Year)	62.50	61.25	63.69	68.39	59.64	53.29	60.30
Prescriptiveness = 1							
Frequency	57	62	61	49	113	142	484
Percent (Within-Year)	37.50	38.75	36.31	31.61	40.36	46.71	39.70
Panel B: Prescriptiveness (Not Voted On)							
Prescriptiveness = 0							
Frequency	50	49	47	52	44	33	275
Percent (Within-Year)	44.64	51.58	37.01	36.36	44.00	50.77	42.83
Prescriptiveness = 1							
Frequency	62	46	80	91	56	32	367
Percent (Within-Year)	55.36	48.42	62.99	63.64	56.00	49.23	57.17
Panel C: Total (Voted and Not Voted On)							
Prescriptiveness = 0							
Frequency	145	147	154	158	211	195	1,010
Percent (Within-Year)	54.92	57.65	52.20	53.02	55.53	52.85	54.27
Prescriptiveness = 1							
Frequency	119	108	141	140	169	174	851
Percent (Within-Year)	45.08	42.35	47.80	46.98	44.47	47.15	45.73
All Proposals (Prescriptiveness = 0 or 1)							
Frequency	264	255	295	298	380	369	1,861
Percent (Across-Years)	14.19	13.70	15.85	16.01	20.42	19.83	100.00

Note: This Table illustrates the frequencies and within-year percentages of prescriptive proposals over time. In Panels A and B, proposals are categorized based on whether they advance to a vote. Panel C then combines these categories, while also providing frequencies and across-year percentages for all proposals (regardless of whether they are prescriptive or not). Prescriptive proposals are indicated by the header “Prescriptiveness = 1”, while non-prescriptive proposals are indicated by the header “Prescriptiveness = 0”.

Figure 2: Proportion of Prescriptive Proposals over Time

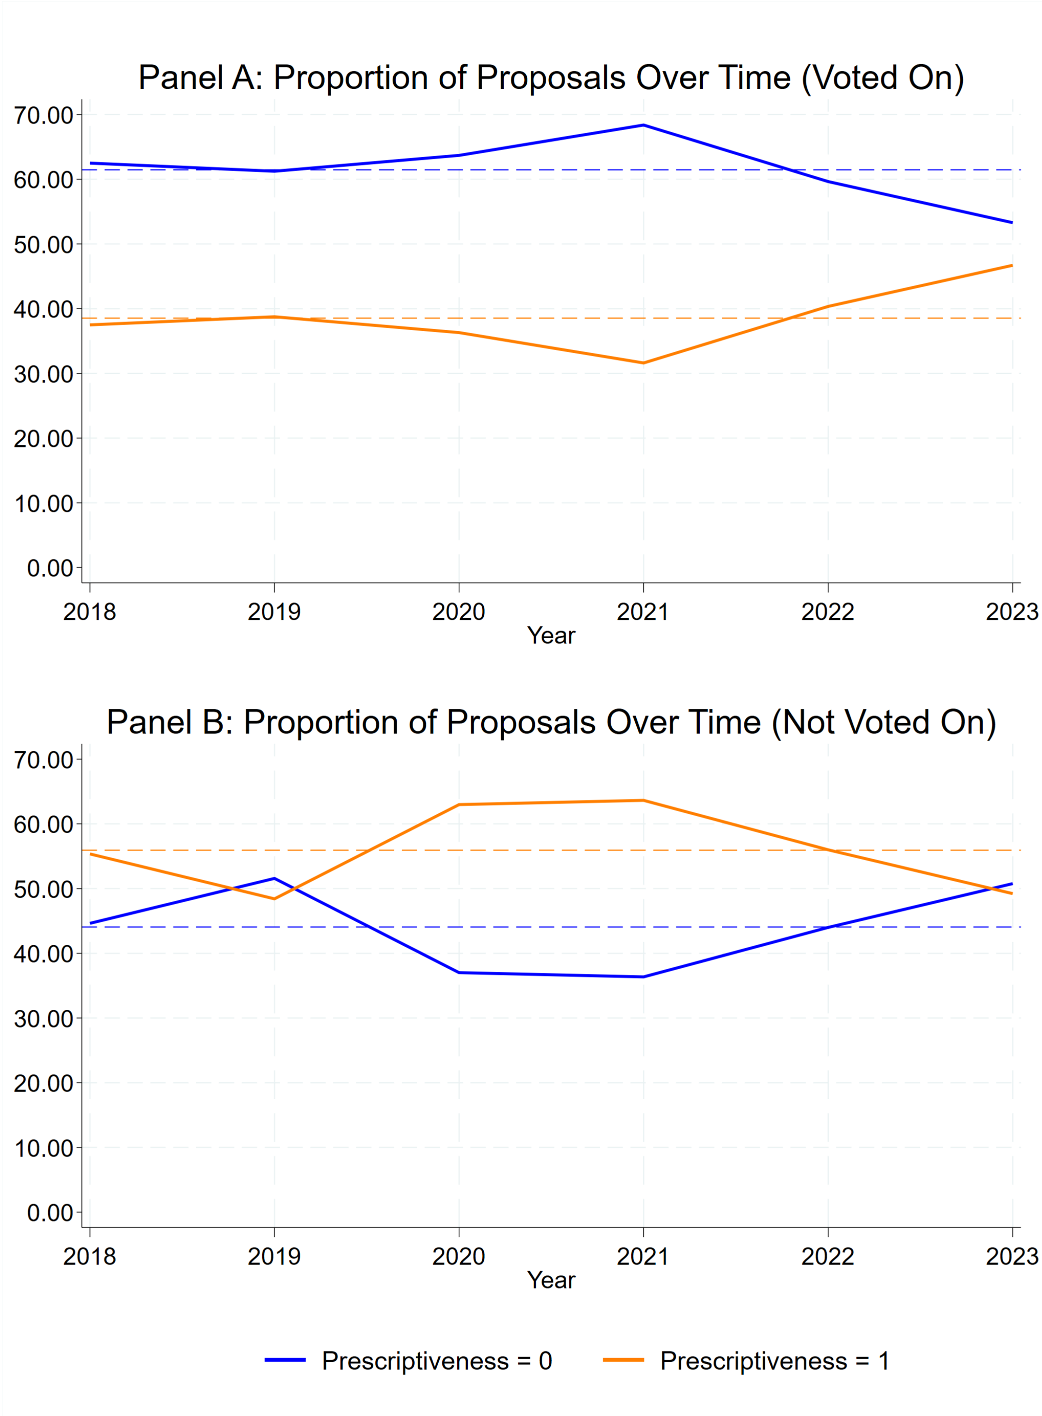

Note: This figure highlights the evolving trends over time in the percentages of prescriptive shareholder proposals. Panel A charts the proportion of shareholder proposals which proceed to a vote, where the orange line represents prescriptive proposals and the blue line denotes non-prescriptive ones. Panel B, on the other hand, illustrates the percentage of shareholder proposals that do not advance to a vote, encompassing proposals that are withdrawn, settled, or excluded. Like Panel A, the orange line represents prescriptive proposals, while the blue line denotes non-prescriptive proposals.

Table 4: Panel Regressions of Voting Support on Prescriptiveness

	Votes For As % Votes Cast				Votes For As % Shares Out	Votes For As % Yes & No
	(1) E & S	(2) E & S	(3) Environmental	(4) Social	(5) E & S	(6) E & S
Prescriptiveness	-5.219***	-6.340***	-11.158***	-4.741***	-3.750***	-5.382***
	(0.000)	(0.000)	(0.006)	(0.000)	(0.000)	(0.000)
Observations	1082	1180	205	831	1082	1080
Firm FE	Yes	No	Yes	Yes	Yes	Yes
Industry FE	No	Yes	No	No	No	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Proponent-Type FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.403	0.302	0.515	0.442	0.423	0.405
F Statistic	3.087	5.466	11.979	2.099	3.334	3.155

Note: *** p≤0.01, ** p≤0.05, * p≤0.10. P-values are reported in parentheses. In specifications (1) through (4), the dependent variable relates to the percentage of affirmative votes out of the total votes cast. For specification (5), it relates to the percentage of affirmative votes out of all outstanding shares. In specification (6), it relates to the percentage of affirmative votes out of the total of affirmative and negative votes. Additionally, specifications (1), (2), (5), and (6) apply to all E&S proposals. Specification (3) specifically addresses environmental proposals, and specification (4) focuses on social proposals. we suppress reporting of the constant term and firm-proposal controls. Standard errors are clustered at the meeting-level.

Table 5: Changes in Prescriptiveness Post Treatment

	Votes For As % Votes Cast				Votes For As % Shares Out	Votes For As % Yes & No
	(1) E & S	(2) E & S	(3) Env	(4) Social	(5) E & S	(6) E & S
Prescriptiveness × Post	-8.476***	-11.000***	-19.466***	-5.299**	-6.605***	-8.497***
	(0.000)	(0.000)	(0.001)	(0.023)	(0.000)	(0.000)
Prescriptiveness	-0.777	-0.784	0.371	-1.985	-0.289	-0.947
	(0.626)	(0.606)	(0.943)	(0.235)	(0.808)	(0.559)
Observations	1082	1180	205	831	1082	1080
Firm FE	Yes	No	Yes	Yes	Yes	Yes
Industry FE	No	Yes	No	No	No	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Proponent-Type FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.414	0.321	0.543	0.447	0.435	0.416
F Statistic	4.256	8.036	12.144	2.288	4.549	4.256

Note: *** p≤0.01, ** p≤0.05, * p≤0.10. P-values are reported in parentheses. In specifications (1) through (4), the dependent variable relates to the percentage of affirmative votes out of the total votes cast. For specification (5), it relates to the percentage of affirmative votes out of all outstanding shares. In specification (6), it relates to the percentage of affirmative votes out of the total of affirmative and negative votes. Additionally, specifications (1), (2), (5), and (6) apply to all E&S proposals. Specification (3) specifically addresses environmental proposals, and specification (4) focuses on social proposals. we suppress reporting of the constant term and firm-proposal controls. Standard errors are clustered at the meeting-level.

Figure 3: Voting Support for Prescriptive Proposals over Time

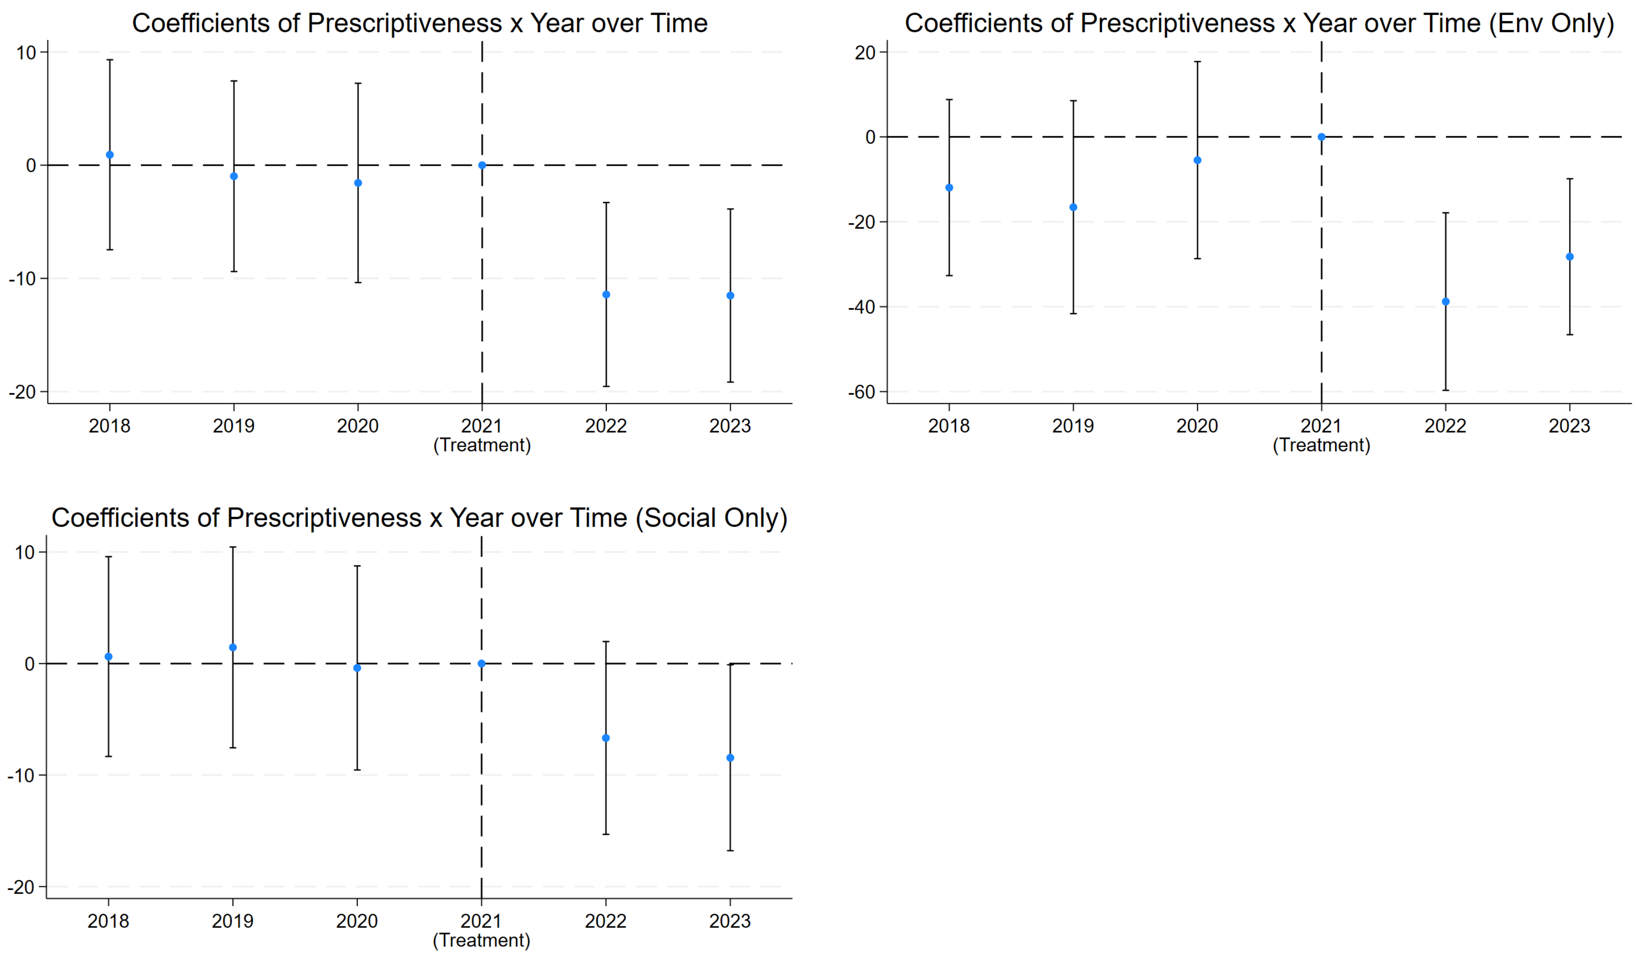

Note: This figure displays the coefficients of interaction terms which combine year-specific indicators with a binary variable distinguishing prescriptive from non-prescriptive proposals. In estimating these coefficients, we include firm-proposal controls along with fixed effects for firm, year, and the type of proponent. we adopt a baseline year of 2021, corresponding to the year when treatment occurred. The dependent variable relates to the percentage of affirmative votes out of the total votes cast. Confidence intervals are drawn at the 95% level. Panel A of the figure illustrates the coefficients for E&S proposals, Panel B showcases those for environmental proposals, and Panel C presents the coefficients for social proposals.

Table 6: Changes in Prescriptiveness Post Treatment: New vs Existing Proponents

	Baseline	Proponent FE			New Proponents	
	(1)	(2)	(3)	(4)	(5)	(6)
Prescriptiveness × Post	-8.476***	-6.587***	-8.635***	-6.501***	-7.315***	-11.561***
	(0.000)	(0.008)	(0.000)	(0.008)	(0.008)	(0.000)
Prescriptiveness	-0.777	1.634	0.689	1.411	-1.966	-0.898
	(0.626)	(0.367)	(0.689)	(0.455)	(0.363)	(0.682)
Share of Prescriptive Proposals				0.733		
				(0.780)		
First Apperance					1.713	4.299**
					(0.362)	(0.035)
Prescriptiveness × First Appearance					2.853	0.223
					(0.299)	(0.936)
Post × First Apperance					0.418	-5.385*
					(0.877)	(0.051)
Prescriptiveness × Post × First Apperance					-2.398	3.049
					(0.593)	(0.478)
Observations	1082	923	1011	923	1082	1180
Firm FE	Yes	Yes	No	Yes	Yes	No
Industry FE	No	No	Yes	No	No	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Proponent-Type FE	Yes	No	No	No	Yes	Yes
Proponent FE	No	Yes	Yes	Yes	No	No
Firm-Proposal Controls	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.414	0.527	0.457	0.527	0.415	0.323
F Statistic	4.256	2.861	4.425	2.666	3.608	6.443

Note: *** $p \leq 0.01$, ** $p \leq 0.05$, * $p \leq 0.10$. P-values are reported in parentheses. The dependent variable in all specifications relates to the proportion of affirmative votes out of the total votes cast. In specifications (2) and (3), we replace proponent-type fixed effects with proponent fixed effects. In specifications (4) and (5), we include an additional variable in a triple DID specification, "First Appearance", which denotes when a proposal is first submitted by a new proponent not observed in prior years. We suppress reporting of the constant term and firm-proposal controls. Standard errors are clustered at the meeting-level.

Figure 4: Proportion of E&S Proposals submitted to Stable Firms

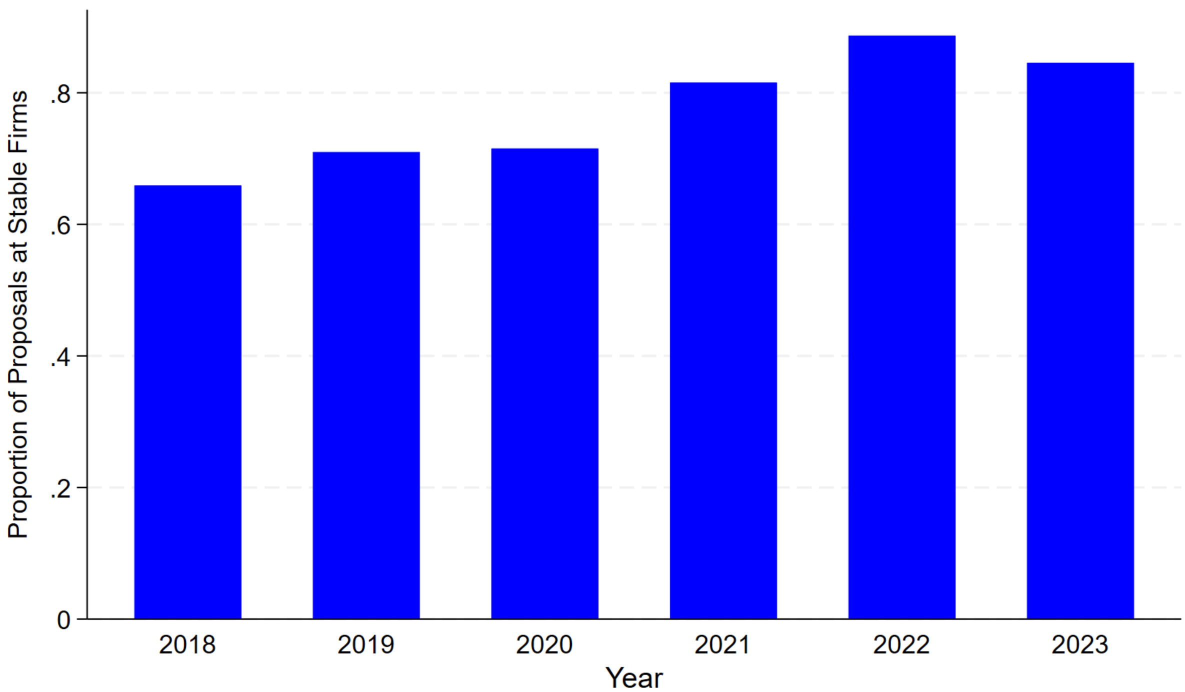

Note: This figure provides an annual breakdown of the proportion of E&S proposals submitted to "Stable Firms", which we define as firms present in our sample both before and after the 2021 Guidance. On average, 78.34% of all proposals in our sample were submitted to stable firms. Specifically, 65.91% of proposals in 2018, 70.98% in 2019, 71.53% in 2020, 81.54% in 2021, 88.68% in 2022, and 84.55% in 2023 were directed toward these firms.

Table 7: Summary Statistics: Individual Fund Votes on E&S Proposals

	Proxy Category		
	Environmental Issues (N=185,294)	**Social Issues Related (N=714,926)**	**Total (N=900,220)**
Binary Fund Vote	0.4466 (0.4971)	0.3990 (0.4897)	0.4088 (0.4916)
Ordered Fund Vote	0.4197 (0.4797)	0.3748 (0.4714)	0.3841 (0.4735)
% of Security owned by Fund	0.0004 (0.0022)	0.0004 (0.0020)	0.0004 (0.0020)
Security as % of Fund's Total Assets	0.0114 (0.0191)	0.0139 (0.0212)	0.0134 (0.0208)
Total Fund Assets	5.7939 (2.2630)	5.7133 (2.2481)	5.7299 (2.2514)
Mgmt Fees	0.3189 (0.5520)	0.3221 (0.5882)	0.3215 (0.5809)
Expense Ratio	0.0065 (0.0048)	0.0068 (0.0048)	0.0067 (0.0048)
Turnover Ratio	0.6050 (0.8947)	0.6128 (0.8729)	0.6112 (0.8774)
Index Fund	0.4586 (0.4983)	0.4236 (0.4941)	0.4308 (0.4952)

Note: This table provides summary statistics for our dataset at the fund-firm-proposal-year level, omitting information related to the "prescriptiveness" metric and firm-proposal-level data. Mean values for continuous variables are presented without the use of parentheses, whereas their standard deviations are enclosed in parentheses. In the case of factor or binary variables, the frequencies of these variables are provided without parentheses, while the percentages of factor variables are indicated within parentheses.

Table 8: Panel Regressions of Individual Fund Votes on Prescriptiveness

	Binary Fund Vote								Ordered Fund Vote	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Prescriptiveness	-0.092***	-0.094***	-0.085***	-0.092***	-0.085***	-0.085***	-0.084***	-0.084***	-0.081***	-0.082***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
% of Security owned by Fund	-4.557***	-4.668***	-4.572***	-2.949***	-2.986***	-3.158***			-2.977***	-4.567***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)			(0.000)	(0.000)
Security as % of Fund's Total Assets	-0.950***	-0.857***	-0.949***	-0.117	-0.112	-0.080			-0.131*	-0.950***
	(0.000)	(0.000)	(0.000)	(0.112)	(0.122)	(0.237)			(0.057)	(0.000)
Total Fund Assets	-0.027***	-0.027***	-0.027***	0.000	0.001				0.005***	-0.027***
	(0.000)	(0.000)	(0.000)	(0.906)	(0.758)				(0.008)	(0.000)
Mgmt Fees	0.008***	0.008***	0.008***	0.008	0.009				0.008	0.008***
	(0.002)	(0.002)	(0.003)	(0.119)	(0.107)				(0.123)	(0.001)
Expense Ratio	7.679***	7.602***	7.680***	-0.529	-0.464				-0.277	7.903***
	(0.000)	(0.000)	(0.000)	(0.645)	(0.683)				(0.798)	(0.000)
Turnover Ratio	-0.001	-0.000	-0.001	0.004	0.004				0.003	-0.001
	(0.392)	(0.829)	(0.389)	(0.165)	(0.175)				(0.288)	(0.611)
Index Fund	0.045***	0.046***	0.045***	-0.017	-0.016			-0.018***	-0.014	0.041***
	(0.000)	(0.000)	(0.000)	(0.115)	(0.136)			(0.000)	(0.197)	(0.000)
Observations	582861	582861	582861	582509	582509	802593	875372	848795	582509	582861
Firm FE	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	No	Yes	No	No	No	No	No	No	No	No
Fund FE	No	No	No	Yes	Yes	Yes	Yes	No	Yes	No
Proponent-Type FE	No	No	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.152	0.113	0.160	0.363	0.371	0.384	0.385	0.127	0.375	0.163
F Statistic	143.256	141.521	144.250	4.200	4.151	4.818	3.147	7.188	5.025	139.536

Note: *** p≤0.01, ** p≤0.05, * p≤0.10. P-values are reported in parentheses. In specifications (1) through (8), the dependent variable relates to the variable "Binary Fund Vote", which takes on a value of 1 when a specific fund votes in favor of a proposal, and 0 for all other outcomes. In specifications (9) and (10), the dependent variable relates to the variable "Ordered Fund Vote", which is assigned a value of 1 for a "yes" vote, 0.5 for an "abstained" vote, and 0 for all other outcomes. We suppress reporting of the constant term and firm-proposal controls. Standard errors are clustered at the meeting-level.

Table 9: Fund Voting Support after the 2021 Guidance

	Binary Fund Vote					Ordered Fund Vote	
	(1) Uncorrected	(2) Uncorrected	(3) Uncorrected	(4) Uncorrected	(5) IPTW	(6) Uncorrected	(7) IPTW
Prescriptiveness × Post	-0.108***	-0.081**	-0.109***	-0.111***	-0.110***	-0.111***	-0.109***
	(0.002)	(0.025)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Prescriptiveness	-0.033	-0.041	-0.031	-0.030	-0.032	-0.027	-0.031
	(0.249)	(0.180)	(0.272)	(0.277)	(0.269)	(0.333)	(0.254)
% of Security owned by Fund	-4.690***	-4.560***	-17.243***		-4.547***	-4.664***	-4.479***
	(0.000)	(0.000)	(0.000)		(0.000)	(0.000)	(0.000)
Security as % of Fund's Total Assets	-0.847***	-0.947***	-0.494***		-0.812***	-0.851***	-0.811***
	(0.000)	(0.000)	(0.000)		(0.000)	(0.000)	(0.000)
Total Fund Assets	-0.027***	-0.027***			-0.027***	-0.027***	-0.026***
	(0.000)	(0.000)			(0.000)	(0.000)	(0.000)
Mgmt Fees	0.008***	0.008***			0.008***	0.008***	0.008***
	(0.002)	(0.003)			(0.002)	(0.001)	(0.000)
Expense Ratio	7.601***	7.670***			7.249***	7.839***	7.463***
	(0.000)	(0.000)			(0.000)	(0.000)	(0.000)
Turnover Ratio	-0.000	-0.001			-0.001	0.000	-0.000
	(0.857)	(0.383)			(0.638)	(0.901)	(0.844)
Index Fund	0.046***	0.045***			0.044***	0.042***	0.040***
	(0.000)	(0.000)			(0.000)	(0.000)	(0.000)
Constant	0.795***	1.374**	0.723***	0.722***	0.824***	0.811***	0.836***
	(0.000)	(0.033)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	582861	582861	803030	898820	582861	582861	582861
Firm FE	No	Yes	No	No	No	No	No
Industry FE	Yes	No	Yes	Yes	Yes	Yes	Yes
Proponent-Type FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fund Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.125	0.161	0.096	0.090	0.129	0.127	0.120
F Statistic	134.010	134.979	111.062	5.597	122.601	132.670	119.653

Note: *** $p\leq0.01$, ** $p\leq0.05$, * $p\leq0.10$. P-values are reported in parentheses. In specifications (1) through (5), the dependent variable relates to the variable "Binary Fund Vote", which takes on a value of 1 when a specific fund votes in favor of a proposal, and 0 for all other outcomes. In specifications 6 and (7), the dependent variable relates to the variable "Ordered Fund Vote", which is assigned a value of 1 for a "yes" vote, 0.5 for an "abstained" vote, and 0 for all other outcomes. Additionally, in specifications (5) and (7), we implement the Inverse Probability of Treatment Weighting (IPTW) model introduced in Section 9.4. We suppress reporting of the constant term and firm-proposal controls. Standard errors are clustered at the meeting-level.

Table 10: Changes in Prescriptiveness Post Treatment: Fund-Level (Subsets)

	(1) Big Three	(2) Blackrock	(3) Active	(4) 5th Quintile AUM	(5) 1st Quintile AUM	(6) 5th Quintile Concentration	(7) 1st Quintile Concentration	(8) ESG Fund	(9) ESG Fund (Non-ES Family)	(10) ESG Fund (ES Family)
Prescriptiveness × Post	-0.073**	-0.106***	-0.153***	-0.078***	-0.129***	-0.105***	-0.091**	-0.068*	-0.078**	-0.058
	(0.012)	(0.003)	(0.001)	(0.007)	(0.002)	(0.000)	(0.049)	(0.078)	(0.045)	(0.184)
Prescriptiveness	-0.020	-0.006	-0.042	-0.018	-0.039	-0.003	-0.064	-0.052	-0.040	-0.064*
	(0.356)	(0.818)	(0.282)	(0.455)	(0.246)	(0.872)	(0.104)	(0.103)	(0.232)	(0.069)
Constant	0.904***	0.653***	0.637***	0.721***	0.693***	0.844***	0.752***	0.601***	0.324	0.454*
	(0.000)	(0.002)	(0.003)	(0.000)	(0.001)	(0.000)	(0.002)	(0.002)	(0.148)	(0.051)
Observations	94292	44224	44455	133572	127960	141372	118520	37290	12517	24763
Firm FE	No	No	No	No	No	No	No	No	No	No
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Proponent-Type FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fund Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.183	0.314	0.129	0.132	0.114	0.146	0.112	0.163	0.118	0.260
F Statistic	19.554	4.648	19.258	71.929	15.919	67.248	15.006	43.367	9.620	31.072

Note: *** p≤0.01, ** p≤0.05, * p≤0.10. P-values are reported in parentheses. The dependent variable in all specifications relates to the proportion of affirmative votes out of the total votes cast. In specification (1), we estimate the model for Big Three funds; in specification (2), for Blackrock funds; in specification (3), for Active funds (defined in Section 5.3.2); in specification (4), for the top quintile of funds sorted by Assets under Management (AUM); in specification (5), for the bottom quintile of funds sorted by AUM; in specification (6), for the top quintile of funds sorted by their concentration of holdings; in specification (7), for the bottom quintile of funds sorted by their concentration of holdings; in specification (8), for ESG funds; in specification (9), for ESG funds in non-ES families (defined in Section 5.3.1); and in specification (10), for ESG funds in ES families (defined in Section 5.3.1). We drop all anti-ESG proposals which we identify in Section 6.3. We suppress reporting of the constant term and firm-proposal controls. Standard errors are clustered at the meeting-level.

Table 11: Changes in Prescriptiveness Post Treatment: Heterogeneity amongst ESG Funds

	Binary Fund Vote				Ordered Fund Vote	
	(1)	(2)	(3)	(4)	(5)	(6)
Prescriptiveness × Post × ESG Fund	0.054**				0.056**	
	(0.021)				(0.016)	
Prescriptiveness × Post × ESG Fund (ES Family)		0.066**				0.066**
		(0.032)				(0.032)
Prescriptiveness × Post × ESG Fund (Non-ES Family)			0.017			
			(0.472)			
Prescriptiveness × Post × ESG Fund (Large-ES Family)				0.091**		
				(0.011)		
Prescriptiveness × Post	-0.100***	-0.100***	-0.096***	-0.097***	-0.104***	-0.103***
	(0.004)	(0.004)	(0.005)	(0.005)	(0.002)	(0.002)
ESG Fund	0.245***				0.242***	
	(0.000)				(0.000)	
ESG Fund (ES Family)		0.402***				0.396***
		(0.000)				(0.000)
ESG Fund (Non-ES Family)			-0.121***			
			(0.000)			
ESG Fund (Large-ES Family)				0.529***		
				(0.000)		
Prescriptiveness	-0.025	-0.025	-0.026	-0.026	-0.019	-0.019
	(0.363)	(0.354)	(0.344)	(0.342)	(0.474)	(0.462)
Observations	528153	528153	528153	528153	528153	528153
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Proponent-Type FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes	Yes	Yes
Fund Controls	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.121	0.130	0.115	0.122	0.124	0.133
F Statistic	162.434	196.758	118.711	319.507	155.840	185.752

Note: *** p≤0.01, ** p≤0.05, * p≤0.10. P-values are reported in parentheses. The dependent variable in specifications (1) to (4) relates to the variable “Binary Fund Vote”, which takes on a value of 1 when a specific fund votes in favor of a proposal, and 0 for all other outcomes. In specifications (5) and (6), the dependent variable relates to the variable “Ordered Fund Vote”, which is assigned a value of 1 for a “yes” vote, 0.5 for an “abstained” vote, and 0 for all other outcomes. We drop all anti-ESG proposals which we identify in Section 6.3. We suppress reporting of the constant term, firm-proposal controls, fund controls, as well as all interaction terms in specification (5) (e.g., $Post_t \times FundCat_m$) which are unreported in this Table. Standard errors are clustered at the meeting-level.

Table 12: Prescriptiveness Post Treatment: Heterogeneity amongst Big Three and Active Funds

	Binary Fund Vote				Ordered Fund Vote			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Prescriptiveness × Post × Big Three	0.052				0.051			
	(0.107)				(0.111)			
Prescriptiveness × Post × Blackrock		0.016				0.017		
		(0.679)				(0.668)		
Prescriptiveness × Post × Active (Measure 1)			-0.058***				-0.057***	
			(0.005)				(0.007)	
Prescriptiveness × Post × Active (Measure 2)				-0.046***				-0.046***
				(0.007)				(0.006)
Prescriptiveness × Post	-0.104***	-0.098***	-0.092***	-0.088***	-0.107***	-0.101***	-0.096***	-0.091***
	(0.005)	(0.006)	(0.006)	(0.008)	(0.003)	(0.004)	(0.004)	(0.005)
Big Three	-0.263***				-0.268***			
	(0.000)				(0.000)			
Blackrock		-0.297***				-0.282***		
		(0.000)				(0.000)		
Active (Measure 1)			-0.055***				-0.054***	
			(0.000)				(0.000)	
Active (Measure 2)				-0.047***				-0.052***
				(0.000)				(0.000)
Prescriptiveness	-0.027	-0.027	-0.026	-0.026	-0.022	-0.021	-0.020	-0.019
	(0.355)	(0.340)	(0.336)	(0.328)	(0.445)	(0.444)	(0.446)	(0.442)
Post × Big Three	0.024				0.021			
	(0.225)				(0.285)			
Observations	528153	528153	528153	528153	528153	528153	528153	528153
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Proponent-Type FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fund Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.140	0.132	0.115	0.115	0.144	0.133	0.117	0.117
F Statistic	112.885	116.944	116.907	117.140	112.033	114.169	115.126	116.114

Note: *** p≤0.01, ** p≤0.05, * p≤0.10. P-values are reported in parentheses. The dependent variable in specifications (1) to (4) relates to the variable "Binary Fund Vote", which takes on a value of 1 when a specific fund votes in favor of a proposal, and 0 for all other outcomes. In specifications (5) through (8), the dependent variable relates to the variable "Ordered Fund Vote", which is assigned a value of 1 for a "yes" vote, 0.5 for an "abstained" vote, and 0 for all other outcomes. We drop all anti-ESG proposals which we identify in Section 6.3. We suppress reporting of the constant term, firm-proposal controls, fund controls, as well as all interaction terms in specification (5) (e.g., $Post_t \times FundCat_m$) which are unreported in this Table. Standard errors are clustered at the meeting-level.

9 Online Appendix

9.1 Estimating Changes in E&S Support Using Synthetic Difference-in-Differences

As mentioned in Section 1, determining the impact of the SEC's 2021 Guidance on E&S proposal support is challenging due to the ostensible absence of a suitable "control" group for establishing a counterfactual scenario without the treatment.[60] Given the consistently high levels of support for governance proposals over time (see Figure 1), using governance proposals as a control group would violate the "parallel trends" assumption required for a Difference-in-Differences (DID) analysis. Indeed, the identifying assumption behind a DID analysis requires that the differences between control and treatment groups remain constant over time in the absence of the treatment. If the control and treatment groups were to have different pre-existing trends, treatment effects which are estimated from such models could be biased, as changes in the outcome variable that are due to pre-existing trends might be incorrectly attributed to the treatment (Rambachan and Roth (2023)).

As an initial strategy to address these concerns, we employ the Synthetic Difference-in-Differences (SDID) methodology introduced by Arkhangelsky et al. (2021). This approach is consistent with the plausible assumption that the 2021 Guidance, which determines the excludability of proposals based on whether they raise "significant policy issues," would primarily impact the level of prescriptiveness of E&S proposals while having minimal or no effect on governance proposals (Tallarita (2022); Gibson-Dunn (2022); Gibson-Dunn (2023)).[61] As Arkhangelsky et al. (2021) point out, the SDID methodology is particularly effective in addressing this issue because it allows for treated and control units to trend on entirely different levels before a regulatory shock. Specifically, the presence of unit-fixed effects in SDID allows for the matching of treated and control units based on pre-treatment trends rather than requiring similarity in both pre-treatment trends and levels. Since SDID necessitates the use of a balanced panel, we aggregate the data at the firm-year level, designating a firm as treated if more than 50% of the proposals it faces in a given year relate to E&S proposals relative to governance proposals.[62] All other firms are labeled

[60]Note that the caveats outlined in Section 4.3.2 (regarding the fact that each proposal is observed only once in the dataset) also apply to this setting. To the extent that these assumptions fall short of supporting causal inference, we refrain from asserting causality in this section.

[61]To the extent that these assumptions fall short of supporting causal inference, we refrain from asserting causality in this section.

[62]To ensure the resulting panel dataset is balanced, we backfill and forward-fill all missing values with the most recent available data. Following the approach of Arkhangelsky et al. (2021), we also assume that once a unit is treated,

as untreated. We then proceed to estimate the following parameters:

$$\left(\hat{\tau}^{\text{SDID}}, \hat{\mu}, \hat{\alpha}, \hat{\beta}\right) = \underset{\tau,\mu,\alpha,\beta}{\operatorname{argmin}} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} \left(y_{it}^{res} - \mu - \theta_i - \nu_t - W_{it}\tau \right)^2 \hat{\omega}_i^{\text{SDID}} \hat{\lambda}_t^{\text{SDID}} \right\} \tag{7}$$

where μ is a constant, i indexes firms, and t indexes years. θ_i and ν_t represent firm and year fixed effects, respectively. The variable y_{it}^{res} represents the residuals obtained after regressing a measure of voting support y_{it} (e.g., the percentage of votes in favor relative to total votes cast) on a vector of firm-proposal controls, X_{it} (aggregated at the firm-year level), where $y_{it}^{res} = y_{it} - X_{it}\hat{\beta}$. W_{it} is a categorical variable indicating whether a firm is treated, determined by whether it faces more than 50% of E&S (or, alternatively, environmental or social proposals) relative to governance proposals in a given year. Finally, we harness the weights $\hat{\omega}_i^{\text{SDID}}$ and $\hat{\lambda}_t^{\text{SDID}}$ as defined by Arkhangelsky et al. (2021) to match treated and control units based on pre-treatment trends.

In column (1) of Table A12, we present the results for specification (7), where the dependent variable relates to the percentage of votes in favor relative to the total votes cast, controlling for firm-proposal controls, as well as firm and time fixed effects. Our analysis focuses on the estimated parameter $\hat{\tau}^{\text{SDID}}$, which represents the average treatment effect on the treated (ATT). Despite the consistently high levels of support for governance proposals both before and after the 2021 Guidance, our estimates indicate a 4.11% decrease in support for treated firms (those facing a higher proportion of E&S proposals) following the shock, compared to untreated firms. In column (2), we observe a similar decline in support, around 4.72%, when treated firms are defined as those exposed to a higher proportion of environmental proposals.[63] Column (3) presents a parallel analysis for firms facing a relatively higher proportion of social proposals, showing that these firms received approximately 3.77% more support than governance proposals prior to 2021.

In column (4), we estimate an alternative specification that designates firms facing a higher proportion of environmental proposals than social proposals as the treatment group and those facing a higher proportion of social proposals than environmental proposals as the control group. This approach is based on the assumption that commitments to environmental reform are more costly for firms than social reforms (Balogh and Yonker (2024)). Alternatively, this specification could be motivated by the assumption that the 2021 Guidance had a more significant impact on environmental proposals, as it explicitly stated that proposals "adopting timeframes or targets to address climate change" would no longer be excluded. Although environmental and social pro-

it remains exposed to the treatment indefinitely.

[63]Note that we exclude all social proposals from this analysis before aggregating the data to the firm-year level. Similarly, we exclude all environmental proposals for a similar analysis in column (3).

posals have followed similar trends over time, we observe a 4.97% decrease in support for firms classified as treated compared to those in the control group. Finally, columns (5) and (6) report the results of specification 7 where different variants of the dependent variable are used: column (5) considers the percentage of affirmative votes relative to outstanding shares, and column (6) looks at the percentage of affirmative votes as all votes for and against the proposal. The findings in these last two columns closely mirror those in column (1).

To visually illustrate how support for E&S proposals has changed over time following the creation of a synthetic control that matches the parallel trends of treated units (while accounting for a comprehensive set of firm-proposal controls), we present a time series of the weighted voting support for E&S and governance proposals at the firm-year level in Figure A1.[64] As shown in the figure, the weights used in the SDID methodology allow treated and control units to follow different levels prior to the 2021 Guidance. Panel A demonstrates that firms with a higher proportion of governance proposals consistently received significantly more voting support than firms with a higher proportion of E&S proposals, both before and after the regulatory shock in 2021. Panels B and C provide similar illustrations, where treated firms are defined as those facing a greater proportion of environmental or social proposals, respectively. Figure A1 also highlights the negative ATTs reported in Table A12. Specifically, after 2021, firms exposed to a relatively larger number of E&S proposals (Panel A) experienced a much sharper decline in voting support compared to control firms exposed to more governance proposals. Similar trends are evident in Panels B and C.

9.2 BERT: Supervised Machine Learning

9.2.1 Data and Model Set-up

Supervised machine learning methods inherently rely on a "labeled" or "training" dataset to guide the classification of new data. As explained in Section 4.2, we utilize a specific subset of contested proposals from our dataset, specifically those involving Rule 14a-8(i)(7) prior to the 2022, as a foundational basis for labeling the proposals in our dataset. When resolving a disputed proposal, the SEC either supports the company's management by agreeing to exclude the proposal, or backs the proposal's advocate by denying the exclusion request. The key assumption made here is that the SEC, guided by the "ordinary business exception" in Rule 14a-8(i)(7), tends to exclude

[64] We harness a measure of voting support where the dependent variable is the percentage of votes in favor relative to the total votes cast.

proposals that exhibit a greater level of prescriptiveness. This assumption, as detailed in Section 4.2, is motivated by the observation that prescriptive proposals often venture into the details of a company's day-to-day business operations, which usually fall under the domain of the company's board or management (Bainbridge (2016)).

Focusing on the subset of all contested proposals from 2001 to 2021, we assign all excluded proposals under Rule 14a-8(i)(7) (favoring management) with a prescriptiveness indicator value of 1, and all precluded proposals under Rule 14a-8(i)(7) (favoring the proponent) with an indicator value of 0. This results in the creation of a "training and validation" set—what we term our "training dataset"—consisting of 1,158 proposals. In constructing our training dataset, we weigh the advantage of providing the algorithm with a larger pool of examples to learn from against the risk of incorrect classifications due to shifts in SEC policy.

As is common in the literature, we randomly split our 1,158 proposals into a training set (80%) and a validation set (20%) (see Joseph (2022)).[65] The validation set includes proposals drawn randomly across different years, allowing us to test the model on data it has not encountered during training. This random split ensures robust out-of-sample validation, helping to confirm that the algorithm generalizes well to new, unseen data.

Harnessing our training algorithm, we then classify all E&S proposals that were either uncontested or withdrawn (including settlements) between 2018 and 2021, as well as those from 2022 to 2023, employing Google's BERT (Bidirectional Encoder Representations from Transformers) algorithm. This approach is consistent with the methodologies used in Michaely et al. (2023), Rajan et al. (2023), and Liu and Lapata (2019). The BERT model is pre-trained on approximately 3.2 billion words from Wikipedia and 11,000 books from a variety of genres, which allows it to generate a large number of embeddings (numerical weights assigned to words) which are context specific (Liu and Lapata (2019)).

The BERT algorithm employs multiple steps to achieve classification, each of which will be outlined below:

1. **Tokenization:** Each text sample (i.e., an individual proposal) in our labeled training set is first broken down into smaller components called tokens (words or subwords). These tokens serve as the basic representation of the input text. Each token is then converted into an embedding vector, which assigns numerical weights to various aspects of the token, including its type, position in the text, and any associated segments.

[65] Accordingly, note that the "training set" is a distinct subset of the "training dataset", which comprises of both the "training set" and the "validation set".

2. **Classification Layers and Fine-Tuning:** The BERT model we use is already pre-trained on large text corpora (about 3.2 billion words from Wikipedia and 11,000 books across various genres), giving it a strong foundation for understanding language. However, to adapt it to our specific classification task, we add a new classification layer initialized with random weights. During fine-tuning, both the pre-trained layers of BERT and this newly added classification layer are updated based on the labeled training data. The goal of this process is to minimize classification loss (in this case, binary cross-entropy), ensuring that the model can accurately classify proposals as either prescriptive or not.

3. **Validation:** To ensure that the model generalizes well and avoids over-fitting to the training data, we reserve 20% of the labeled proposals as a validation set. This set is used exclusively to test the model's performance during the fine-tuning process. By evaluating the model's accuracy on this unseen validation data, we can detect any signs of over-fitting and adjust the model as needed. Standard metrics like accuracy, precision, recall, and F1-score are used to assess how well the model performs on the validation set.

4. **Prediction on Unlabeled Data:** Once the model has been fine-tuned and validated, we apply it to entirely new, unlabeled data. This includes all uncontested and withdrawn proposals in our dataset from 2018 to 2021, as well as all proposals from 2022 to 2023. The model processes these new proposals through the BERT architecture and classification layer, generating a prediction score (typically a probability) indicating the likelihood of whether the SEC would have excluded the proposal prior to 2021 under the "ordinary business exception". As is standard in the literature, we assign a value of 1 for probabilities ≥ 0.5 and a value of 0 for probabilities < 0.5.

In training our classification model, we safeguard the training process against overfitting by combining early stopping with model checkpointing. Early stopping halts training when the validation loss fails to improve for a predetermined number of epochs, thereby avoiding needless iterations. Model checkpointing stores the set of weights that achieves the lowest validation loss, so the final model always corresponds to the best performance observed during training.

Consistent with common practice, we pre-commit to the principal hyperparameters in our baseline model: a maximum of 10 training epochs, an initial learning rate of 3e-5, and a patience length of 3 epochs for early stopping. These settings mirror those widely used in the transformer literature (see, for example, Mosbach et al. (2021)).

9.2.2 Model Performance (Baseline Model)

Since the training dataset reflects SEC decisions on whether a shareholder proposal may be omitted from the ballot, the principal descriptive statistic is the label distribution. In this sample, 56.8% of proposals are coded as prescriptive (i.e., the SEC barred them from a vote), while 43.2 percent are classified as non-prescriptive.

We report the performance metrics of our baseline model in Table A13. As the Table illustrates, our baseline model has a binary accuracy of 0.741, precision of 0.880, recall of 0.669, F1 score of 0.760, and an area under the ROC curve (AUC) of 0.835. Although an accuracy of 0.741 may appear modest, the figure must be interpreted in light of three considerations: the small size of the training set, the performance of alternative models, and the model's heterogeneous predictive strength across proposals. In particular, the classifier discriminates well when predicted probabilities cluster near either 0 or 1, which is where most observations lie.

In Figure A2, we plot the Receiver Operating Characteristic (ROC) curve for our benchmark single-fold model, the weight set used throughout the study. The curve traces how the true-positive rate rises against the false-positive rate as the classification threshold is swept from 0 to 1. A 45-degree line indicates random guessing, while a curve that traces the upper-left border (an inverted "L") marks perfect discrimination vis-à-vis the validation set. The shape observed here shows strong separation when predicted probabilities cluster near the extremes of 0 or 1. Because evidence presented elsewhere in the paper (see Section 4.3.4) indicates that our treatment effects are driven chiefly by highly prescriptive proposals—observations the model almost always assigns probabilities close to 1—the classifier's predictive utility exceeds what the overall accuracy figure alone might suggest.

9.2.3 Type I vs Type II Errors (Baseline Model)

Misclassification in our supervised model biases the estimated voting effect by pulling the coefficient toward zero, irrespective of whether the errors are Type I or Type II. A Type I error places a non-prescriptive (high-support) proposal in the treated group, raising the treated mean and narrowing the treated–control gap. A Type II error places a prescriptive (low-support) proposal in the control group, lowering the control mean and again narrowing the gap. Either imbalance therefore reduces the absolute magnitude of the Prescriptiveness × Post coefficient, although the amount of attenuation depends on the relative frequencies of the two error types. This attenuation result for binary treatments is an extension of Aigner et al. (1973)'s analysis of misclassified binary regressors.[66]

To evaluate the potential bias introduced by misclassification, we conduct a sensitivity analysis, demonstrating that any imbalance between Type I and Type II errors is of second-order importance. Using the validation set—a random 20% sample of proposals defined in Sections 9.2.1 and 9.2.2—we first compute Type I and Type II error rates under the baseline rule that classifies a proposal as prescriptive when its predicted probability is at least 0.50. We then recalculate these error rates with lower decision thresholds of 0.40 and 0.30; a proposal is labeled prescriptive whenever the BERT model assigns it a probability that meets or exceeds the chosen cut-off. Lowering the threshold increases Type I errors because more non-prescriptive proposals are incorrectly labeled as prescriptive, while Type II errors fall. We repeat the same procedure with higher thresholds of 0.60 and 0.70. Raising the threshold has the opposite effect, increasing Type II errors by classifying a larger share of prescriptive proposals as non-prescriptive, while reducing Type I errors.

Table A14 summarizes these classification errors at each probability threshold, reporting the Type I error rate, the Type II error rate, and both error-rate ratios (Type I : Type II and Type II : Type I). Because each threshold generates a distinct set of observations coded as "prescriptive," we re-estimate the baseline regression (column (1)) from Table 5 under each coding scheme. Table A15 reports the coefficients on the Prescriptiveness × Post interaction—estimated with the full set of controls, fixed effects, and clustered standard errors.

Collectively, Tables A14 and A15 show that misclassification imbalance does not materially affect our findings. At the baseline threshold of 0.50, the Type II-to-Type I error ratio is 1.03, indicating a balanced pattern of errors. Varying the threshold alters this ratio substantially, shifting it from 0.36 to 4.12. Despite this change, the coefficient on the Prescriptiveness × Post term remains consistently negative, between −7.26 and −9.54%, and is statistically significant at the 1% level across all specifications. This stability suggests that the relative imbalance of Type I and Type II errors is a second-order concern for our headline results.

Figure A3 plots the Prescriptiveness × Post coefficients against the corresponding Type II-to-Type I error ratios. The estimates are tightly clustered, showing declines in voting support of roughly 7.26 to 9.54% across all thresholds. Because the coefficients vary only modestly across a broad range of error ratios (from 0.36 to 4.12), the figure reinforces the robustness of our baseline results.

To see why the imbalance between Type II and Type I errors is not a key driver of our results,

[66] However, as Lewbel (2007) notes, the bias toward zero persists only under symmetric, outcome-independent misclassification. For a contemporary analysis of such biases created by machine-learning algorithms, see Zhang (2021).

consider Figure A2, which plots the receiver-operating-characteristic (ROC) curve for the benchmark model evaluated with the 0.50 threshold. Although the classifier is not perfect, its predictions are highly reliable whenever the estimated probability that a proposal is prescriptive lies near 0 or 1. Indeed, 84.38% of the out-of-sample observations are assigned predicted probabilities of prescriptiveness that are either ≥ 0.70 or ≤ 0.30.[67] Because only a small fraction of proposals fall in the ambiguous middle range, changing the classification threshold affects relatively few observations; those proposals lack sufficient weight to materially shift the estimated treatment effect.

9.2.4 Model Performance (K-Fold, LLMs, and Hyperparameter Optimization)

9.2.4.1 K-Fold Validation

To assess the robustness of the baseline classifier, we conduct a k-fold cross-validation exercise. The full training dataset is randomly partitioned into k equal folds; the model is trained on $k - 1$ folds and evaluated on the remaining fold, which we label the validation set. We use $k = 5$, thereby maintaining an 80 : 20 split between training and validation data in each iteration. Table A13 reports the Accuracy, Precision, Recall, F1 Score, and AUC for each of the four fitted models.[68]

The cross-validation results align closely with the performance of the single baseline model, indicating that the paper's main results are not driven by idiosyncratic data splits. The baseline achieves an Accuracy of 0.741, Precision of 0.880, Recall of 0.669, F1 Score of 0.760, and AUC of 0.835. In comparison, averaging across the four folds yields an Accuracy of 0.721, Precision of 0.865, Recall of 0.661, F1 Score of 0.789, and AUC of 0.835. Figure A4 plots the individual ROC curves for each fold, with their mean ROC highlighted in blue.

9.2.4.2 Large Language Models

Large language models (LLMs) are not, by design, conventional supervised learning algorithms that ingest a fixed labelled corpus before generating predictions solely from those fitted param-

[67]Note that this does not indicate a particular level of binary accuracy. Probabilities of prescriptiveness are computed for the entire set of "other" E&S proposals in our dataset—including those uncontested or withdrawn between 2018 and 2021, as well as all proposals from 2022 to 2023—whereas binary accuracy is reported only for the validation sample, a randomly drawn 20 percent subset of the training data.

[68]We omit the baseline model because it uses one of the folds as its held-out validation set. In other words, the baseline model is one of the "5 folds" used in the k-fold validation exercise.

eters. Recent methodological work, however, indicates that advancements in model architecture have enabled LLMs to execute supervised classification tasks with competitive accuracy (Gunel et al. (2020)).

The literature now relies on two main strategies for using LLMs to classify text corpora. The first approach is often known as "fine-tuning". This process involves updating an LLM's internal weights—starting, for example, from OpenAI's GPT-4.1-2025 and retraining on our shareholder-proposal corpus—to embed domain-specific knowledge. The second approach relates to "prompt-based" classifications. Here, the model receives a carefully engineered prompt instructing it to label each proposal as prescriptive or non-prescriptive. This method is constrained by the model's context window, which currently tops out at roughly 100,000–200,000 tokens (Chen et al. (2023)). Because the weights cannot be revised across sequential prompts, the LLM cannot "see" the entire training set at once. To mitigate this constraint, researchers commonly employ "few-shot prompting": a small, randomly selected set of labeled exemplars is embedded in a single, length-optimized prompt so that the entire input remains within the model's context window (Wang et al. (2020); Parnami and Lee (2022)).

To evaluate whether an LLM could outperform our dedicated BERT classifier, we replicated both methods on the identical train–validation split described in Sections 9.2.1 and 9.2.2. For the few-shot approach, we prompted ChatGPT-4.1-2025 with a length-optimized instruction set (Figure A5), supplemented by 30 randomly selected, labeled examples from the training data—just enough to remain within the model's context window. Each run required roughly 15 million input tokens, making the exercise substantially more resource-intensive than fine-tuning BERT. As reported in Table A13, the prompting method nevertheless yielded weaker metrics than the BERT baseline: an Accuracy of 0.634, Precision of 0.672, Recall of 0.847, and F1 Score of 0.749.[69] These results accord with prior evidence that zero-shot and few-shot prompting methods may lag behind purpose-built supervised models in classification tasks (see, for example, Wang et al. (2024)).

Turning to the fine-tuning approach, we retrained OpenAI's GPT-4.1-2025 on the identical shareholder-proposal corpus used for our baseline BERT model. Each run processed roughly 11 million tokens, so the procedure was likewise computationally demanding.[70] As reported in Table A13, the fine-tuned GPT model delivers higher recall and F1 scores relative to the BERT

[69]The notably high recall likely reflects our prompt design in Figure A5, which instructs the LLM to return a value of 0 whenever its classification confidence is low, an approach consistent with best-practice recommendations in the literature (Fisch et al. (2022); Wen et al. (2025)).

[70]Section 9.2.4.4 outlines some of these trade-offs involved in choosing between LLM-based classifiers and our BERT model.

baseline, although its binary accuracy and precision are lower.[71]

9.2.4.3 Hyperparameter Optimization

As explained in Section 9.2.1, the baseline BERT model fixes its hyperparameters ex ante, with a maximum of 10 epochs, learning rate 3e-5, and an early-stopping patience of 3 epochs. Because our dataset is modest in size, we did not expect substantial loss reductions from hyperparameter optimization. Nevertheless, as a robustness check, we carried out a limited search, retaining the same training set, but reallocating the overall training dataset to an 80/10/10 train–test–validation split.

We conducted 15 trials, each aimed at minimizing cross-entropy loss on the training set. The best specification selected a maximum of 6 epochs, a learning rate of 2.427e-05, and an early-stopping patience of 4 epochs, yielding a validation loss of 0.5202 versus 0.5625 for the baseline. The tuned model's performance, reported in Table A13, has an Accuracy of 0.7241, Precision of 0.8246, Recall of 0.6812, F1 Score of 0.7857, and AUC of 0.7757. In our view, these performance metrics represent only marginal differences from the baseline and do not constitute a meaningful improvement.

9.2.4.4 Robustness of Baseline Results

We regard the proprietary nature of commercial LLM weights and the substantial financial costs of training/fine-tuning them as tipping the cost–benefit calculus for this study in favor of a lighter-weight model such as BERT. Moreover, robustness checks confirm that our headline findings are insensitive to the choice of classifier.

Table A16 replicates the baseline specification (column (1)) from Table 5, successively replacing the prescriptiveness measure with outputs from each alternative model. Row (1) of Table A16 restates the original estimate— -8.48 with firm and year fixed effects—reported in Table 5. Row (2) substitutes the "median" BERT model from the five-fold exercise (model 3 with its distinct validation set); here, the Prescriptiveness × Post coefficient remains negative and significant at the 1 percent level. Row (3) uses prescriptiveness probabilities derived from the fine-tuned OpenAI GPT-4.1-2025 model, Row (4) employs the few-shot prompting classifier based on OpenAI's

[71] We do not report AUC for the LLM models because they output only binary labels—prescriptive or non-prescriptive—without the continuous probabilities needed to trace an ROC curve.

ChatGPT-4.1-2025 model, and Row (5) applies predictions from the BERT model with optimized hyperparameters. Despite the differences in training splits, hyperparameter settings, and initial weights, every specification yields a negative coefficient on Prescriptiveness × Post that is statistically significant at the 1% level, underscoring the robustness of our main results and, by extension, the analyses that build upon them throughout the paper.

9.3 Unsupervised Machine Learning: Topic Modeling

While the initial supervised machine learning approach offers a preliminary means of identifying prescriptive proposals, it may not encompass the full range of such proposals. This limitation arises because the training dataset is drawn from proposals clustered near the SEC's decision boundary, which can misclassify proposals whose language or context lies far outside that training manifold (i.e., the surface in embedding-space traced out by the training points in machine learning).[72] Essentially, the firms, proponents, and investors behind contested proposals may exhibit attributes the model never encountered during training—attributes absent from uncontested or withdrawn proposals—so the classifier must extrapolate, increasing the likelihood of misclassification.[73]

To formalize the source of misclassification, let $\phi : X \to \mathbb{R}^d$ be an embedding function. The set of n training embeddings is $\mathcal{S}_{\text{train}} = \left\{\phi(x_i)\right\}_{i=1}^{n}$. For any new proposal $x \in X$, we define its out-of-distribution (OOD) distance as the Euclidean gap to its nearest neighbor in the training set:

$$d_{\text{OOD}}(x) = \min_{z \in \mathcal{S}_{\text{train}}} \left\|\phi(x) - z\right\|_2$$

Classical margin theory shows that, while test inputs remain on the same manifold as the training data, their misclassification probability decreases as their geometric margin to the learned deci-

[72] Additionally, the training dataset employed is relatively small, which heightens the likelihood of misclassifying proposals.

[73] Unsupervised machine learning transformers, such as Mistral 7B, are particularly well-suited for tasks that involve uncovering latent structures or relationships in data without the need for labeled examples. These models utilize self-supervised learning objectives to generate high-dimensional embeddings that effectively capture the contextual nuances of text. This makes them ideal for exploratory tasks like clustering or dimensionality reduction, where the aim is to identify patterns/groupings that are not predefined. In contrast, supervised transformers like BERT are optimized for leveraging labeled data in task-specific fine-tuning, enabling precise and interpretable predictions. However, their reliance on labeled data can limit their applicability in settings where such data is scarce or unavailable. We selected Mistral 7B for our analysis as it was one of the leading models on HuggingFace's leaderboards at the time of writing.

sion boundary widens (Belkin et al., 2006). As a mathematical illustration of this extrapolation risk, our central assumption concerns the ideal behavior the classifier aims to learn: we assume the Bayes-optimal decision function is L-Lipschitz. This means that for the true, perfect classifier, the change in its output probability is mathematically bounded by a constant, L, multiplied by the change in its input. We also define ϵ as the classifier's baseline training error. Based on these assumptions, the "trust radius" around the training data can be defined as $r = \epsilon / L$.

For any new proposal x, the probability of error is bounded by its out-of-distribution distance, $d_{\text{OOD}}(x)$. Once x is outside the trust radius (i.e., $d_{\text{OOD}}(x) > r$), the minimum possible error increases linearly with its distance from the known data. This relationship can be expressed as:

$$P_{\text{error}}(x) \geq \max\left(0, L \cdot d_{\text{OOD}}(x) - \epsilon\right)$$

This expression predicts that once the distance-based risk ($L \cdot d_{\text{OOD}}(x)$) outweighs the baseline error (ϵ), the error bound grows proportionally with the distance—a behavior characteristic of dataset shift (Quiñonero-Candela et al., 2022).

Empirical studies further show that the model's reported "confidence percentages"—the probabilities obtained after rescaling raw class scores to sum to one—remain spuriously high even in this high-risk regime, leading to systematically over-confident predictions (Hendrycks and Gimpel, 2016). Because the SEC-labelled corpus spans only a narrow linguistic region, proposals whose wording, issuer profile, or proponent type diverge markedly from that region have large d_{OOD} values and are therefore susceptible to misclassification.

To mitigate the risk of misclassification, we implement a "Topic Modeling" strategy in a secondary step to identify prescriptive proposals (Grootendorst (2022)). This method seeks to identify groups of proposals that share common themes related to "prescriptive content", like the implementation of particular policies, thereby offering a more nuanced insight into the proposals' characteristics. Given the probable differences in content between environmental and social proposals, we apply our topic modeling algorithms to each set of environmental and social proposals independently. Like the techniques employed in Section 9.2, our algorithm utilizes a series of steps to discern distinct topics, which are detailed as follows:

1. **Embedding**: In this step, we employ an embedding model from Mistral (Mistral 7B) which assigns numerical weights to the words, phrases, and sentence structures in our proposals (Jiang et al. (2023)). Mistral 7B's strength lies in its pre-training on an extensive range of internet-scale data, encompassing 7 billion parameters. This vast foundation enables Mistral

7B to assign context-specific weights to textual elements, enhancing the model's ability to represent the underlying semantic relationships within the data.

2. **Dimensionality Reduction**: Subsequently, we utilize a UMAP (Uniform Manifold Approximation and Projection) algorithm to reduce the textual data's dimensionality, aiming to retain the most critical attributes of environmental or social proposals. UMAP applies mathematical principles from topology to condense complex, high-dimensional data, such as Mistral 7B embeddings. UMAP is based on the concept that data points can be represented as a connected graph in high-dimensional space, and seeks to preserve these connections when projecting the data into a lower-dimensional space. The algorithm uses Riemannian geometry to adjust local metrics, ensuring that dense regions do not dominate the layout. This allows UMAP to maintain the inherent topological features of the dataset.

3. **Vectorization**: During this stage, we implement a filtering process to remove common stop-words — words that are frequent in the language but typically carry little information relevant to the analysis (e.g., "the", "is", "and"). Eliminating these words focuses the analysis on more meaningful, content-specific words, thereby enhancing the effectiveness of text-based models.

4. **Clustering**: We use the HDBSCAN (Hierarchical Density-Based Spatial Clustering of Applications with Noise) clustering algorithm to facilitate the identification of distinct proposal clusters. HDBSCAN starts by modifying the dataset through a mutual reachability distance that incorporates density into the distance measure between data points, thereby accentuating dense regions. A minimum spanning tree is then constructed to highlight these dense areas, from which a hierarchy of clusters is derived. This hierarchy is condensed based on cluster stability across different density levels, with the final step being the pruning of this condensed tree to extract significant clusters. This approach allows HDBSCAN to detect meaningful clusters while disregarding less significant ones as noise.

The application of topic modeling algorithms to our dataset uncovers specific topic clusters that align with characteristics identified by the SEC, legal practitioners and institutional investors. In Figure A6, we provide a two-dimensional representation of environmental proposals.[74] Amongst these proposals, there emerges a clear cluster advocating for companies to set "time-bound" or "company-wide" emissions targets (Topic 8: "GHG Emissions Management Goals Adoption policy in Corporations").[75] Another group of proposals calls for companies to "adopt a [specific] policy"

[74]Note that we repeat the same process for all social proposals.

[75]For instance, one of these resolutions provides that "Resolved: Shareholders request Illinois Tool Works, Inc. (ITW) adopt time-bound, quantitative, company-wide, science-based targets for reducing greenhouse gas (GHG)

(or similar expressions like "implementing," "adopting," or "committing to a policy"), such as phasing out fossil fuel exploration and development (Topic 23: "Finance Commitment to Net Zero Emissions by 2050" and Topic 28: "Fossil Fuel Phase-Out Policies by Major Banks").[76] we identify these clusters of proposals that are ostensibly "prescriptive" in nature, before assigning these proposals with a prescriptiveness indicator of 1.[77]

9.4 Selection Bias Models

There are two types of selection bias that could potentially distort the observed treatment effects. The first relates to sample selection bias, characterized by endogenous or non-random selection into the sample. For example, the SEC's 2021 Guidance might lead company management to refrain from contesting proposals. Specifically, Bebchuk et al. (2020) have noted that larger firms face more severe reputational risks if a contested proposal is decided in favor of the proponent. Should these managers alter their actions in response to the SEC's 2021 Guidance, this could lead to variations in voting support that are not directly tied to changes in the prescriptiveness of E&S proposals.

To mitigate the possibility of sample selection bias, we employ a Heckman selection model (Heckman (1979)), which accounts for the potential selection bias arising from proposals that were either excluded by the SEC or withdrawn by the proponent (Zytnick (2022); Brav et al. (2024)).[78] To implement the Heckman selection model, we first estimate the probit specification:

$$\Pr(S_n = 1) = \Phi\left(Z_{in}\gamma\right)$$

where S_n is a binary variable indicating whether a proposal is excluded/withdrawn ($S_n = 0$) or

emissions, consistent with the goals of the Paris Climate Agreement, and report annually, at reasonable cost and omitting proprietary information, on its plans and progress towards achieving these targets."

[76]For instance, one of these resolutions provides that "Resolved: Shareholders request that the Board of Directors adopt a policy for a time-bound phase-out of BAC's lending and underwriting to projects and companies engaging in new fossil fuel exploration and development."

[77]Notably, a topic-modeling approach reduces manual workload: rather than read all 1,891 proposals, an analyst need examine only the model's 29 clusters—each auto-labeled by an LLM from its top keywords. For example, proposals urging boards to align policies with Net-Zero targets (highlighted as prescriptive in SEC SLB 14L) appear as a single, easily identified cluster.

[78]Note that shareholder proposals must be contested by firm management prior to being excluded by the SEC. Bebchuk et al. (2020) detail a large number of firm and proponent characteristics that may influence whether a proposal is contested. These factors include the activist's stake, insider ownership, the target firm's share class structure, the firm's performance, historical success rates in past engagements, and the board's structure, among others.

whether it has proceeded to a vote ($S_n = 1$), Z_{in} is a vector of observed firm-proposal controls,[79] Φ represents the cumulative distribution function of the standard normal distribution, and γ is a vector of parameters to be estimated. Given the estimated parameters $\hat{\gamma}$, we compute the Inverse Mills Ratio (IMR) λ_n for each proposal:

$$\lambda_n = \frac{\phi(Z_{in}\hat{\gamma})}{\Phi(Z_{in}\hat{\gamma})}$$

where $\phi(\cdot)$ denotes the probability density function of the standard normal distribution.[80]

In a second stage, we estimate the outcome equation as per specification (2), adjusting for selection bias from the first step by including the IMR in the specification:

$$y_{ijktn} = \alpha + p_{ijktn}\beta + (p_{ijktn} \times Post_t)\gamma + X_{ijktn}\xi + \theta_i + \eta_j + \psi_k + \nu_t + \lambda_n\delta + \varepsilon_{ijktn}$$

where ε_{ijktn} is assumed to be independent from the error term in the first-stage probit specification.

The second form of selection bias concerns the possibility of being selected for treatment, which, in this context, relates to the prescriptiveness of a specific proposal. For instance, the 2021 Guidance might encourage proponents to direct more prescriptive proposals towards larger firms, under the belief that E&S proposals at such entities post-2021 have a better chance of proceeding to a vote (Era et al. (2021); Bebchuk et al. (2020)). In such a scenario, if the size of the firm influences voting support, neglecting to adjust for this bias could lead to biased estimates of the treatment effects.

To ameliorate these concerns, we calculate propensity scores for both prescriptive (the treatment group) and non-prescriptive proposals (the control group). These scores represent the likelihood of a proposal being assigned to the treatment group, based on an array of observable characteristics. By incorporating these propensity scores into the analysis, we aim to ensure that the treatment and control groups are essentially equivalent, with no systematic differences between them aside from their levels of "prescriptiveness". In a first stage, we estimate a probit regression to predict the likelihood of treatment:

[79]Since proposals are observed at a more granular level when compared to firms, proposals within the same firm will have similar firm-level controls. Further information about these variables can be found in Table A1.

[80]Note that the dot product of Z_{in} and $\hat{\gamma}$ is a scalar, so λ_n may be computed as a ratio of scalars.

$$\Pr(T_n = 1) = \Phi\left(Z_{in}\gamma\right)$$

where T_n is a binary variable indicating whether a proposal is prescriptive ($T_n = 1$) or whether it is non-prescriptive ($T_n = 1$), Z_{in} is a vector of observed firm-proposal controls,[81] Φ represents the cumulative distribution function of the standard normal distribution, and γ is a vector of parameters to be estimated. After deriving the propensity scores $\tau(Z_{in}) = \Phi(Z_{in}\hat{\gamma})$ by estimating the coefficients $\hat{\gamma}$ in the aforementioned specification, we follow Rosenbaum (1987) in computing an inverse probability of treatment weight (IPTW), where the IPTW w_n is defined as:

$$w_n = \frac{T_n}{\tau(Z_{in})} + \frac{1 - T_n}{1 - \tau(Z_{in})}$$

As the IPTW is equal to the inverse of the observation's probability of receiving the treatment, weighting a regression model with IPTWs allows for a specification that consistently estimates the true treatment effect (Joffe et al. (2004)).

9.5 Data Cleaning Procedures: Fund-Level Matching

As noted in Section 3.2, matching funds between the ISS dataset and the CRSP mutual fund database is non-trivial due to the lack of common identifiers in both datasets. To tackle this challenge, we begin by extracting the "Series Name" (i.e., fund name), the "Series ID," and the CIK (Central Index Key) linked to each N-PX identifier (obtained from the ISS dataset) from the SEC's EDGAR database. This process enables us to associate a CIK identifier, an N-PX file identifier, and a Series Name identifier from EDGAR with each voting record observed in the ISS dataset.

Unfortunately, these identifiers do not uniquely identify the voting records observed in the ISS dataset, as multiple funds are linked with each N-PX identifier documented in the ISS dataset. To address this issue, we perform fuzzy-matching between funds *within* an N-PX filing (identified by its "Series Name") and funds in the ISS dataset (identified by its "Fund Name") pursuant to a procedure indicated by Moskalev (2019), who matches funds with similar Levenshtein distances. For matches with Levenshtein distances of 3 or smaller (where 0 corresponds to a perfect match) we assume that we assume that funds in both datasets represent the same fund. Additionally, for any unmatched funds with a minimum distance of 4 or greater, we perform a hand matching

[81] Since proposals are observed at a more granular level when compared to firms, proposals within the same firm will have similar firm-level controls. Further information about these variables can be found in Table A1.

process (aided by sorting the within-N-PX filing funds based on their similarity to the fund in question). If no suitable match is identified, we drop the fund from our dataset.

Following the matching of ISS records with EDGAR data, we utilize the "CRSP CIK MAP" sourced from WRDS (Wharton Research Data Services). This dataset connects pairs of CIK identifiers and "Series ID" identifiers from EDGAR to CRSP fund numbers within the CRSP mutual fund database. This enables the matching of funds in ISS to their corresponding CRSP fund numbers. Finally, we merge fund characteristics, which may be associated with CRSP portfolio numbers, into this dataset using either CRSP fund or portfolio numbers along with the nearest record date provided in the ISS dataset. When dealing with portfolio numbers, we rely on CRSP's mapping between fund numbers and portfolio numbers to facilitate this matching process.

While the CRSP mutual fund database offers ownership data for certain funds in our dataset, it does not provide comprehensive coverage for all included funds. To ameliorate these gaps in coverage, we follow Brav et al. (2024) by incorporating data from the TR (Thomson Reuters) S12 mutual fund database. Notably, while the TR S12 database provides data at a quarterly frequency, the CRSP mutual fund database operates at a monthly frequency. Consequently, we include only March, June, September, and December holdings from the CRSP mutual fund database to create a comprehensive set of mutual fund holdings at the quarterly frequency.

To match the CRSP mutual fund database to the TR S12 database, we use the MFLINKS tables from WRDS to link each fund in the CRSP to the Thomson Reuters S12 data, using the provided link between a CRSP portfolio number and an S12 fund number. For funds in our dataset linked to an S12 fund number, we utilize ownership data from the TR S12 database. Conversely, for funds in our dataset lacking links to an S12 fund number, we rely on ownership data from the CRSP mutual fund database.

9.6 Alternative Indicators for Prescriptiveness

While we have outlined the need for a more complex approach to develop a measure of prescriptiveness for proposals in our dataset, it is reasonable to consider whether a simpler proxy could effectively capture a proposal's "prescriptiveness." For example, one might consider a binary indicator for whether a proposal was contested under Rule 14a-8(i)(7)'s "ordinary business exception." However, this simpler approach has significant drawbacks. As discussed in Section 4.2, proposals with greater prescriptiveness are less likely to reach a vote due to a higher likelihood of being excluded by the SEC. Therefore, limiting the analysis to only contested proposals could substantially under-represent highly prescriptive proposals. Furthermore, a large majority (63.68%) of the proposals in our dataset are uncontested. This is possibly because managers often have strong incentives to avoid contesting shareholder proposals, perhaps due to career concerns (Gantchev and Giannetti (2021); Matvos and Ostrovsky (2010)), the direct (e.g., legal expenses) and indirect (e.g., the impact on firm value from the uncertainty brought about by no-action letters) costs of contesting proposals (Matsusaka et al. (2021)), or the risks involved in challenging the recommendations of proxy advisory firms (Gantchev and Giannetti (2021)). These incentives to avoid contesting proposals might be further influenced by the SEC's 2021 Guidance, with firm managers potentially reluctant to oppose an outcome that would likely favor the proponent.

In Table A17, we present preliminary evidence supporting these claims by replacing the binary prescriptiveness indicator (p_{ijkt}) with a binary variable indicating whether a proposal was contested under Rule 14a-8(i)(7). While the coefficients on the interaction terms of interest are negative, they generally lack statistical significance across different specifications, including adjustments for industry fixed effects and the selection models outlined in Section 9.4. These results suggest that the potential negative association between the proposed proxy variable and the 2021 Guidance may not be robust.

9.7 Variance Decomposition

To justify the choice of fixed effects in our main analysis, we consider the hierarchical nature of our fund-level dataset (which includes various layers such as firms, industries, proponents, funds, and years) and follow Zytnick (2022) by decomposing the sources of variation in voting. To this end, we estimate the following specification:

$$y_{ijktnmc} = \alpha + \theta_i + \eta_j + \psi_k + \kappa_m + \delta_c + \nu_t + \varepsilon_{ijktnmc} \tag{8}$$

where α is a constant, i indexes firms, j indexes industries, k indexes proponent-types, m indexes funds, c indexes proponents, t indexes years, and n indexes proposals.[82] Meanwhile, $y_{ijktnmc}$ relates to the binary indicator, "Binary Fund Vote", while θ_i, η_j, ψ_k, κ_m, δ_c and ν_t represent firm, industry, proponent-type, fund, proponent, and year fixed effects, respectively.

Table A18 presents the outcomes from applying specification (8). In column (1) of Panel A, a baseline model incorporating only year fixed effects is estimated, yielding an expectedly low R^2 value

[82]Note that the notation utilized here deviates slightly from what was elucidated in specification (3).

of 0.022. Column (2) introduces proponent-type fixed effects alongside year fixed effects. Column (3) revisits the baseline model but includes industry fixed effects, whereas column (4) incorporates firm fixed effects. Column (5) adds proponent fixed effects, and column (6) includes fund fixed effects in the analysis. The findings illustrate that the between-fund variation in our dataset is substantial (as evidenced by an R^2 value of 0.286 in column (6)), suggesting that fund-related differences account for a significant portion of the variability in voting behavior when compared to other factors. In terms of the hierarchy of voting variation contributors, fund characteristics emerge as the most significant, followed by proponents, firms, industries, proponent-types, and years, in decreasing order of impact. In Panels B and C of Table A18, the same analysis is conducted for environmental and social proposals respectively, yielding results that closely mirror those from Panel A.

Employing fixed effects for a given set of entities (e.g., firms) involves a balancing act–while they account for unobserved, time-invariant attributes of the entities under study, they also remove variation between these entities. This may potentially weaken the statistical power of the analysis (Bai (2009)). Specifically, incorporating firm fixed effects eliminates the variation across firms, which could explain a considerable amount of the variation in voting behavior. In the analyses that follow, we consistently apply year and proponent-type fixed effects where applicable. However, to preserve some between-entity variation, we opt for industry fixed effects instead of firm fixed effects in various iterations of our findings.

Table A19 demonstrates the aforementioned tradeoffs by applying specification (4) to both environmental and social proposals independently.[83] Since environmental proposals represent a smaller share of the dataset than social proposals, preserving some degree of variation between entities is of increased importance. This necessity is underscored by the tendency of environmental proposals to target specific firms, notably within sectors like the fossil fuel industry (Tallarita (2022)).[84] In columns (1) and (3) of Table A19, the larger standard errors associated with the prescriptiveness coefficient, p_{ijktnm}, is evident when firm fixed effects are included. On the other hand, columns (2) and (4) show consistent and statistically significant negative coefficients at the 5% level when industry fixed effects replace firm fixed effects. For social proposals, as shown in columns (5) to (8), these methodological considerations do not appear to have a substantial impact.

[83] The dependent variable here relates to the binary indicator, "Binary Fund Vote".

[84] Both factors suggest that the introduction of firm fixed effects would consume relatively more of the available degrees of freedom, relative to social proposals.

9.8 Additional Tables and Figures

In this section, we provide additional tables and figures that supplement the main analysis presented in the paper. These tables offer alternative specifications and robustness checks that support our findings.

Table A1: List of Variables

Variable	Definition
Firm-Proposal Controls	
Log Mkvalt	Logarithm of the firm’s market capitalization
Firm Size	Logarithm of the firm’s total assets
Return on Assets (RoA)	Net income scaled by total assets
Leverage Ratio	Total debt scaled by total assets
Tobin’s Q	Market capitalization of equity plus total debt, divided by total assets
HHI	Sum of squared market shares for each firm within the firm’s 4-digit SIC industry code
Inst Own	% of firm equity owned institutional investors (i.e., 13F investors)
Inst HHI	Sum of squared ownership shares for each investor within the firm
Has No Action Letter Sought	Binary variable denoting whether the proposal is subject to a No Action Letter request
Fund Controls	
% of Security owned by Fund	Percentage of a firm’s outstanding shares owned by the fund
Security as % of Fund’s Total Assets	Percentage of the fund’s total assets represented by a firm’s holdings
Total Fund Assets	Total assets under management by the fund
Mgmt Fees	Management fees charged by the fund, including fee waivers/reimbursements
Expense Ratio	Ratio of total investment that shareholders pay for the fund’s operating expenses, which include 12b-1 fees
Turnover Ratio	Ratio of the fund’s aggregated sales/purchases of securities over a year, divided by the fund’s average annual total net assets
Index Fund	Binary variable coded as 1 if the fund is passively managed

Table A2: Panel Regressions of Voting Support on Prescriptiveness:
Heckman Selection and IPTW Variants

	Heckman Selection Models			IPTW Models		
	(1) E & S	(2) Environmental	(3) Social	(4) E & S	(5) Environmental	(6) Social
Prescriptiveness	-5.218***	-11.417***	-4.750***	-5.188***	-10.777***	-4.711***
	(0.000)	(0.006)	(0.000)	(0.000)	(0.007)	(0.000)
Inverse Mills Ratio	-84.427	-1139.885	97.707			
	(0.826)	(0.741)	(0.791)			
Observations	1856	461	1395	1082	205	831
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	No	No	No	No	No	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Proponent-Type FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.402	0.513	0.442	0.405	0.508	0.449
F Statistic				2.975	11.992	2.077
Chi-Square	25.097	19.140	20.728			

Note: *** p≤0.01, ** p≤0.05, * p≤0.10. P-values are reported in parentheses. The dependent variable in all specifications relates to the proportion of affirmative votes out of the total votes cast. In specifications (1) through (3), we implement the Heckman Selection model introduced in Section 9.4, while in specifications (4) through (6), we implement the Inverse Probability of Treatment Weighting (IPTW) model introduced in the same Section. we suppress reporting of the constant term and firm-proposal controls. Standard errors are clustered at the meeting-level.

Table A3: Changes in Prescriptiveness Post Treatment:
Supervised Machine-Learning Measure

	Votes For As % Votes Cast		Votes For As % Shares Out	Votes For As % Yes & No
	(1)	(2)	(3)	(4)
Prescriptiveness (Log-Transformed)	1.720	0.845	1.358	1.816
	(0.538)	(0.751)	(0.521)	(0.524)
Prescriptiveness (Log-Transformed) × Post	-8.701**	-7.115*	-6.497**	-8.830**
	(0.040)	(0.068)	(0.037)	(0.041)
Observations	1082	1180	1082	1080
Firm FE	Yes	No	Yes	Yes
Industry FE	No	Yes	No	No
Year FE	Yes	Yes	Yes	Yes
Proponent-Type FE	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.390	0.281	0.411	0.391
F Statistic	1.492	2.452	1.741	1.531

Note: *** $p \leq 0.01$, ** $p \leq 0.05$, * $p \leq 0.10$. P-values are reported in parentheses. In specifications (1) and (2), the dependent variable relates to the percentage of affirmative votes out of the total votes cast. For specification (3), it relates to the percentage of affirmative votes out of all outstanding shares. In specification (4), it relates to the percentage of affirmative votes out of the total of affirmative and negative votes. The independent variable of interest (i.e., our prescriptiveness measure) in this table is initially derived as a raw probability from our supervised machine-learning algorithms. We apply a logarithmic transformation to these values and normalize them to a 0-1 range to mitigate heteroskedasticity in the residuals. We suppress reporting of the constant term and firm-proposal controls. Standard errors are clustered at the meeting-level.

Table A4: Changes in Prescriptiveness Post Treatment:
Heckman Selection and IPTW Variants

	Heckman Selection Models			IPTW Models		
	(1) E & S	(2) Env	(3) Social	(4) E & S	(5) Env	(6) Social
Prescriptiveness × Post	-8.483***	-19.311***	-5.297**	-8.507***	-19.050***	-5.133**
	(0.000)	(0.009)	(0.030)	(0.000)	(0.001)	(0.026)
Prescriptiveness	-0.773	0.064	-1.995	-0.738	0.616	-2.046
	(0.654)	(0.992)	(0.279)	(0.639)	(0.903)	(0.217)
Inverse Mills Ratio	-92.632	-947.227	95.704			
	(0.798)	(0.776)	(0.813)			
Observations	1856	461	1395	1271	205	831
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	No	No	No	No	No	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Proponent-Type FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.414	0.541	0.446	0.417	0.539	0.453
F Statistic				4.149	12.251	2.278
Chi-Square	42.864	25.764	24.317			

Note: *** p≤0.01, ** p≤0.05, * p≤0.10. P-values are reported in parentheses. The dependent variable in all specifications relates to the proportion of affirmative votes out of the total votes cast. In specifications (1) through (3), we implement the Heckman Selection model introduced in Section 9.4, while in specifications (4) through (6), we implement the Inverse Probability of Treatment Weighting (IPTW) model introduced in the same Section. we suppress reporting of the constant term and firm-proposal controls. Standard errors are clustered at the meeting-level.

Table A5: Changes in Prescriptiveness Post Treatment
(without Anti-ESG Proposals)

	Votes For As % Votes Cast				Votes For As % Shares Out	Votes For As % Yes & No
	(1) E & S	(2) E & S	(3) Env	(4) Social	(5) E & S	(6) E & S
Prescriptiveness × Post	-7.412***	-10.503***	-22.734***	-4.240*	-5.871***	-7.381***
	(0.001)	(0.000)	(0.000)	(0.095)	(0.001)	(0.002)
Prescriptiveness	-0.754	-0.550	2.770	-1.518	-0.266	-0.930
	(0.634)	(0.719)	(0.529)	(0.370)	(0.824)	(0.563)
Observations	983	1082	195	750	983	981
Firm FE	Yes	No	Yes	Yes	Yes	Yes
Industry FE	No	Yes	No	No	No	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Proponent-Type FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.419	0.301	0.560	0.428	0.428	0.420
F Statistic	2.997	6.541	11.696	1.408	3.312	2.987

Note: *** p≤0.01, ** p≤0.05, * p≤0.10. P-values are reported in parentheses. In specifications (1) through (4), the dependent variable relates to the percentage of affirmative votes out of the total votes cast. For specification (5), it relates to the percentage of affirmative votes out of all outstanding shares. In specification (6), it relates to the percentage of affirmative votes out of the total of affirmative and negative votes. Additionally, specifications (1), (2), (5), and (6) apply to all E&S proposals. Specification (3) specifically addresses environmental proposals, and specification (4) focuses on social proposals. We drop all anti-ESG proposals which we identify in Section 6.3. We suppress reporting of the constant term and firm-proposal controls. Standard errors are clustered at the meeting-level.

Table A6: Changes in Prescriptiveness Post Treatment
(Subset of Stable Firms)

	Votes For As % Votes Cast				Votes For As % Shares Out	Votes For As % Yes & No
	(1) E & S	(2) E & S	(3) Env	(4) Social	(5) E & S	(6) E & S
Prescriptiveness × Post	-7.995***	-8.227***	-19.954***	-5.164**	-6.152***	-8.014***
	(0.000)	(0.000)	(0.002)	(0.031)	(0.000)	(0.000)
Prescriptiveness	-0.989	-1.587	1.386	-1.767	-0.469	-1.149
	(0.554)	(0.334)	(0.814)	(0.315)	(0.707)	(0.498)
Observations	943	957	174	736	943	941
Firm FE	Yes	No	Yes	Yes	Yes	Yes
Industry FE	No	Yes	No	No	No	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Proponent-Type FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.413	0.338	0.524	0.441	0.422	0.414
F Statistic	4.652	5.545	3.628	2.285	4.332	4.627

Note: *** $p \leq 0.01$, ** $p \leq 0.05$, * $p \leq 0.10$. P-values are reported in parentheses. In specifications (1) through (4), the dependent variable relates to the percentage of affirmative votes out of the total votes cast. For specification (5), it relates to the percentage of affirmative votes out of all outstanding shares. In specification (6), it relates to the percentage of affirmative votes out of the total of affirmative and negative votes. Additionally, specifications (1), (2), (5), and (6) apply to all E&S proposals. Specification (3) specifically addresses environmental proposals, and specification (4) focuses on social proposals. We exclude from our sample all firms identified as non-Stable Firms, as defined in Section 4.3.4. We suppress reporting of the constant term and firm-proposal controls. Standard errors are clustered at the meeting-level.

Table A7: Changes in Prescriptiveness Post Treatment: Heterogeneity amongst ESG Funds (Governance-Families)

	Binary Fund Vote			Ordered Fund Vote	
	(1)	(2)	(3)	(4)	(5)
Prescriptiveness × Post × ESG Fund	0.054**				
	(0.021)				
Prescriptiveness × Post × ESG Fund (G Family)		0.049**		0.052**	
		(0.045)		(0.029)	
Prescriptiveness × Post × ESG Fund (Non-G Family)			0.053**		0.053**
			(0.041)		(0.038)
Prescriptiveness × Post	-0.100***	-0.098***	-0.098***	-0.101***	-0.101***
	(0.004)	(0.004)	(0.004)	(0.003)	(0.002)
ESG Fund	0.245***				
	(0.000)				
ESG Fund (G Family)		0.226***		0.224***	
		(0.000)		(0.000)	
ESG Fund (Non-G Family)			0.228***		0.224***
			(0.000)		(0.000)
Prescriptiveness	-0.025	-0.026	-0.025	-0.020	-0.019
	(0.363)	(0.344)	(0.358)	(0.451)	(0.468)
Observations	528153	528153	528153	528153	528153
Industry FE	Yes	Yes	Yes	Yes	Yes
Proponent-Type FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes	Yes
Fund Controls	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.121	0.117	0.117	0.120	0.119
F Statistic	162.434	160.620	137.576	155.279	133.675

Note: *** p≤0.01, ** p≤0.05, * p≤0.10. P-values are reported in parentheses. The dependent variable in specifications (1) to (3) relates to the variable "Binary Fund Vote", which takes on a value of 1 when a specific fund votes in favor of a proposal, and 0 for all other outcomes. In specifications (4) and (5), the dependent variable relates to the variable "Ordered Fund Vote", which is assigned a value of 1 for a "yes" vote, 0.5 for an "abstained" vote, and 0 for all other outcomes. We drop all anti-ESG proposals which we identify in Section 6.3. We suppress reporting of the constant term, firm-proposal controls, fund controls, as well as all interaction terms in specification (5) (e.g., $Post_t \times FundCat_m$) which are unreported in this Table. Standard errors are clustered at the meeting-level.

Table A8: Changes in ESG Fund Flows Post Treatment

	(1)	(2)	(3)	(4)	(5)
ESG Fund × Post	0.149	0.125			
	(0.725)	(0.794)			
ESG Fund (ES Family) × Post			0.216		
			(0.690)		
ESG Fund (Non-ES Family) × Post				0.036	
				(0.956)	
ESG Fund (Large-ES Family) × Post					0.488
					(0.661)
ESG Fund	-0.118	0.138			
	(0.616)	(0.890)			
ESG Fund (ES Family)			-0.141		
			(0.629)		
ESG Fund (Non-ES Family)				-0.069	
				(0.858)	
ESG Fund (Large-ES Family)					-0.180
					(0.758)
Observations	523510	523456	523510	523510	523510
Fund FE	No	Yes	No	No	No
Year FE	Yes	Yes	Yes	Yes	Yes
Fund Controls	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.000	-0.004	0.000	0.000	0.000
F Statistic	0.376	0.159	0.376	0.348	0.368

Note: *** p≤0.01, ** p≤0.05, * p≤0.10. P-values are reported in parentheses. The dependent variable in all specifications relates to the monthly fund flows of a given fund as defined in Section 6.2. ESG fund variables are defined in Section 5.3.1. We suppress reporting of the constant term and fund-proposal controls. Standard errors are clustered at the meeting-level.

Table A9: Changes in Prescriptiveness Post Treatment: Heterogeneity in Fund Flows amongst ESG Funds

	Binary Fund Vote			Ordered Fund Vote
	(1)	(2)	(3)	(4)
Prescriptiveness × Post × Btm Decile (Fund Flow)	-0.030			-0.038
	(0.586)			(0.492)
Prescriptiveness × Post × Btm Quintile (Fund Flow)		-0.013		
		(0.737)		
Prescriptiveness × Post × Btm Quartile (Fund Flow)			-0.022	
			(0.550)	
Prescriptiveness × Post	-0.049	-0.049	-0.048	-0.053
	(0.185)	(0.187)	(0.206)	(0.145)
Btm Decile (Fund Flow)	-0.134***			-0.140***
	(0.000)			(0.000)
Btm Quintile (Fund Flow)		-0.060***		
		(0.001)		
Btm Quartile (Fund Flow)			-0.066***	
			(0.000)	
Prescriptiveness	-0.035	-0.037	-0.038	-0.030
	(0.246)	(0.224)	(0.210)	(0.320)
Observations	33463	33463	33463	33463
Industry FE	Yes	Yes	Yes	Yes
Proponent-Type FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes
Fund Controls	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.133	0.132	0.132	0.132
F Statistic	38.243	35.752	36.834	41.761

Note: *** p≤0.01, ** p≤0.05, * p≤0.10. P-values are reported in parentheses. The dependent variable in specifications (1) to (3) relates to the variable "Binary Fund Vote", which takes on a value of 1 when a specific fund votes in favor of a proposal, and 0 for all other outcomes. In specification (4), the dependent variable relates to the variable "Ordered Fund Vote", which is assigned a value of 1 for a "yes" vote, 0.5 for an "abstained" vote, and 0 for all other outcomes. Observations not linked to ESG-fund votes, as well as all anti-ESG proposals identified in Section 6.3, are excluded. We suppress reporting of the constant term, firm-proposal controls, fund controls, as well as all interaction terms in specification (5) (e.g., $Post_t \times FundCat_m$) which are unreported in this Table. Standard errors are clustered at the meeting-level.

Table A10: Summary Statistics: Anti-ESG Proposals

	Year						
	2018	2019	2020	2021	2022	2023	Total
Panel A: Prescriptive Proposals							
Non-Anti-ESG							
Frequency	57	62	61	49	113	142	484
Percent (Within-Year)	98.28	87.32	98.39	100.00	86.92	85.54	90.30
Anti-ESG							
Frequency	1	9	1	0	17	24	52
Percent (Within-Year)	1.72	12.68	1.61	0.00	13.08	14.46	9.70
Panel B: Non-Prescriptive Proposals							
Non-Anti-ESG							
Frequency	95	98	107	106	167	162	735
Percent (Within-Year)	97.94	100.00	97.27	95.50	92.27	88.52	94.23
Anti-ESG							
Frequency	2	0	3	5	14	21	45
Percent (Within-Year)	2.06	0.00	2.73	4.50	7.73	11.48	5.77
Panel C: Total							
Non-Anti-ESG							
Frequency	152	160	168	155	280	304	1,219
Percent (Within-Year)	98.06	94.67	97.67	96.88	90.03	87.11	92.63
Anti-ESG							
Frequency	3	9	4	5	31	45	97
Percent (Within-Year)	1.94	5.33	2.33	3.13	9.97	12.89	7.37

Note: This Table illustrates the frequencies and within-year percentages of anti-ESG proposals which proceed to a vote. Panels A and B classify these proposals by whether they are prescriptive or non-prescriptive, respectively. Panel C then combines the data from Panels A and B.

Table A11: Changes in Anti-ESG Proposals Post Treatment

	Votes For As % Votes Cast				Votes For As % Shares Out	Votes For As % Yes & No
	(1) Uncorrected	(2) Uncorrected	(3) IPTW	(4) IPTW	(5) Uncorrected	(6) Uncorrected
Anti ESG	-14.505***	-16.541***	-13.442***	-15.896***	-10.002***	-14.822***
	(0.002)	(0.000)	(0.003)	(0.000)	(0.002)	(0.002)
Anti ESG × Post	-3.865	-1.762	-4.833	-2.049	-3.076	-3.973
	(0.418)	(0.645)	(0.301)	(0.580)	(0.359)	(0.414)
Observations	1082	1180	1082	1180	1082	1080
Firm FE	Yes	No	Yes	No	Yes	Yes
Industry FE	No	Yes	No	Yes	No	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Proponent-Type FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm and Proposal Controls	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.446	0.337	0.448	0.344	0.460	0.447
F Statistic	11.411	15.770	11.714	15.451	12.155	11.448

Note: *** p≤0.01, ** p≤0.05, * p≤0.10. P-values are reported in parentheses. In specifications (1) through (4), the dependent variable relates to the percentage of affirmative votes out of the total votes cast. For specification (5), it relates to the percentage of affirmative votes out of all outstanding shares. In specification (6), it relates to the percentage of affirmative votes out of the total of affirmative and negative votes. In specifications (3) and (4), we implement the Inverse Probability of Treatment Weighting (IPTW) model introduced in the Section 9.4. We suppress reporting of the constant term and firm-proposal controls. Standard errors are clustered at the meeting-level.

Table A12: SDID Estimates: E&S vs Governance Proposals

	Votes For As % Votes Cast				Votes For As % Shares Out	Votes For As % Yes & No
	(1)	(2)	(3)	(4)	(5)	(6)
	E&S	Environmental	Social	Environmental (With S as Control)	E&S	E&S
Average Treatment Effect on the Treated (ATT)	-4.106***	-4.716***	-3.764***	-4.972***	-3.060***	-3.962***
	(0.000)	(0.006)	(0.009)	(0.005)	(0.000)	(0.000)
Observations	3084	3234	3162	1752	3084	3066
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes	Yes	Yes

Note: *** $p \leq 0.01$, ** $p \leq 0.05$, * $p \leq 0.10$. P-values are reported in parentheses. All data is aggregated at the firm-year level, with firms classified as treated if they face more than 50% of E&S (or, alternatively, Environmental or Social proposals) proposals relative to governance proposals, with the exception of column (4) where social proposals are used as a control instead. In specifications (1) through (4), the dependent variable of interest relates to the percentage of affirmative votes out of the total votes cast. For specification (5), it relates to the percentage of affirmative votes out of all outstanding shares. In specification (6), it relates to the percentage of affirmative votes out of the total of affirmative and negative votes. We compute the standard errors using placebo simulations, which involve randomly reassigning the treatment status across units to generate a distribution of placebo estimates; these simulations are repeated 50 times to account for sampling variability.

Figure A1: Weighted Voting Support for E&S vs Governance Proposals

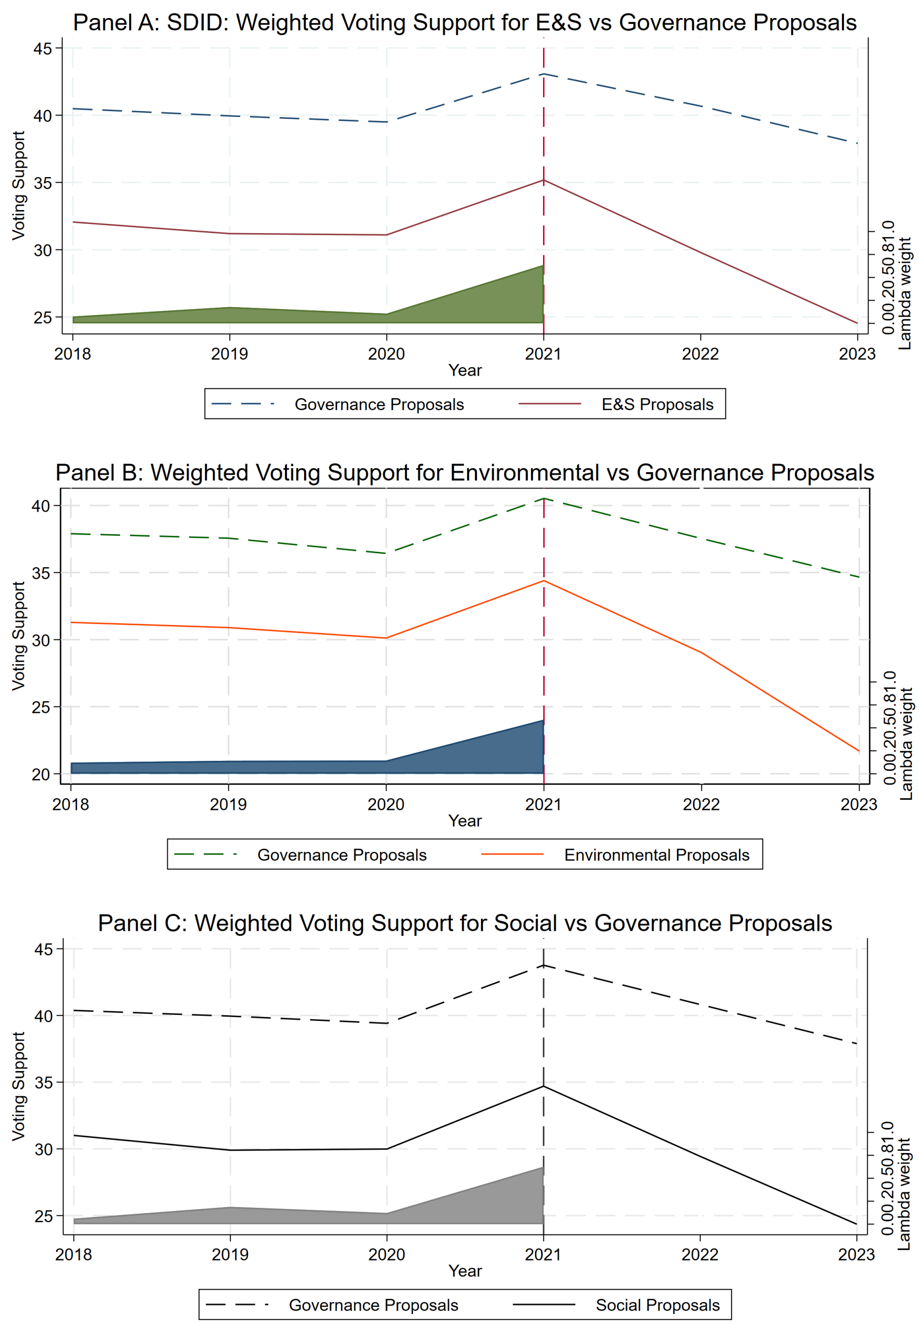

Note: This figure presents the time series of weighted voting support, defined as the percentage of votes in favor relative to the total votes cast, for E&S and governance proposals at the firm-year level. Firms are classified as treated if they face more than 50% of E&S proposals (or, alternatively, Environmental or Social proposals) relative to governance proposals and are depicted by the solid lines, while control firms are represented by the dashed lines. Following specification (7), we harness the weights $\hat{\omega}_i^{\text{SDID}}$ and $\hat{\lambda}_t^{\text{SDID}}$ as defined by Arkhangelsky et al. (2021) to match treated and control units based on pre-treatment trends. The distribution of the weights $\hat{\lambda}_t^{\text{SDID}}$ over time is represented by the shaded areas in each panel, with the intensity of the weights indicated by the legend on the right side of the panels.

Table A13: Model Performance

Model	Accuracy	Precision	Recall	F1 Score	AUC
Baseline Model	0.7414	0.8796	0.6690	0.7600	0.8353
Fold 1	0.7000	0.8552	0.6526	0.7709	0.8031
Fold 2	0.7621	0.8427	0.7853	0.8299	0.8393
Fold 3	0.7163	0.8252	0.6743	0.8024	0.8431
Fold 4	0.7059	0.9375	0.5325	0.7533	0.8379
Mean	0.7211	0.8651	0.6612	0.7891	0.8313
OpenAI Fine-Tuned	0.7198	0.7576	0.8333	0.7937	N.A.
OpenAI Prompt	0.6336	0.6720	0.8467	0.7493	N.A.
HyperParameter Tuned	0.7241	0.8246	0.6812	0.7857	0.7757

Note: This table presents the performance of the "supervised" models discussed in Section 9.2. For each model we report accuracy, precision, recall, F1 score, and AUC, and we give the corresponding mean values from the K-fold validation exercise. We do not report AUC for the LLM models, as they produce only binary labels and therefore lack the continuous probabilities required to plot a ROC curve. The Fine-Tuned model was trained on OpenAI's GPT-4.1-2025-04-14 weights, and the OpenAI Prompt model likewise uses GPT-4.1-2025-04-14.

Figure A2: ROC Curve: Benchmark Model

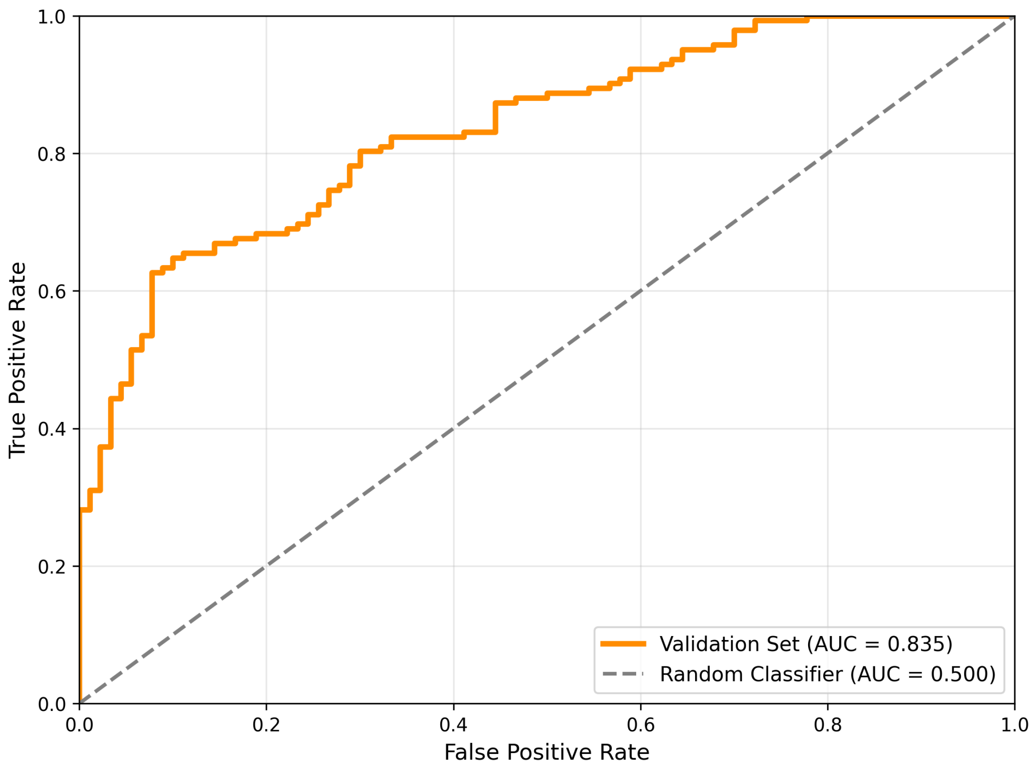

Note: This Figure plots the Receiver Operating Characteristic (ROC) curve for our benchmark single-fold model, the weight set used throughout the study. The curve traces how the true-positive rate rises against the false-positive rate as the classification threshold is swept from 0 to 1. A 45-degree line indicates random guessing, while a curve that traces the upper-left border (an inverted "L") marks perfect discrimination vis-à-vis the validation set. The shape observed here shows strong separation when predicted probabilities cluster near the extremes of 0 or 1.

Table A14: Threshold Error Rates

Thresholds	Type I Errors	Type II Errors	Ratio I/II	Ratio II/I
0.3	0.4444	0.1620	2.7440	0.3644
0.4	0.3222	0.1972	1.6341	0.6119
0.5	0.2667	0.2746	0.9709	1.0299
0.6	0.1667	0.3310	0.5035	1.9859
0.7	0.0889	0.3662	0.2427	4.1197

Note: This table reports the Type I and Type II error rates calculated on the validation set described in Sections 9.2.1 and 9.2.2. For each probability threshold used to classify a proposal as prescriptive (i.e., when its predicted probability is at least the threshold, such as 0.50), we present the Type I error rate, the Type II error rate, and the corresponding error-rate ratios (Type I : Type II and Type II : Type I).

Table A15: Robustness of Baseline Results to Prescriptiveness Thresholds

	(1) ≥ 0.3	(2) ≥ 0.4	(3) ≥ 0.5	(4) ≥ 0.6	(5) ≥ 0.7
Prescriptiveness × Post	-7.260***	-7.981***	-8.476***	-9.542***	-8.230***
	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
Prescriptiveness	-1.373	-0.933	-0.777	-0.759	-2.416
	(0.389)	(0.571)	(0.626)	(0.633)	(0.123)
Observations	1082	1082	1082	1082	1082
Firm FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Proponent-Type FE	Yes	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.412	0.413	0.414	0.420	0.422
F Statistic	4.263	4.159	4.256	5.182	5.207

Note: *** p≤0.01, ** p≤0.05, * p≤0.10. P-values are reported in parentheses. The dependent variable in all specifications relates to the variable "Binary Fund Vote", which takes on a value of 1 when a specific fund votes in favor of a proposal, and 0 for all other outcomes. Each column represents a probability threshold, used to classify a proposal as prescriptive (i.e., when its predicted probability is at least the threshold, such as 0.50). We suppress reporting of the constant term, firm-proposal controls and fund controls. Standard errors are clustered at the meeting-level.

Figure A3: Coefficient Plot of Prescriptiveness × Post across Thresholds

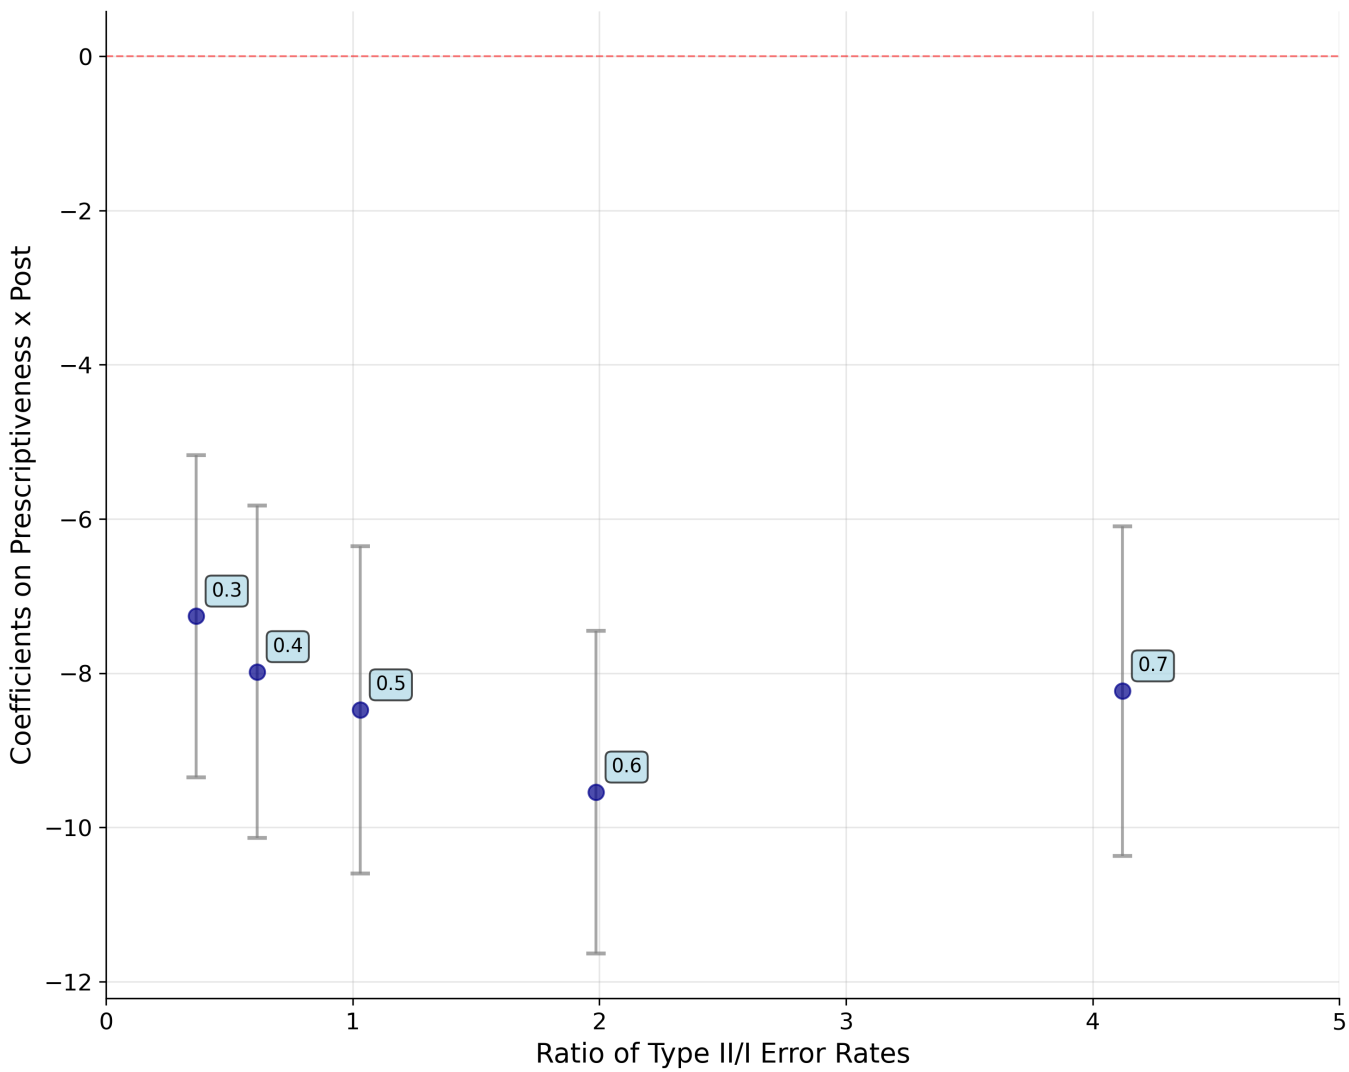

Note: This figure plots the coefficients from Table A15 against the corresponding Type II-to-Type I error-rate ratios for each threshold, as reported in Table A14.

Figure A4: ROC Curves: K-Fold Validation

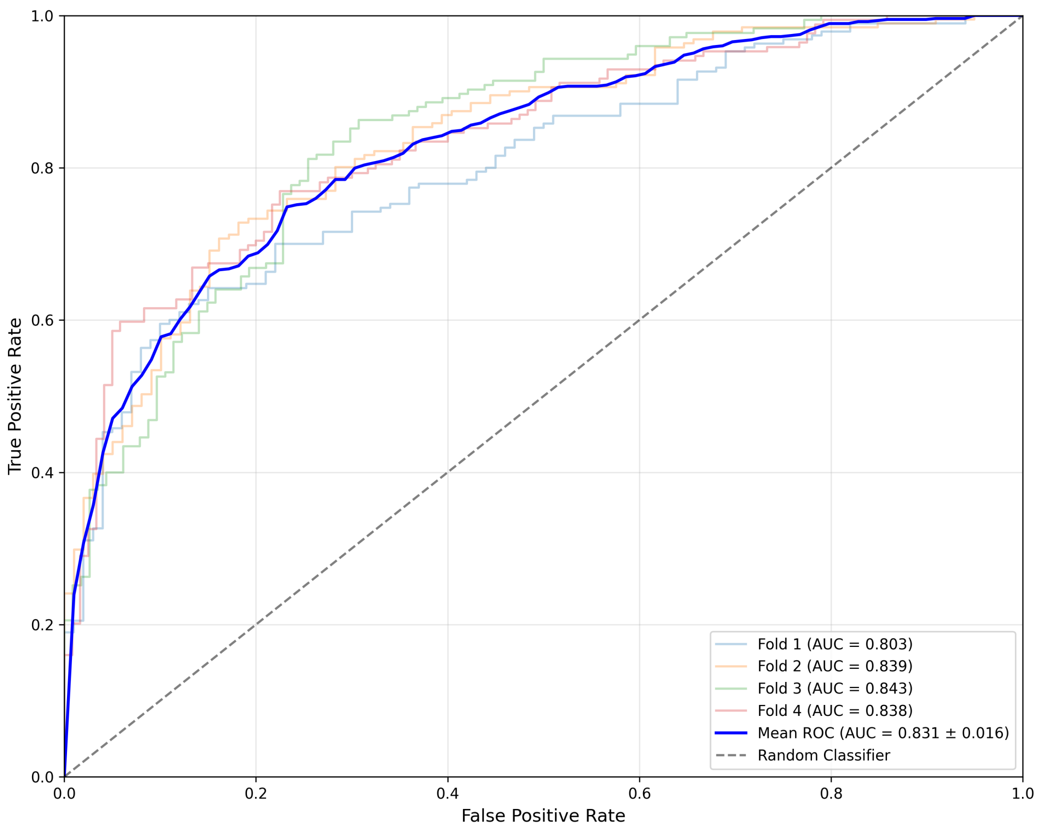

Note: This figure displays the Receiver Operating Characteristic (ROC) curves from the k-fold validation exercise. Four individual model curves appear in teal, and their mean ROC is shown in bright blue. The curves trace how the true-positive rate rises against the false-positive rate as the classification threshold is swept from 0 to 1. A 45-degree line indicates random guessing, while a curve that traces the upper-left border (an inverted "L") marks perfect discrimination vis-à-vis the validation set. The shape observed here for the mean ROC curve exhibits strong separation when predicted probabilities cluster near the extremes of 0 or 1. We omit the baseline model because it uses one of the folds as its held-out validation set (i.e., the baseline model is one of the "5 folds" used in the k-fold validation exercise).

Table A16: Alternative Classifiers

	(1)	(2)	(3)	(4)	(5)
Prescriptiveness × Post (Baseline)	-8.476***				
	(0.000)				
Prescriptiveness × Post (K-Fold (Median))		-7.345***			
		(0.001)			
Prescriptiveness × Post (OpenAI Fine-Tuned)			-5.049**		
			(0.012)		
Prescriptiveness × Post (OpenAI Prompt)				-6.819***	
				(0.004)	
Prescriptiveness × Post (Hyperparameter Optimized)					-6.836***
					(0.001)
Prescriptiveness (Baseline)	-0.777				
	(0.626)				
Prescriptiveness (K-Fold (Median))		-2.437			
		(0.146)			
Prescriptiveness (OpenAI Fine-Tuned)			-5.696***		
			(0.000)		
Prescriptiveness (OpenAI Prompt)				0.320	
				(0.863)	
Prescriptiveness (Hyperparameter Optimized)					-2.325
					(0.123)
Observations	1082	1082	1082	1082	1082
Firm FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Proponent-Type FE	Yes	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.414	0.418	0.432	0.398	0.415
F Statistic	4.256	4.719	6.712	2.527	4.800

Note: *** p≤0.01, ** p≤0.05, * p≤0.10. P-values are reported in parentheses. The dependent variable in all specifications relates to the variable "Binary Fund Vote", which takes on a value of 1 when a specific fund votes in favor of a proposal, and 0 for all other outcomes. Each of the first 5 rows presents results from a distinct supervised model described in Section 9.2. When the prescriptiveness variable is generated by an LLM, the Fine-Tuned classifier is trained on OpenAI's GPT-4.1-2025-04-14 weights, and the OpenAI Prompt classifier likewise uses GPT-4.1-2025-04-14. We suppress reporting of the constant term, firm-proposal controls and fund controls. Standard errors are clustered at the meeting-level.

Figure A5: System Prompt

```
You are a corporate-governance analyst.

Label definitions
=================
• **PRESCRIPTIVE**  → 1
  A proposal is prescriptive when it **seeks intricate detail** or **requests a
  specific strategy, method, action, outcome, or timeline** for addressing an
  issue—effectively supplanting the judgment of management and the board.
  Thus, any proposal that sets explicit timeframes or methods for implementing
  complex policies is prescriptive.
  *Example:* “Report annually on short-, medium- and long-term greenhouse-gas
  targets aligned with the Paris Agreement goals to keep the global
  temperature rise well below 2 °C and pursue efforts to limit it to 1.5 °C.”
  This is prescriptive because it directs the adoption of time-bound targets.

  **Indicators include (non-exhaustive):**
  ▸ verbs such as “adopt”, “implement”, “amend”, “eliminate”, “separate”,
    “require”, “set”, “establish”, “cease”, “phase out”;
  ▸ explicit methods (“use 100 % renewable electricity”, “link bonuses to X”);
  ▸ measurable timelines or numeric goals (“by 2028”, “reduce GHG 50 %”).

• **NON-PRESCRIPTIVE**  → 0
  A proposal is non-prescriptive when it **does not seek intricate detail** and
  **does not specify a strategy, method, action, outcome, or timeline**.
  Instead, it defers to management’s discretion—for example, by asking the
  company to *consider*, *evaluate* or *discuss* the feasibility of a course
  of action.
  *Example:* “Prepare a report describing if, and how, the company plans to
  reduce its total contribution to climate change and align its operations and
  investments with the Paris Agreement’s goal of keeping global temperatures
  well below 2 °C.”
  This is non-prescriptive because it leaves management free to decide
  whether—and how—to reduce the carbon footprint.

  **Indicators include (non-exhaustive):**
  ▸ requests to “consider,” “review,” “assess,” “evaluate,” “discuss”;
  ▸ broad calls for disclosure or feasibility studies without prescribing
    how results must be acted upon.

Output format
=============
Respond **only** with JSON formatted exactly as below:

{
  "label_int": 0 | 1,
  "label_text": "non-prescriptive" | "prescriptive"
}

Rules
-----
1. If you are genuinely uncertain, default to **label_int = 0**.
2. Do **not** output anything except the JSON object.
```

Note: This Figure shows the length-optimized instruction set employed in our few-shot prompting approach. The prompt directs the LLM to classify proposals as prescriptive or non-prescriptive according to the criteria it contains and is accompanied by 30 randomly selected, labeled examples from the training set to remain within the model’s context window.

Figure A6: Two-Dimensional Representation of Environmental Proposals.

Note: In this figure, topic labels have been condensed to three keywords for conciseness. Nevertheless, we utilize a representation model from OpenAI to relabel topics based on their key words. For example, Topic 1 is linked with the label "Corporate Sustainability Reporting on ESG Metrics."

Table A17: Ordinary Business Exception Proxy for Prescriptiveness

	Uncorrected		Heckman Selection Models		IPTW Models	
	(1)	(2)	(3)	(4)	(5)	(6)
14-a(8)(i)(7) × Post	-4.363	-5.167*	-4.364	-5.167*	-0.167	2.164
	(0.153)	(0.066)	(0.111)	(0.051)	(0.955)	(0.465)
Post			-1.360	-2.193		
			(0.533)	(0.186)		
14-a(8)(i)(7)	0.010	-0.323	0.010	-0.323	0.453	-2.248
	(0.997)	(0.878)	(0.996)	(0.871)	(0.854)	(0.324)
Observations	1082	1180	1856	1856	1082	1180
Firm FE	Yes	No	Yes	No	Yes	No
Industry FE	No	Yes	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Proponent-Type FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.389	0.281			0.502	0.418
F Statistic	1.482	3.018			0.426	0.940
Chi-Square			0.1367	0.3395		

Note: *** p≤0.01, ** p≤0.05, * p≤0.10. P-values are reported in parentheses. The dependent variable in all specifications relates to the proportion of affirmative votes out of the total votes cast. In specifications (3) and (4), we implement the Heckman Selection model introduced in Section 9.4, while in specifications (5) and (6), we implement the Inverse Probability of Treatment Weighting (IPTW) model introduced in the same Section. We suppress reporting of the constant term and firm-proposal controls. Standard errors are clustered at the meeting-level.

Table A18: Decomposition of Variation in Voting for E&S Proposals

	Panel A: E & S Proposals					
	(1) Year FE Only	(2) Year and Proponent-Type FE	(3) Year and Industry FE	(4) Year and Firm FE	(5) Year and Proponent FE	(6) Year and Fund FE
Observations	877227	877227	877227	877227	877227	876760
R-Sq	0.022	0.037	0.071	0.112	0.140	0.286
Firm FE	No	No	No	Yes	No	No
Industry FE	No	No	Yes	No	No	No
Fund FE	No	No	No	No	No	Yes
Proponent FE	No	No	No	No	Yes	No
Proponent-Type FE	No	Yes	No	No	No	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

	Panel B: Environmental Proposals					
	(1) Year FE Only	(2) Year and Proponent-Type FE	(3) Year and Industry FE	(4) Year and Firm FE	(5) Year and Proponent FE	(6) Year and Fund FE
Observations	173584	173584	173584	173584	173584	172789
R-Sq	0.044	0.062	0.161	0.209	0.211	0.305
Firm FE	No	No	No	Yes	No	No
Industry FE	No	No	Yes	No	No	No
Fund FE	No	No	No	No	No	Yes
Proponent FE	No	No	No	No	Yes	No
Proponent-Type FE	No	Yes	No	No	No	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

	Panel C: Social Proposals					
	(1) Year FE Only	(2) Year and Proponent-Type FE	(3) Year and Industry FE	(4) Year and Firm FE	(5) Year and Proponent FE	(6) Year and Fund FE
Observations	703643	703643	703643	703643	703643	703150
R-Sq	0.021	0.038	0.067	0.106	0.135	0.297
Firm FE	No	No	No	Yes	No	No
Industry FE	No	No	Yes	No	No	No
Fund FE	No	No	No	No	No	Yes
Proponent FE	No	No	No	No	Yes	No
Proponent-Type FE	No	Yes	No	No	No	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Note: This table reports the R^2 statistics associated with regressing the dependent variable "Binary Fund Vote" on a variety of fixed effects and a constant. we suppress reporting of the constant term. Panel A outlines these R^2 statistics for E&S proposals, Panel B for environmental proposals specifically, and Panel C for social proposals.

Table A19: Panel Regressions of Individual Fund Votes on Prescriptiveness: Environmental/Social Distinction

	Environmental				Social			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Prescriptiveness	-0.086**	-0.112***	-0.116**	-0.123***	-0.065***	-0.065***	-0.063***	-0.066***
	(0.045)	(0.008)	(0.015)	(0.004)	(0.005)	(0.002)	(0.005)	(0.001)
Observations	106252	106252	106252	106252	421901	421901	421901	421901
Firm FE	Yes	No	Yes	No	Yes	No	Yes	No
Industry FE	No	Yes	No	Yes	No	Yes	No	Yes
Fund FE	No	No	No	No	No	No	No	No
Proponent-Type FE	No	No	Yes	Yes	No	No	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Proposal Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fund Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-Sq	0.238	0.189	0.246	0.201	0.140	0.107	0.144	0.111
F Statistic	39.011	41.996	42.915	42.250	125.279	128.946	124.783	126.747

Note: *** $p \leq 0.01$, ** $p \leq 0.05$, * $p \leq 0.10$. P-values are reported in parentheses. The dependent variable in all specifications relates to the variable "Binary Fund Vote", which takes on a value of 1 when a specific fund votes in favor of a proposal, and 0 for all other outcomes. Specifications (1) to (4) focus on environmental proposals, whereas specifications (5) to (8) target social proposals. we suppress reporting of the constant term, firm-proposal controls and fund controls. Standard errors are clustered at the meeting-level.