Shareholder Democracy and the Market for Voting Advice *

Odilon Câmara — *University of Southern California*
John Matsusaka — *University of Southern California*
Chong Shu — *University of Utah*

November 2, 2025

Abstract

In corporate elections, most votes are cast based on recommendations provided by for-profit proxy advisory companies. We develop a model of the voting advice market to explore how competition and demand for advice shape the slant of the advice offered. In the model, advisory firms compete via prices and advice procedures. Shareholders have heterogeneous goals, differing in the weight they place on financial returns versus nonfinancial ("social") returns. In equilibrium, advisory firms tailor their advice to the preference of the average investor, which can skew voting outcomes away from those that would prevail if investors had full information.

*This paper incorporates some ideas from the retired working paper, "A Theory of Proxy Advice when Investors Have Social Goals" (2021). For helpful suggestions, we thank discussant Benjamin Noble and participants at the 2024 SoCal Political Economy and Institutions Conference (Caltech), 2024 SITE Conference (Stanford), 2023 Latin American Workshop in Law and Economics, and workshops at LSE, Harvard University, Hong Kong Baptist University, Nanyang Technological University, National University of Singapore, and USC.

I. Introduction

American corporations hold approximately 250,000 elections each year to choose directors and corporate policies. Over 70 percent of votes are cast by institutional investors, primarily mutual funds and pension funds.[1] These investors, with their broad portfolios, face a daunting task: Vanguard, for example, reported voting on 184,521 items across 13,490 companies in 2022 (Vanguard 2023). Because acquiring information about so many election items is prohibitively expensive, most institutional investors rely on voting advice from for-profit proxy advisory companies. The effectiveness of shareholder democracy thus hinges on the quality of advice offered by these advisory companies. This is cause for concern because only two advisory firms, Institutional Shareholder Services (ISS) and Glass Lewis, control over 90 percent of the market, and their recommendations appear to be slanted toward outcomes favored by socially responsible investors – a relatively narrow segment of the investor base.[2]

This paper develops a model of the proxy advice market to understand how recommendations are formed in a competitive environment, and how those recommendations affect corporate elections. While the literature has focused on markets with a monopolist advisor, we are particularly interested in the role of competition. An important feature of our model is that investors have heterogeneous goals: they differ in the weights they place on financial returns versus nonfinancial ("social") returns, such as reductions in carbon emissions, an assumption grounded in survey evidence (Riedl and Smeets 2017). The model can also be interpreted in terms of other forms of preference heterogeneity, such as preference for short-run vs. long-run returns. We assume that investors want their votes to express their underlying values even if their votes are not pivotal ("expressive preferences"), and that some investors care more than others about their voting record. For example, a public pension fund may care more about its voting record because it is scrutinized by politicians, labor unions, and media, than a small index fund.

[1] As of 2019, institutional investors held 70 percent of corporate equity, compared to 30 percent held by retail investors (Broadridge + PwC, 2019).

[2] Shu (2024) finds that ISS and Glass Lewis together had 91 percent of the mutual fund business as of 2021. Bolton et al. (2020) estimate the spatial location of ISS and Glass Lewis recommendations compared to the preferences of various types of institutional investors.

In the model, investors vote on many proposals, each characterized by a level of financial and social returns that are determined by an unknown state. In order to cast a vote that is aligned with their interests, investors purchase recommendations from proxy advisory firms that compete on prices and recommendation policies. A recommendation policy is a binary "experiment" that recommends voting for or against a proposal based on which action provides the highest expected utility given a predetermined relative weight on social versus financial returns. The relative weight is called the "slant" of the policy. Upon receiving advice, investors Bayesian update their beliefs and vote. In equilibrium, they choose to follow the slanted recommendation. Customers have preferences over the amount of slant in the advice they receive, which depends on their own tradeoff between financial and social returns. We describe advice as "biased" if its slant is different from what the majority of shareholders would like it to be. Advice prices are set through discriminatory price competition.

According to proxy advisors, their advice policies are "an aggregation of institutional investor preferences" over opinions that differ (Hayne and Vance 2019, p. 975). We characterize how this aggregation occurs: proxy advisors maximize profit by choosing recommendation policies that target the average preference of their customers, weighted by how much investors care about their voting reputations. This has two implications. First, advice is slanted toward the preferences of funds with the highest expressive value from voting. If, as we argue, socially responsible investment (SRI) funds have particularly high expressive benefits (because their pro-SRI voting record is part of their marketing strategy), the model predicts that voting recommendations are slanted toward their preferences, and their interests end up being overweighted in corporate elections. Second, because the recommendation policy preferred by the average voter is generally different from the policy preferred by the median voter, proxy advice does not produce the democratic outcome of majority rule and therefore is biased in the sense we are using the term.

One of our contributions is to show how slanted advice can emerge from profit maximization. Proxy advisors are often accused of offering biased advice, and empirical estimates indicate that their recommendations occupy ideological positions distant from many of their customers (Bolton et al. 2020). The prevailing explanation for bias is conflicts of interest: proxy advisors distort their recommendations to favor companies or activists that buy their consulting services (Li 2018; Ma and Xiong 2021; Business Roundtable 2025). Another ex-

planation is that proxy advisors make recommendations based on the personal ideologies of their managers or owners (American Accountability Foundation 2023; U. S. House of Representatives Committee on the Judiciary 2025). In contrast, our model shows how bias can emerge through profit maximization without requiring conflicts of interest or managerial biases.[3] The model also produces predictions about the nature and direction of advice slant as a function of the distribution of customer preferences and demand for voting reputation.

The paper also contributes to the literature by analyzing how competition affects proxy advice and corporate elections. In equilibrium, advisory firms segment the market, with each firm choosing an advice policy that maximizes the weighted average preference of its consumers. While differentiation in advice policies increases the aggregate utility (from voting) of investors, it does not necessarily make corporate elections more democratic. Corporate elections can become less democratic if competition results in the median investor being offered advice that is less aligned with its preferences. We identify conditions under which competition improves versus deteriorates the information of the median voter, and thus the "informed representation" of corporate elections. We also show that as the number of firms grows to the number of customer types, the market approaches a "perfectly competitive" outcome: firms earn zero profits, each firm sells to one type of investor, investor utility is maximized, and election outcomes become fully informed and democratic.

After laying out the core model and developing its main implications, we explore several extensions. First, we consider customized advice: each advisor can offer a limited number of policies. This extension is based on the observed practices of ISS and Glass Lewis (see Section VII.A for a description of their policy offerings). We interpret these policies as different values of α. The key constraint is that each advisor is only able to offer limited customization, in the form of a certain number of different policies. We then consider other extensions, such as internal information collection by funds. The extensions are of substantive interest and help to show that the main forces we identify are fairly robust. We also discuss the model's empirical and policy implications. Empirically, we contrast the implications of our theory of

[3]Ma and Xiong (2021) develop a model in which proxy advisors offer advice that diverges from value maximization if their homogeneous customers do not want to maximize value or have mistaken beliefs. Advice in their model is not biased in our sense of being different from what the majority of customers desire. (They define bias to mean a slant that differs from value maximization.)

advice bias with existing theories, offer a new empirical fact regarding disagreement between proxy advisors, and provide a perspective on robo-voting. In terms of policy, we discuss recent antitrust concerns that stem from the ISS-Glass Lewis duopoly, and analyze the effect of regulations that require funds to vote.

Our paper is related to multiple research areas. In terms of the finance literature, our focus on competition and heterogeneous investors differs considerably from previous theoretical research on proxy advice. Malenko and Malenko (2019), and Buechel et al. (2024) develop models of a monopolist advisor facing customers who would agree on how to vote if they had full information, in order to study the incentives to acquire information. Ma and Xiong (2021) study a monopolist advisor that sells advice to homogeneous investors but also has a related consulting business that creates a conflict of interest; they also study a case where investors have nonfinancial goals and mistaken beliefs. Malenko et al. (2025) study a monopolist that biases its recommendations in order to increase the probability of close elections. All of these models assume a monopoly advisory firm, abstracting away from the market structure issues that we explore, issues that have attracted considerable scrutiny recently from policy makers. These models also abstract away from conflicts between investors: they assume that investors have homogeneous preferences, in the sense that voters would agree on how to vote if they knew the realized state, which make them unable to address emerging issues relating to conflicts between investors over social goals.

Our paper is also related to the literature on information markets and elections, in particular the role of mass media (see Prat and Strömberg 2013 for a literature review).[4] Some studies focus on a single provider of information (e.g., Duggan and Martinelli 2011) or competition between outlets that seek to sway the election (e.g., Chan and Suen 2009); our focus is on understanding the effects of competition between profit-maximizing firms without any policy preferences. Other studies consider competition between profit-maximizing media that sell advertising (Bernhardt, Krasa, and Polborn 2008; Chan and Suen 2008) or sell news to guide private decisions (Strömberg 2004). We are concerned with firms that sell information specifically to guide voting decisions, which is the central purpose of the proxy

[4]There are several papers that model information markets but do not study how they affect election outcomes (e.g., Gentzkow and Shapiro 2006, and Mullainathan and Shleifer 2005). See Bergemann and Bonatti (2019) for a literature review.

advice market.

The closest paper to ours is Perego and Yuksel (2022), which studies competition between profit-maximizing information intermediaries that sell to consumers with heterogeneous preferences. We draw on their insight that competition can be modeled as a spatial game, and like them, provide a microfoundation based on institutional features of the proxy advice market, which allows exploitation of results from the spatial economics (Hotelling) and industrial organization literatures, especially Lederer and Hurter (1986). A key difference is that in Perego and Yuksel (2022), media outlets sell a noisy point estimate of the value of a proposal, which the consumer then decodes to make the voting decision, which is a natural fit for political markets in which voters face a small number of proposals. In corporate elections, on the other hand, because investors must vote on tens of thousands of different proposals in a short time, it's not feasible for them to decode complex information election by election; consequently, the information is transacted in the form of binary approve/reject recommendations. In some respects, this simplifies the analysis, but it also allows us to work with general functional forms while producing a very tractable model. Allowing asymmetric distributions permits exploration of the wedge between the median and average voter. As we show, the fact that an advisor targets the average voter instead of the median voter has important welfare implications. The model's tractability allows it to be naturally extended to a number empirical and policy applications.

Although our model draws from and contributes to several related strands of literature, its primary focus is on the corporate advice market. To that end, the model embeds a number of assumptions that are based on specific features of that market, and may not apply to other markets, such as the reliance on binary recommendations and perhaps the presence of expressive voters. We believe our theoretical framework advances the study of advice markets that share these institutional features, while at the same time offering a new perspective on the market for proxy advice. By showing how key features of that market can be embedded in a model with a number of standard features, we are able to produce an economically grounded framework that can be used to speak to empirical and policy debates. The model may also be useful for studying more conventional political markets in which voters rely on binary recommendations (such as endorsement, as in Lupia 1994) because they face a large number of races. In California, for example, citizens are asked to

vote on a long list of candidates for federal, state, county, city, and judicial races, as well as dozens of ballot propositions.

II. Model

A. Investors and Proposals

There is a measure one of investment funds (investors or shareholders), indexed by $i \in [0,1]$, and a measure one of proposals to be voted on, indexed by $j \in [0,1]$. Shareholders vote "yes" (approve) or "no" (reject) on each proposal, using a majority voting rule. Variable $v_{ij} \in \mathbb{R}$ captures the *policy payoff*: if proposal j is approved by shareholders, then investor i attains a policy payoff v_{ij}; if the proposal is rejected, then investor i receives a policy payoff $-v_{ij}$. Therefore, investor i would like shareholders to approve proposals with a positive value and reject proposals with a negative value. The absolute value of v_{ij} indicates the importance of this proposal for the investor.

In practice, the value of a proposal is a complex function of many variables. To keep the model tractable and focus on the tradeoff between social return and financial return, we assume that a proposal's value takes the form:

$$v_{ij} \quad = \quad \theta_i s_j + (1-\theta_i) r_j, \tag{1}$$

where $r_j \in \mathbb{R}$ captures the financial return, and $s_j \in \mathbb{R}$ captures the social return. These returns are specific to a proposal and common across all funds. Preference parameter $\theta_i \in [0,1]$ is the relative weight that fund i attaches to the social versus financial return. Because weights differ, even if funds know all of the facts about a proposal, they may differ on whether they favor approval or rejection. The model could also be interpreted in terms of other types of preference heterogeneity, such as differences in investment horizon, tax exposure, and liquidity needs. In the Online Appendix, we consider the general case in which each proposal has many different dimensions.

Funds know their preference parameters, but don't know the realized values of s_j and r_j. They only know the joint distribution, defined by the cdf (cumulative density function) $F(s,r)$ and pdf (probability density function) $f(s,r)$. We assume that f has no atoms, full support on some subset of $\mathbb{R}^2$, and $(0,0)$ belongs to the interior of this set.

B. *Utility from Voting*

Funds are assumed to vote for expressive reasons, following the expressive voting theory (Fiorina 1976). Expressive voting is the idea that voters derive a direct utility from expressing their preferences in their vote, even if their vote is not pivotal. It contrasts with the instrumental voting theory in which investors value their votes only if they are pivotal. We assume expressive voting in order to sidestep secondary issues related to learning from pivotality — for the same reason, Perego and Yuksel (2022) consider individuals who behave as if they were voting expressively. It's also the case that instrumental voting is descriptively questionable for many corporation elections, in which there is almost no chance of being pivotal — in only 2.7 percent of elections is the margin of victory less than 5 percent, and the typical margin is 67 percent.[5]

If $a_{ij} \in \{-1, 1\}$ is a fund's vote, where $a_{ij} = 1$ is approval and $a_{ij} = -1$ is rejection, then the *expressive utility* of voter i is

$$u_i(a_{ij}, v_{ij}) \quad = \quad a_{ij}\lambda_i v_{ij}, \tag{2}$$

where the preference parameter $\lambda_i \in (0, \overline{\lambda}]$ is the importance that the fund attaches to casting a vote that expresses its preferences (voting for a proposal with a positive value and against a proposal with a negative value). Abusing notation, we can rewrite the expressive utility of fund (λ, θ) voting on proposal (s, r) as

$$u(\lambda, \theta, s, r, a) \quad = \quad a\lambda\left[\theta s + (1-\theta)r\right]. \tag{3}$$

It is optimal to vote yes if $[\theta s + (1-\theta)r] > 0$, and vote no otherwise. Fund preference

[5]The reported vote margin numbers are our calculations using data from 489,657 elections during 2003-2018. Because we assume a continuum of voters, in our analysis a fund's probability of casting a pivotal vote is zero, implying that its willingness to pay for information comes entirely from its expressive utility. Note that expressive voting does *not* mean funds are indifferent to financial returns when voting – they always cast their votes for the outcome that delivers the highest utility according to their weight on financial and social returns. In the Online Appendix, we consider an alternative scenario — instead of assuming that voters are never pivotal, we assume that voters are always pivotal and electoral outcomes are probabilistic. We also consider an extension in which funds are willing to pay a premium for information when they believe that elections will be close.

parameters are distributed according to the cdf $G(\lambda, \theta)$ and continuous pdf $g(\lambda, \theta)$, with variances σ_λ^2 and σ_θ^2, covariance $cov_{\lambda\theta}$, and Pearson correlation coefficient $\rho_{\lambda\theta}$.

We have in mind several reasons why some funds may care more than others about expressing their preferences through their votes. One is that a fund may advertise its voting record as a way to attract fund flow. Trillium Asset Management, a well-known SRI fund, declares on its website: "[W]e're proud of the responsibility we've taken to develop and communicate to clients our proxy voting policies, and we take that voting seriously," immediately below which it provides its proxy votes for the most recent 13 years.[6] Also, a fund's votes may be scrutinized by third parties, such as a public pension fund that is under pressure to align its votes with the preferences of elected officials in the state, or very large funds that receive considerable attention from mass media and the general public. Our model assumes heterogeneity across funds, capturing the idea that some funds may care a lot about their voting records, while others may care very little, such as perhaps a small index fund.

C. Proxy Advisory Firms

There are $N \geq 1$ proxy advisory firms that produce and sell recommendations about proposals. In practice, advisory firms do not collect detailed and complete information on the consequences of each proposal, nor do their reports provide deep, comprehensive analyses for each recommendation, in part because the costs to investors of reviewing that amount of detail would be prohibitive. We capture the coarseness of advice by assuming that an advisory firm creates a binary approve/reject test. Following the information design literature, we assume that the firm can choose any test that maps the realized state—the return (s, r) of proposal j—to an approve/reject recommendation. Also following this literature, we abstract from fixed and marginal costs, and instead constrain each firm to produce at most one type of advice. We consider multiple policies per firm in Section VII.A.

In a multifirm model where firms compete on prices and by designing flexible advice policies for heterogeneous consumers, one might expect a nonexistent or intractable equilibrium. Our model is quite tractable because, as we show in the Online Appendix, it is

[6] https://www.trilliuminvest.com/esg/advocacy-policy, accessed September 20, 2020. SEC rules require funds to report their votes, so this information is relatively easy for investors to track.

always optimal for firms to choose a *targeted policy*, which is defined by a single parameter $\alpha \in [0,1]$: it recommends voting yes if $\alpha s + (1-\alpha)r > 0$ and voting no otherwise.[7] We call it a targeted policy because the optimal advice for a fund with preference θ is $\alpha = \theta$: such advice allows the fund to cast its votes as if it had full information.

We assume that each firm offers a single advice policy. We imagine that, in order to credibly and accurately implement a specific policy α, the firm must employ a team of specialists that applies protocols that partition recommendations based on this particular α. It is important to emphasize that the advisory firm does not directly estimate r and s for each proposal. The firm only generates a binary yes/no recommendation. We have in mind that the firm does not collect the raw information on r and s; rather its researchers search for critical pieces of information that reveal if a proposal lies in the accept or reject region (what observers call "checking boxes").

Each advisory firm makes a take-it-or-leave-it price offer to each fund. It also squares with the fact that advisory firms do not post flat prices on their websites, and with Shu's (2024) finding, based on Public Records Act requests, that the annual fees paid by the largest public pension funds ranged from \$50,000 to \$500,000. This is not an anonymous market: sellers have considerable information about their potential customers, such as their investment strategy, size, and past voting record. Our use of Bertrand competition with first-degree price discrimination also avoids technical issues with equilibrium existence.

In summary, the timing of the game is the following:

1. Advisory firms simultaneously select policies $\{\alpha_1, \ldots, \alpha_N\}$.
2. After observing each others' policies, advisory firms simultaneously offer prices to each fund.
3. Each fund selects one policy to purchase, or does not purchase advice.
4. Advisory firms run their tests and issue their recommendations.
5. Funds observe the recommendation they purchased and then vote.

We focus on subgame perfect equilibria with pure strategies, which we refer to as equilibria.

[7] This is an example of a "linear disclosure" experiment. Interestingly, Smolin (2023) proves that linear disclosure experiments are also optimal in his mechanism design model, which is quite different than ours.

III. The Value of Advice

A. *Basic Properties*

To simplify presentation, we assume that the expected values of s and r are zero, and $E[r|s > 0] \geq 0$ and $E[s|r > 0] \geq 0$. These assumptions mean that, without additional information about a proposal, all funds are indifferent between voting yes and no, and the outside option of acquiring no information yields an expressive voting payoff of zero. They also imply that the value of advice is positive. In the Online Appendix, we prove that our main equilibrium features continue to hold without these assumptions — the only difference is that some funds might not purchase an advice.

Consider an advisory firm that produces recommendations according to an advice policy α. Because recommending in favor of a proposal implies that $\alpha s + (1-\alpha)r > 0$ and a recommendation against implies the opposite, fund i can use Bayesian updating to compute the expected value of v_{ij} conditional on the recommendation. The fund chooses to vote for a proposal if this conditional expectation is positive, and against it if the conditional expectation is negative. Lemma 1 shows that funds benefit from voting according to the recommendation provided. Hence, funds rationally choose to acquire and follow advice in this setup, even knowing that it may be slanted.

Next we compute the value of advice and how much funds are willing to pay for it. Consider a fund that receives recommendations from an advisory firm, generated according to some advice policy α. By following the recommendation, the fund's payoff from the social dimension is s when $s > \frac{-(1-\alpha)}{\alpha} r$ and $-s$ when $s \leq \frac{-(1-\alpha)}{\alpha} r$. Integrating over all possibilities, the expected return from the social dimension given advice policy α is

$$V_s(\alpha) = \int_{-\infty}^{\infty} \left[\int_{-\infty}^{-\frac{(1-\alpha)}{\alpha} r} -s f(s,r) ds + \int_{-\frac{(1-\alpha)}{\alpha} r}^{\infty} s f(s,r) ds \right] dr, \tag{4}$$

and the expected return from the financial dimension is

$$V_r(\alpha) = \int_{-\infty}^{\infty} \left[\int_{-\infty}^{\frac{-\alpha}{1-\alpha} s} -r f(s,r) dr + \int_{\frac{-\alpha}{1-\alpha} s}^{\infty} r f(s,r) dr \right] ds. \tag{5}$$

The value of advice α for fund (λ, θ) is the expected value of voting according to the recommendations generated by α, expressed as:

$$U_{\lambda\theta}(\alpha) = \lambda \left[\theta V_s(\alpha) + (1-\theta) V_r(\alpha) \right]. \tag{6}$$

We sometimes refer to this as the "expected utility" from policy α. The following lemma summarizes some useful properties of the value of information.

Lemma 1. *Fund (λ, θ) benefits from voting according to the recommendation provided by advice policy α, and is willing to pay $U_{\lambda\theta}(\alpha)$ to purchase this advice. This value is maximized at $\alpha = \theta$, and strictly decreases as α moves away from θ. Formally, $\forall \alpha \in (0,1)$ and $\theta \in [0,1]$, we have $V_s'(\alpha) > 0$, $V_r'(\alpha) < 0$, $U_{\lambda\theta}(\alpha) \geq 0$, $U_{\lambda\theta}'(\alpha) > 0$ if $\alpha < \theta$, and $U_{\lambda\theta}'(\alpha) < 0$ if $\alpha > \theta$.*

The opposite signs on the derivatives $V_s'(\alpha)$ and $V_r'(\alpha)$ imply a tradeoff: increasing the weight that an advisor places on the social dimension increases the quality of the decision in that dimension but decreases the quality of the decision in the financial dimension. If the fund is more socially inclined than the advice policy ($\alpha < \theta$), then the fund strictly benefits from a marginal increase in α, which brings it closer to the fund's optimal policy θ, and conversely when $\alpha > \theta$.

B. An Example for Intuition

Suppose $F(s,r)$ represents a uniform distribution on a disk with radius $R > 0$, that is, $f(s,r) = 1/(\pi R^2)$ if $s^2 + r^2 \leq R^2$, and $f(s,r) = 0$ otherwise. Consider an advisor using policy $\alpha = 0.5$. The dots in Figure 1 depict a set of proposals drawn from F, where a proposal's financial return is its position on the horizontal r axis and its social return is its position on the vertical s axis. Policy $\alpha = 0.5$ is captured by the dashed black line with slope $-(1-\alpha)/\alpha = -1$ going through the origin. The advisor recommends approval of all policies above the dashed line (blue points) and rejection of all policies below the line (red points). A decrease in α is represented by a clockwise rotation of the dashed line.

The solid line in Figure 1 represents the indifference curve of a fund with preference $\theta = 0.25$, passing through the origin with a slope $-(1-\theta)/\theta = -3$. The fund would like to vote yes on proposals above the solid line and vote no on proposals below the line.

Independent of α and θ, the fund always votes ex-post correctly when the proposal is in quadrants I or III. Proposals in the north-east quadrant I deliver non-negative values in both dimensions, hence every policy $\alpha \in [0,1]$ recommends approval and every fund $\theta \in [0,1]$ prefers to vote for approval. Proposals in the south-west quadrant III deliver non-positive values in both dimensions, hence every policy recommends rejection and every fund prefers

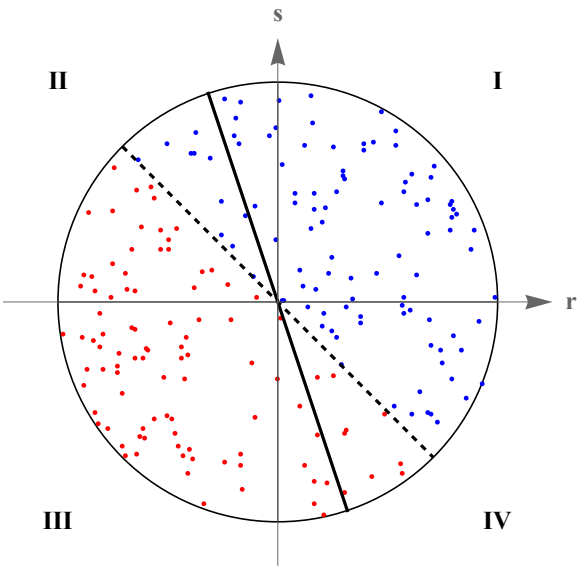


Figure 1: Proposals uniformly distributed on a disk. Each dot represents a proposal. The dashed line represents the advice policy $\alpha = 0.5$. The solid line represents a fund with $\theta = 0.25$.

to vote for rejection. Conflict may arise in quadrants II and IV. In quadrant II, proposals have a positive social return but a negative financial return. For proposals in quadrant II between the solid line and the dashed line, the fund would like to vote no but the advisor recommends approval: $\alpha > \theta$ implies that the advice policy overweights the social dimension relative to the fund's ideal policy. This misalignment of preferences generates a loss for the fund; the fund casts a yes vote that is ex post the wrong vote. The fund would benefit from a marginal decrease in policy α, bringing it closer to θ. Conversely, for all proposals in quadrant IV between the solid line and the dashed line, the fund would like to vote yes but the advisor recommends reject. Again, the fund would benefit from a marginal decrease in policy α. Recall that the fund only observes the approval/rejection recommendation and so is unable to distinguish these different cases. In equilibrium, the fund rationally updates its belief and optimally chooses to follow the advice it purchased, even knowing it is slanted.

To illustrate the payoff functions, consider two funds that have the same λ but different preferences over policies: fund $\theta_1 = 0.25$ places less weight on the social dimension than fund $\theta_2 = 0.75$. Figure 2 plots the funds' expected utility as a function of different policies. The expected utility functions are strictly quasiconcave, maximized at a fund's preferred policy $\alpha = \theta$, and strictly decreasing as the policy moves away from this bliss point. The expected

utility is strictly positive, hence funds strictly benefit from following the advice.

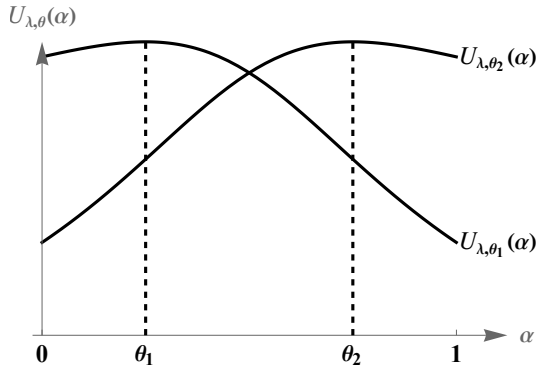


Figure 2: Utility as a function policy α. Proposals are uniformly distributed on a disk, and the two funds have preferences $\theta_1 = 0.25$ and $\theta_2 = 0.75$, with the same λ.

C. Advice and Shareholder Democracy

Different structures of the advice market produce different voting decisions and thus different distributions of election outcomes (and ultimately different corporate policies, although that is only implicit in our model). Here we develop a method to compare and rank the different election outcomes. We later discuss the tradeoffs of our approach versus others.

Each fund is assumed to own a fraction $\tau_i > 0$ of the issuing company's stock, where $\int_0^1 \tau_i di = 1$; note that τ_i is also the fraction of votes that the fund casts. For notational convenience, label the funds in increasing order according to their preferences: that is, if $\theta_i < \theta_{i'}$, then $i < i'$. If more than one fund has the same θ, within this group order funds from smallest to largest according to their ownership share.

Election outcomes are determined by majority voting, so that a proposal is approved if a majority of funds vote yes, where votes are weighted by the fraction of shares owned by a fund. We assume that ties are broken in favor of rejection. Define the median fund as the fund i_m such that $\int_0^{i_m} \tau_i di = 0.5$, and let $\theta_m \in (0, 1)$ be that fund's preference parameter. Note that the median fund is not the median from the distribution G, but the median weighted by the ownership shares.

Remark 1. *If all funds vote and are fully informed about all proposals, then the median fund is decisive. That is, if the median fund prefers to approve a proposal, then a majority of funds also prefer to approve the proposal, and conversely.*[8]

This application of the median voter theorem implies that the median fund's preference is a sufficient statistic to determine the election outcomes under full information. Less obviously, the median's preference is also a sufficient statistic to define how a majority of funds rank different election outcomes under partial information. To see this, consider a market structure k in which different funds may have different amounts of information. This market structure encompasses the number of avisors, the equilibrium strategies of all players, the information available to voters, and the election outcomes. Let the function $q_k(s,r)$ represent the election outcomes in market k, where $q_k(s,r) = 1$ if proposal (s,r) is approved by a majority of votes, and $q_k(s,r) = -1$ if it is rejected. Recall from Section II.A that investor i receives a policy payoff v_{ij} if shareholders approve proposal j and $-v_{ij}$ if they reject it. Therefore, for fund θ, the expected policy payoff from election outcomes q_k is

$$\hat{U}_\theta(q_k) \quad = \quad \theta E[s \cdot q_k(s,r)] + (1-\theta)E[r \cdot q_k(s,r)], \tag{7}$$

where the expectation is taken over all possible values of (s,r) according to F. Comparing across two outcome functions, fund (λ,θ) weakly prefers q_A to q_B if it delivers a higher expected policy payoff, that is, if

$$\theta E[s \cdot q_A(s,r)] + (1-\theta)E[r \cdot q_A(s,r)] \geq \theta E[s \cdot q_B(s,r)] + (1-\theta)E[r \cdot q_B(s,r)]. \tag{8}$$

Hence, we say that fund (λ,θ) weakly prefers market q_A whenever (8) holds.

Remark 2. *Consider two market structures, with election outcomes q_A and q_B. If the median fund weakly prefers q_A to q_B, then a majority of funds weakly prefer q_A to q_B.*[9]

[8] If the median fund prefers to approve a proposal, then $\theta_m s + (1-\theta_m)r > 0$. If $s \geq r$, then for all funds $\theta > \theta_m$, $\theta s + (1-\theta)r > 0$; and if $s < r$, then for all funds $\theta < \theta_m$, $\theta s + (1-\theta)r > 0$. The logic is similar when the median fund prefers to reject a proposal.

[9] If q_A delivers a weakly higher policy payoff for the median fund, $\theta_m E[s \cdot q_A(s,r)] + (1-\theta_m)E[r \cdot q_A(s,r)] \geq \theta_m E[s \cdot q_B(s,r)] + (1-\theta_m)E[r \cdot q_B(s,r)]$, then $\theta_m \{E[s \cdot q_A(s,r)] - E[s \cdot q_B(s,r)]\} + (1-\theta_m)\{E[r \cdot q_A(s,r)] - E[r \cdot q_B(s,r)]\} \geq 0$. The result then follows from the same logic as Remark 1.

If we imagine asking funds to choose between q_A and q_B, the outcome favored by the median fund would attract majority support. Because the information available to voters determines the election outcomes, if a majority of funds prefer q_A to q_B, then we can say that a majority of voters prefer the information that is collectively provided in market A to B. This idea motivates our approach to comparing market structures. For each market structure k with election outcomes q_k, the expected policy payoff for the median voter is

$$\Psi(q_k) \quad = \quad \theta_m E[s \cdot q_k(s,r)] + (1-\theta_m)E[r \cdot q_k(s,r)], \tag{9}$$

where the expectation is taken over all possible values of (s,r) according to F.

Remark 3. *The median voter's utility provides a metric to rank the effectiveness of shareholder democracy under different market structures. If $\Psi(q_A) > \Psi(q_B)$, then a majority of funds receive a higher expected policy payoff under q_A than q_B. In such a case, we say that market A provides more effective shareholder democracy than B, or that "informed representation" is higher in A than B.*

The most informed and representative election outcomes occur when every investor has full information. In that case, election outcomes would be determined by the median fund's preference with full information. The median fund's expected payoff under full information is the same as if it followed an advice policy $\alpha = \theta_m$ because knowing the sign of $\theta_m s + (1-\theta_m)r$ is sufficient to vote correctly. Therefore, the median fund's expected payoff in a market where all funds have full information about proposals is

$$\Psi_{max} \quad \equiv \quad \Psi(q_{\text{ Full Information}}) = \theta_m V_s(\theta_m) + (1-\theta_m)V_r(\theta_m). \tag{10}$$

At the other extreme, if funds have no information at all, they are indifferent between approval and rejection, and by assumption vote no, so that $\Psi(q_{\text{ No Information}}) = 0$.

Our ranking concept here employs a procedural concept of democratic effectiveness. We take as given that majoritarian outcomes are the goal, and use the fully informed majoritarian outcome as the optimal benchmark. We then explore which advisory market structures bring election outcomes closer to that standard. Our focus on majoritarian outcomes has some support in the theoretical literature: May's Theorem shows that majority rule is the unique social welfare function that satisfies a small set of normatively desirable procedural conditions in elections with two outcomes, like those we study (May 1952).

An alternative approach would be utilitarianism: the best outcome is the one that maximizes the sum of utilities. Utilitarianism gives more weight to voters with more intense preferences, which majority rule does not. Because both approaches are used in the literature, throughout the paper we consider both total utility and informed representation,[10]

IV. Monopoly

Consider a market with a single proxy advisory firm. If the firm offers a policy α, then each fund (λ, θ) is willing to pay at most $U_{\lambda\theta}(\alpha)$ for its recommendations, and the firm charges each fund a price $U_{\lambda\theta}(\alpha)$. The firm's profit from offering policy α is

$$\text{Profit}(\alpha) = \int_{(\lambda,\theta)} U_{\lambda\theta}(\alpha) dG(\lambda, \theta). \tag{11}$$

The monopolist maximizes profit by choosing the policy that maximizes the total expected expressive utility of the funds. Intuitively, the expected expressive utility is the amount each fund is willing to pay for information, and because the monopolist engages in price discrimination, it wants to maximize the total willingness to pay.

To characterize the optimal policy, rewrite

$$\begin{aligned}\text{Profit}(\alpha) &= \int_0^1 \int_0^{\overline{\lambda}} \lambda \left[\theta V_s(\alpha) + (1-\theta) V_r(\alpha)\right] g(\lambda, \theta) d\lambda d\theta \\ &= \int_0^1 \left[\theta V_s(\alpha) + (1-\theta) V_r(\alpha)\right] \left[\int_0^{\overline{\lambda}} \lambda g(\lambda, \theta) d\lambda\right] d\theta.\end{aligned}$$

Construct the following auxiliary distribution $\hat{g}(\theta)$, which represents the distribution of parameter θ weighted by λ:

$$\hat{g}(\theta) \equiv \frac{\int_0^{\overline{\lambda}} \lambda g(\lambda, \theta) d\lambda}{E[\lambda]}. \tag{12}$$

This new distribution assigns higher weight to values of θ associated with higher values of λ. The profit function can then be expressed as

$$\begin{aligned}\text{Profit}(\alpha) &= E[\lambda] \int_0^1 \left[\theta V_s(\alpha) + (1-\theta) V_r(\alpha)\right] \hat{g}(\theta) d\theta \\ &= E[\lambda] \left[\hat{E}[\theta] V_s(\alpha) + (1 - \hat{E}[\theta]) V_r(\alpha)\right], \end{aligned} \tag{13}$$

[10] Since our model works at the level of expressive preferences, neither informed representation nor utility maximization capture the downstream effect of policies, such as the health value of cleaning the environment.

where $\hat{E}[\theta] = \int_0^1 \theta \hat{g}(\theta) d\theta$ is the average fund preference, with probability weights adjusted by how much funds care about their voting records.

Because the term $E[\lambda] > 0$ is independent of α, (6) and (13) together imply that profit maximization is equivalent to maximizing the expected expressive utility of a "representative" fund with preference $\hat{E}[\theta]$. The profit maximizing policy targets this average fund, $\alpha^* = \hat{E}[\theta]$. In terms of the original probability distribution,

$$\hat{E}[\theta] = \int_0^1 \theta \hat{g}(\theta) d\theta = \int_0^1 \theta \left[\frac{\int_0^{\overline{\lambda}} \lambda g(\lambda, \theta) d\lambda}{E[\lambda]} \right] d\theta = \frac{E[\lambda\theta]}{E[\lambda]}. \tag{14}$$

By definition, $E[\lambda\theta] = E[\lambda]E[\theta] + cov_{\lambda\theta}$, which leads to our first main result:

Proposition 1. *A monopolist optimally offers an advice policy*

$$\alpha^* \quad = \quad \hat{E}[\theta] = E[\theta] + \frac{cov_{\lambda\theta}}{E[\lambda]}, \tag{15}$$

which produces Profit(α^*)*. This policy maximizes total expected expressive utility of funds. It maximizes informed representation if and only if*

$$E[\theta] + \frac{cov_{\lambda\theta}}{E[\lambda]} = \theta_m.$$

To understand the intuition, suppose $cov_{\lambda\theta} = 0$, so that the advisor selects policy α^* targeting the average fund $E[\theta]$, not the median fund. The advisor is not concerned about obtaining the support of a majority of funds — its concern is to maximize the revenue extracted from funds, which leads to maximizing the average surplus over all funds. Loosely speaking, by slanting its policy from the median fund's ideal policy to the average fund's ideal policy, the firm lowers the surplus of a majority of funds but creates a larger surplus for a minority of funds that are away from the median.[11] Figure 3 depicts an example in which the pdf of θ is right skewed and $\theta_m < E[\theta]$. While the monopolist offers a policy $\alpha^* = E[\theta]$, a majority of funds prefer $\alpha' = \theta_m$, which puts less weight on the social dimension.

When $cov_{\lambda\theta} \neq 0$, because $cov_{\lambda\theta} = \rho_{\lambda\theta}\sigma_\lambda\sigma_\theta$, the equilibrium policy can be written as

$$\alpha^* \quad = \quad E[\theta] + \frac{\rho_{\lambda\theta}\sigma_\lambda\sigma_\theta}{E[\lambda]}. \tag{16}$$

[11] Our result differs from political economy models of elections where candidates target the median voter, and is closer to probabilistic voting models in which candidates target some weighted average of the population. However, unlike probabilistic voting models, our firm is not uncertain about voter preferences.

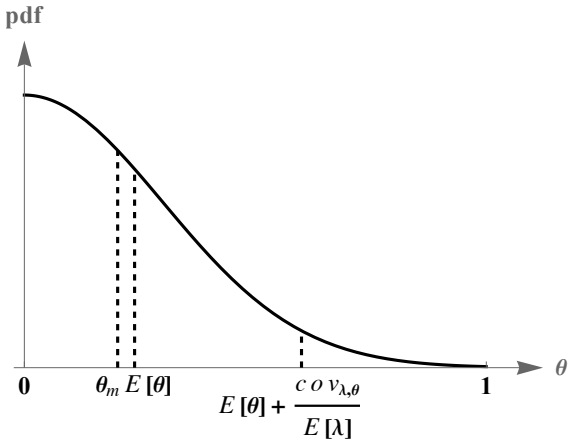


Figure 3: Right-skewed pdf of θ. The median fund has preference θ_m, but the monopolist chooses policy $\alpha^* = E[\theta]$ if preference parameters are independent, and $\alpha^* = E[\theta] + \frac{cov_{\lambda\theta}}{E[\lambda]}$ if $cov_{\lambda\theta} > 0$.

In the example of Figure 3, a positive correlation between θ and λ induces the monopolist to offer a policy that assigns even more weight to the social dimension. The next corollary formalizes this result.

Corollary 1. *A monopolist offers an advice policy α^* that assigns extra weight to the social dimension when socially inclined funds are more likely to care about their voting record ($\rho_{\lambda\theta} > 0$). In such a case, the effect is larger when there is greater heterogeneity in preferences over returns (σ_θ is large), when there is greater heterogeneity in expressive benefits (σ_λ is large), and when expressive benefits are low on average ($E[\lambda]$ is small).*

The corollary highlights an important implication of the model regarding bias in proxy advice. First, it shows how bias – advice recommendations that run against the preferences of a majority of investors – emerges even without managers having a personal preference for bias, and even without conflicts of interest, which are the prevailing explanations for bias. It also links the direction of bias to the distribution of characteristics of advice customers. To the extent that SRI funds are more concerned with their voting record than traditional funds, which focus primarily on financial returns, we expect $\rho_{\lambda\theta} > 0$. This would lead a monopolist to slant its advice policy toward the preferences of the SRI funds.

To characterize the democracy of outcomes in such a market, note that all funds vote according to advice α^*. The equilibrium voting outcome is then $q^*(s,r) = 1$ if $\alpha^* s + (1-\alpha^*)r > 0$, and $q^*(s,r) = -1$ otherwise, implying that

$$\Psi(q^*) \quad = \quad \theta_m V_s(\alpha^*) + (1-\theta_m)V_r(\alpha^*). \tag{17}$$

Because $\Psi(q^*) > \Psi(q_{NoInformation}) = 0$, the voting equilibrium is informative compared to a market in which funds have no information. However, this outcome is different from the full information outcome unless $\alpha^* = \theta_m$, and the loss grows as α^* becomes more distant from θ_m. Finally, from a utilitarian perspective, Proposition 1 shows that a profit-maximizing monopolist chooses the policy that would be selected by a social planner who wants to maximize the total expressive voting utility. By assigning more weight to voters with a higher λ, the monopolist deviates from the full-information majority-rule outcome to an outcome that gives more decision power to voters who care more about their voting record.

V. Duopoly

A. Competitive Equilibrium with Two Proxy Advisory Firms

We now consider two firms competing to provide advice to the market, Firm 1 and Firm 2, choosing advice policies α_1 and α_2. A monopolist that makes a take-it-or-leave-it offer charges each fund its reservation price $U_{\lambda\theta}(\alpha)$. Under a duopoly, advisory firms take into account that funds can purchase advice from the other firm. Clearly, $\alpha_1 = \alpha_2$ cannot be an equilibrium (Bertrand competition would yield a price equal to the marginal cost, which is zero by assumption, resulting in a profit of zero). If the firms offer different policies with $\alpha_1 < \alpha_2$, then there is a cutoff $\tilde{\theta} \in (\alpha_1, \alpha_2)$ such that funds with preference $\tilde{\theta}$ are indifferent between the two advice policies. Using the fact that $U_{\lambda\tilde{\theta}}(\alpha_1) = U_{\lambda\tilde{\theta}}(\alpha_2)$, we can solve for the cutoff value:

$$\tilde{\theta} \quad = \quad \frac{-\left[V_r(\alpha_1) - V_r(\alpha_2)\right]}{\left[V_s(\alpha_1) - V_s(\alpha_2)\right] - \left[V_r(\alpha_1) - V_r(\alpha_2)\right]}. \tag{18}$$

Funds with preference $\theta < \tilde{\theta}$ are willing to pay more for the advice of Firm 1 than Firm 2, and funds with preference $\theta > \tilde{\theta}$ are willing to pay more for the advice of Firm 2 than Firm 1. In a subgame perfect equilibrium, price competition implies that each fund $\theta < \tilde{\theta}$

purchases from Firm 1 and pays a price $U_{\lambda\theta}(\alpha_1) - U_{\lambda\theta}(\alpha_2)$, which is the surplus generated by Firm 1 relative to Firm 2. Each fund $\theta > \tilde{\theta}$ purchases from Firm 2 and pays a price $U_{\lambda\theta}(\alpha_2) - U_{\lambda\theta}(\alpha_1)$. Each advisory firm is able to extract the surplus of its customers to the point that they are indifferent about purchasing the other firm's advice. The ability to switch between advisory firms allows funds to retain some of the surplus.

To characterize policy choices, define the total expressive voting utility in the market as

$$\begin{aligned} \mathcal{U}(\alpha_1, \alpha_2) &\equiv \int_0^1 \int_0^{\overline{\lambda}} \max\{U_{\lambda\theta}(\alpha_1), U_{\lambda\theta}(\alpha_2)\} g(\lambda, \theta) d\lambda d\theta \qquad (19) \\ &= E[\lambda] \int_0^1 \max\{\theta V_s(\alpha_1) + (1-\theta) V_r(\alpha_1), \theta V_s(\alpha_2) + (1-\theta) V_r(\alpha_2)\} \hat{g}(\theta) d\theta, \end{aligned}$$

where $\hat{g}(\theta)$ is defined by (12). This allows us to present the main result of this section:

Proposition 2. *Consider a competitive market with two advisory firms.*

(i) Advice policies (α_1^*, α_2^*) form a competitive equilibrium if and only if

$$\mathcal{U}(\alpha_1^*, \alpha_2^*) \geq \mathcal{U}(\alpha_1, \alpha_2^*) \qquad \text{for all } \alpha_1 \in [0,1], \text{and} \qquad (20)$$

$$\mathcal{U}(\alpha_1^*, \alpha_2^*) \geq \mathcal{U}(\alpha_1^*, \alpha_2) \qquad \text{for all } \alpha_2 \in [0,1]. \qquad (21)$$

Therefore, if the advice policies (α_1^*, α_2^*) maximize total expressive voting utility $\mathcal{U}$, then they also form a competitive equilibrium. Hence, a competitive equilibrium exists.

(ii) If advice policies $\alpha_1^* < \alpha_2^*$ form a competitive equilibrium, then

$$\alpha_1^* = \hat{E}[\theta | \theta \leq \tilde{\theta}] = E[\theta | \theta \leq \tilde{\theta}] + \frac{cov_{\lambda\theta}(\theta \leq \tilde{\theta})}{E[\lambda | \theta \leq \tilde{\theta}]}, \qquad (22)$$

$$\alpha_2^* = \hat{E}[\theta | \theta \geq \tilde{\theta}] = E[\theta | \theta \geq \tilde{\theta}] + \frac{cov_{\lambda\theta}(\theta \geq \tilde{\theta})}{E[\lambda | \theta \geq \tilde{\theta}]}, \qquad (23)$$

$$\tilde{\theta} = \frac{-[V_r(\alpha_1^*) - V_r(\alpha_2^*)]}{[V_s(\alpha_1^*) - V_s(\alpha_2^*)] - [V_r(\alpha_1^*) - V_r(\alpha_2^*)]}. \qquad (24)$$

Moreover, $\alpha_1^* < \alpha^* < \alpha_2^*$, where α^* is the optimal policy of a monopolist.

To understand part (i), fix an advice policy α_2^*. Note that Firm 1 earns $U_{\lambda\theta}(\alpha_1) - U_{\lambda\theta}(\alpha_2^*)$ from each customer. Because $U_{\lambda\theta}(\alpha_2^*)$ is fixed and because Firm 1's profit comes from those funds with $U_{\lambda\theta}(\alpha_1)$ larger than $U_{\lambda\theta}(\alpha_2^*)$, the firm maximizes profit by maximizing the total voting utility of its customers, which implies choosing α_1 that maximizes $\mathcal{U}(\alpha_1, \alpha_2^*)$. The same logic applies to Firm 2.

For part (ii), consider the case $\alpha_1^* < \alpha_2^*$. A monopolist would choose $\alpha^* = \hat{E}[\theta]$ targeted at the weighted average of all funds because it serves the whole market. In a duopoly, Firm 1 serves only those funds with preferences $\theta \in [0, \tilde{\theta}]$, hence its optimal advice policy is $\alpha_1^* = \hat{E}[\theta|\theta \leq \tilde{\theta}]$; and similarly Firm 2 chooses $\alpha_2^* = \hat{E}[\theta|\theta \geq \tilde{\theta}]$. Then $\alpha_1^* < \alpha^* < \alpha_2^*$. Intuitively, while the monopolist targets the average of funds in the entire market, in a duopoly one firm targets the average of funds that place a relatively low value on the social return (funds with $\theta \leq \tilde{\theta}$) while the other targets the average of funds that place a relatively high value on the social return (funds with $\theta \geq \tilde{\theta}$).

In the proof of Proposition 2, we identify conditions on distributions F and G that guarantee equilibrium uniqueness. One case is if s and r are independently drawn from a standard normal distribution and $\hat{g}(\theta)$ is a uniform distribution. When the competitive equilibrium is not unique, there may be an equilibrium in which total voting utility is not maximized, in which case both firms could be trapped at a local maximum that is not a global maximum. We assume from here that the conditions hold for the equilibrium to be unique (or alternatively, that the firms coordinate in the equilibrium that maximizes total expressive voting utility).

A key result in our paper is to show how firms optimally choose locations away from each other. Some Hotelling duopoly models with endogenous location feature too little diversification, with both firms locating at the median, or too much diversification, with firms locating at the extremes. In our model, location choices that maximize expressive voting utility constitute a competitive equilibrium. In this respect, our model parallels that of Lederer and Hurter (1986). They show that given the locations of the other firms, each firm chooses a location that ends up minimizing the total transportation cost, implying that a competitive equilibrium minimizes overall transportation costs. In our model, instead of firms choosing spatial locations and facing transportation costs, firm choose advice policies that endogenously generate different values to different consumers.

B. Intuition for Product Differentiation

To understand how competition drives product differentiation, consider the example in Figure 4, where proposals are distributed on a disk. A monopolist would offer policy α^*. In a

duopoly, if the firms offer policies $\alpha_1^* < \alpha_2^*$, then investors with preference $\theta < \tilde{\theta}$ purchase from Firm 1 and they pay a price equal to the difference in their expected utility from the two policies, and investors with $\theta > \tilde{\theta}$ purchase from Firm 2.

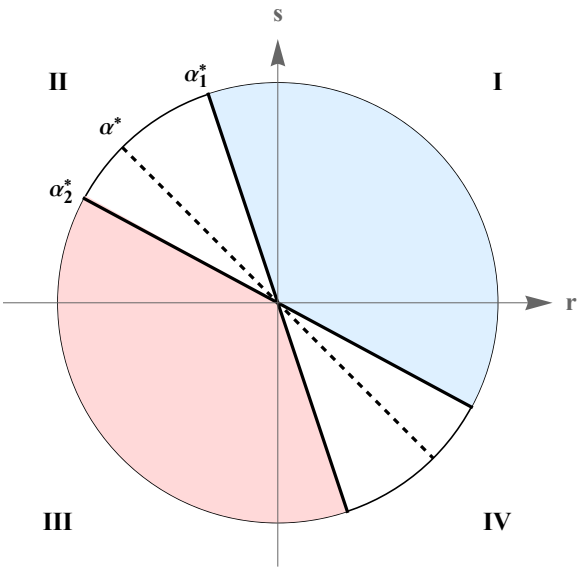


Figure 4: Policies of monopolist (α^*) and duopolists (α_1^*, α_2^*) when proposals are distributed on a disk.

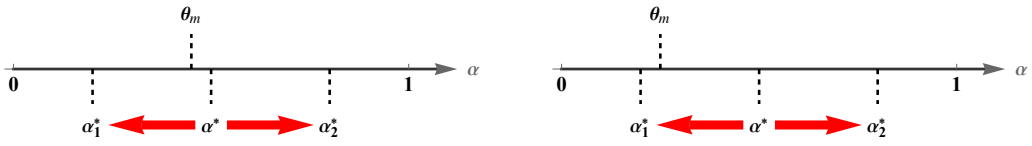
Both firms recommend approval of proposals in the blue-shaded northeast sector, and rejection of proposals in the red-shaded southwest sector. Hence, the firms are not differentiated in these areas — they deliver the same expected utility – and those recommendations add zero to their profit. In the northwest sector between α_1^* and α_2^*, Firm 1 recommends rejection and Firm 2 recommends approval, and the reverse in the southeast sector between α_1^* and α_2^*. Because profits come from offering conflicting recommendations, firms have an incentive to increase the probability of contradictory recommendations. This incentive is offset by by the fact that investor utility declines for the average investor as the policies move apart. The equilibrium reflects a balancing of these two forces.

C. The Effect of Competition on Corporate Elections

How does competition in the advice market affect the outcomes of corporate elections? Are outcomes better informed and more representative with a duopoly or a monopoly? In consumer markets, offering a greater product variety generally increases utility but in this

case differentiated advice can, in some situations, push election outcomes away from the full-information outcome.

To see this, consider the market illustrated by Figure 5a. A monopolist offers a policy α^* very close to the median fund's ideal policy, hence election outcomes closely reflect the median's preference. Under a duopoly, competition drives both advisors to offer policies away from the median. Electoral outcomes are determined by the policy of the advisory firm that commands the largest share of the market, which is now far from the median. Figure 5b shows the other case: with duopolists, the median fund (and a majority of funds) chooses policy α_1^*, which is closer to θ_m than the policy α^* that would be offered by a monopolist. In this case, competition brings election outcomes closer to the fully informed outcome. We can state the following corollary:



(a) Competition deecreases informed representation

(b) Competition increases informed representation

Figure 5: Election Outcomes Under Monopoly vs. Duopoly. The monopolist's policy is α^*; the duopolists' policies are α_1^* and α_2^*; and the median fund's preference is θ_m.

Corollary 2. *Let q^* be the election outcome with a monopolist and q_D^* be the election outcome under duopoly. Elections outcomes more closely correspond to the full-information outcome under duopoly than monopoly if and only if $\Psi(q^*) \leq \Psi(q_D^*)$. Informed representation is higher under duopoly than monopoly if θ_m is sufficiently far from α^*, and lower if θ_m is sufficiently close to α^*.*

The corollary follows from the fact that $\alpha_1^* < \alpha^* < \alpha_2^*$ (see Proposition 2). If the median fund's preference θ_m is sufficiently close to the monopolist's policy α^*, then the median fund strictly prefers the monopolist's policy α^* over both of the duopoly policies α_1^* and α_2^*, and

competition reduces informed representation. Conversely, if θ_m is sufficiently far from α^*, then the median fund strictly prefers one of the duopoly advice policies over the monopolist's policy, and this duopoly firm dictates the voting outcome.

In terms of prices, moving from monopoly to duopoly has two potentially offsetting effects. The existence of a second option limits the price that each advisory firm can charge, but the change in available advice policies alters a fund's willingness-to-pay, which can push its price up or down. For funds with preferences sufficiently close to the monopolist's advice policy, the two effects work in the same direction: they pay a lower price under duopoly because they receive worse advice and benefit from the second option. For funds with preferences further from α^*, they may experience higher or lower prices depending on parameters.

Proposition 2 also implies that changing from a monopoly to a duopoly strictly increases the total expressive voting utility. Consequently, although competition might lower the expressive voting utility of some investors, on average it increases expressive voting utility.

We conclude this section with what could be seen as a paradox. Consider the policies represented in Figure 5a. Suppose that the distribution of θ is such that most of the funds are concentrated around the two extreme points $\{0, 1\}$. When we move from the monopoly to duopoly, the vast majority of funds are able to purchase advice that is closer to their preference. Therefore, each fund close to zero or close to one is able to make a more informed voting decision. Although these funds cast votes that are better aligned with their preferences, as a group the funds make a worse collective decision, in the sense that informed representation deteriorates: a majority of funds would prefer the monopoly market over the duopoly market. This apparent paradox goes back to the classic idea that while information helps a rational individual decision maker, information can be detrimental to a group of individuals in a market, e.g., Hirshleifer (1971). In the voting model of Alonso and Câmara (2016), a single sender strategically uses information to persuade voters and this information can result in voting outcomes that are worse for a majority of voters. In our model, the result is driven not by the actions of one strategic sender, but by competition between the firms providing information and their detrimental product differentiation.

VI. More Than Two Advisory Firms

We now consider the general case of $N \geq 2$ firms to show that our model remains tractable, and qualitative insights from the duopoly case extend to markets with many firms. We also characterize market outcomes in the perfect competition limit, as the number of firms grows large.

In equilibrium, as before, firms never offer identical policies – they segment the market. Consider the policy profile in which firms are ordered by their policies, $\alpha_1 < \ldots < \alpha_N$. Analogous to the duopoly case, between each pair of adjacent firms α_k and α_{k+1}, there is a cutoff fund $\tilde{\theta}_{k,k+1} \in (\alpha_k, \alpha_{k+1})$ that is indifferent between the advice provided by the two firms. Generalize the total expressive utility in the market (19) as:

$$\mathcal{U}(\alpha_1, \ldots, \alpha_N) \equiv \int_0^1 \int_0^1 \max\{U_{\lambda\theta}(\alpha_1), \ldots, U_{\lambda\theta}(\alpha_N)\} g(\lambda, \theta) d\lambda d\theta, \tag{25}$$

and denote $\mathcal{U}(\alpha_k|\alpha_{-k})$ as the total utility as a function of the policy α_k, given a fixed vector α_{-k} of policies of the other firms. We can then characterize the equilibrium:

Proposition 3. *Consider a competitive market with $N \geq 2$ advisory firms.*

(i) Advice policies $(\alpha_1^*, \ldots, \alpha_N^*)$ form a competitive equilibrium if and only if, for all k,

$$\mathcal{U}(\alpha_k^*|\alpha_{-k}^*) \geq \mathcal{U}(\alpha_k|\alpha_{-k}^*) \qquad \text{for all } \alpha_k \in [0,1]. \tag{26}$$

Therefore, if advice policies $(\alpha_1^*, \ldots, \alpha_N^*)$ maximize total expressive voting utility $\mathcal{U}$, then they also form a competitive equilibrium. Hence, a competitive equilibrium exists.

(ii) If advice policies $\alpha_1^* < \ldots < \alpha_N^*$ form a competitive equilibrium, then:

$$\tilde{\theta}_{k,k+1} = \frac{-[V_r(\alpha_k) - V_r(\alpha_{k+1})]}{[V_s(\alpha_k) - V_s(\alpha_{k+1})] - [V_r(\alpha_k) - V_r(\alpha_{k+1})]}, \tag{27}$$

where we define $\tilde{\theta}_{0,1} = 0$ and $\tilde{\theta}_{N,N+1} = 1$; and

$$\alpha_k = \hat{E}[\theta|\tilde{\theta}_{k-1,k} \leq \theta \leq \tilde{\theta}_{k,k+1}] = E[\theta|\tilde{\theta}_{k-1,k} \leq \theta \leq \tilde{\theta}_{k,k+1}] + \frac{cov_{\lambda\theta}(\tilde{\theta}_{k-1,k} \leq \theta \leq \tilde{\theta}_{k,k+1})}{E[\lambda|\tilde{\theta}_{k-1,k} \leq \theta \leq \tilde{\theta}_{k,k+1}]}. \tag{28}$$

(iii) Let α_m^* be the policy of the advisory firm that sells to the median voter.[12] Then the advisory firm α_m^* is decisive for election outcomes. That is, in equilibrium, each proposal

[12] Note that α_m^* is not defined as the median policy in the set $\{\alpha_1^*, \ldots, \alpha_N^*\}$, but as the preferred policy of the median fund θ_m considering the options in the set.

is approved if and only if it receives an approval recommendation by the advisory firm with policy α_m^*.

Figure 6 illustrates a market with four advisory firms. The market is segmented, with segment sizes determined by the distributions F and G, and each advisory firm chooses an advice policy to maximize the weighted expected utility of funds in the segment.

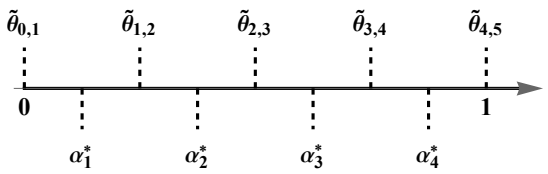


Figure 6: Example of a market with four advisory firms. Firm k serves consumers with preferences $\theta \in [\tilde{\theta}_{k-1,k}, \tilde{\theta}_{k,k+1}]$ and offers policy $\alpha_k^* = \hat{E}[\theta|\tilde{\theta}_{k-1,k} \leq \theta \leq \tilde{\theta}_{k,k+1}]$.

As in the duopoly case, adding another advisory firm always increases the total expressive utility. Entry's effect on the informativeness and representation of corporate elections is ambiguous, however. Entry causes the existing firms to move their advice policies away from the position of the entrant. As a result, entry of a new firm can move the vote of the median fund closer to or away from the fund's fully informed vote, and therefore, entry can increase or reduce the informed representation of corporate elections.

Although the effect of an individual entrant is ambiguous, the effect in the limit is unambiguous. In the limit, the market reaches a "perfectly competitive" outcome: a continuum of advisory firms, with each firm serving one type of fund, and providing that type's ideal advice policy. Perfect competition produces identical election outcomes as a market in which all funds are fully informed, or in our terminology, maximizes informed representation.

Corollary 3. *Under "perfect competition" (as the number of advisory firms reaches the number of fund types), the firm that sells to the median fund offers an advice policy* $\alpha_m^* = \theta_m$, *and informed representation attains its highest possible value,* Ψ_{max}. *Perfect competition also maximizes the expressive voting utility of each fund. All advisory firms charge a price of zero (marginal cost).*

VII. Extensions and Robustness

In the Online Appendix, we extend the model to consider proposals with more than two dimensions and we allow each advisor to choose any binary approve/reject signal. The model remains tractable and our qualitative insights continue to hold. Here we present some extensions that provide additional insights.

A. Customized Advice

One of our findings is that, unless there are many advisory firms, advice is slanted and the results of corporate elections may be different from what would happen if all funds were fully informed. A natural question is whether this would happen if advisory firms offered different advice to different customers, for example, if they sold advice based on two or more different policies. In fact, proxy advisory firms do offer some customized options, although not enough to span the space.[13] Here we extend the model to allow customization.

As before, there are N advisory firms, but now each firm k offers a number Z_k of advice policies. Both the number of firms and the number of policies per firm are fixed exogenously. Let $Z \equiv \sum_{k=1}^{N} Z_k$ be the number of policies in the market; let $\alpha^k = (\alpha_1^k, \ldots, \alpha_{Z_k}^k)$ be the menu of policies offered by firm k, where α_z^k is a generic policy in this set; and let α^{-k} be the menu of policies offered by all firms not including k. Firms simultaneously select and post their menu of policies; each firm then quotes a price for each of its policies to each fund; and each fund then chooses which policy to purchase. This finite set of different policies represent the observed fact that ISS offers seven "specialty policies" (e.g., "Sustainability Policy" and "Socially Responsible Policy"), while Glass Lewis offers nine "thematic policies" (e.g., "Climate Policy" and "Corporate Governance Focused Policy").[14] We continue to assume

[13] While proxy advisors have long claimed that they offer a substantial amount of customization, and Hu et al. (2024) show that many Glass Lewis customers obtain customized advice, many observers have been skeptical, especially with respect to ISS, and McCahery et al. (2016) document that thirty percent of investors believe proxy advisors' advice is too standardized. Glass Lewis' recent announcement to stop issuing single voting positions on issues ("We recognize that a single perspective is not longer sufficient") implicitly concedes that the amount of customization was not high (White 2025).

[14] Each policy comes with a long document explaining how the advisor generates the recommendations. Some of these documents are one hundred pages long. Our implicit assumption is that it is very costly for

that a fund can purchase only a single policy.

In the Online Appendix, we show that a version of Proposition 3 holds for this general market. The key result is the following: given the policies α^{-k} offered by the other firms, firm k chooses the policy menu α^k that maximizes the total expressive voting utility. If the vector of policies $\boldsymbol{\alpha}^* = (\alpha^{1^*}, \ldots, \alpha^{N^*})$ maximizes total expected voting utility, then there is an equilibrium in which the policies offered by the firms equal $\boldsymbol{\alpha}^*$. Each equilibrium policy α_z^k targets the average consumer that purchases this policy, where the average is weighted by how much the funds care about their voting records. The policy purchased by the median fund is decisive for election outcomes.

Because Proposition 3 can be restated in this way, the main implications remain the same under this extension. In particular, in terms of election outcomes, if the policy vector $\boldsymbol{\alpha}^*$ maximizes the total expressive voting utility, then the election outcomes produced in a market with N firms that offer a single policy are identical to those produced by a market with N' firms offering multiple policies that sum to $Z = N$ total policies. There is one important difference, however: when advisory firms offer multiple policies, funds pay a weakly higher price than they pay in a market where advisory firms offer only a single policy. To see this, note that in both markets a fund purchases the policy that yields it the highest expressive utility. The price paid in the single-policy market is the difference between this utility and the utility from the second-best policy option among all of the other $N - 1$ policies available. In the multiple-policy market, the price is the difference between the highest utility and the second-best utility excluding the other policies offered by the firm that sells to the fund. Since this second-best utility is weakly lower, the price is weakly higher. It follows that proliferation of customized policies by a small group of firms is not a substitute for competition between a large number of single-policy firms – customization can yield the same distribution of advice policies but it delivers a larger share of the surplus to suppliers relative to consumers.

advisors to implement and offer many different policies. For instance, in 2024, ISS entered in a partnership with Bowyer Research to offer a new policy using the approach of an "ESG skeptic" (Kerber, 2024). We interpret this as evidence that different αs might require a specialized research team, which can be an in-house team at the advising firm, or an external specialist.

B. Internal Information Production and the Decision to Vote

The benchmark model assumes that funds are uninformed unless they purchase advice from an advisory firm. In practice, some funds produce information internally, either by conducting research specifically about election items, or as a by-product of their normal investment research. Here we extend the model to allow funds to be self-informed, and explore how existence of advisory firms affects shareholder democracy. We abstract away from the crowding-out problem studied by Malenko and Malenko (2019) by assuming that funds are exogenously endowed with internal information; this allows us to bring out a force that has not been recognized previously: existence of an advice market can shift the position of the median voter, and possibly reduce the informed representation of elections.

In this extension, each fund is characterized by a new parameter $\gamma \in [0, 1)$ that represents the fund's internally produced information. Specifically, fund i with internal information γ exogenously learns the value v_{ij} of a measure γ of proposals (recall that there is a unit mass of proposals). We also assume, primarily to simplify the analysis, that if a fund with internal information γ learns the value of a proposal j, then all funds with internal information $\gamma' \geq \gamma$ also learn the value of proposal j. This implies that the set of proposals with values known by a fund with a smaller γ is a subset of the proposals known by a fund with a larger γ. The idea is in the spirit to Garicano (2000), where workers are ordered by their ability to solve problems of different levels of complexity.

Internal information reduces a fund's willingness to pay for advice, but does not change how it ranks different recommendation policies. From the viewpoint of advisory firms, a market in which funds have internal information γ and preference intensity λ is isomorphic to a market in which funds have no internal information but a lower preference intensity $\tilde{\lambda} \equiv (1 - \gamma)\lambda$. With this observation, the original distribution $G(\lambda, \theta, \gamma)$ can be restated as a new distribution $\tilde{G}(\tilde{\lambda}, \theta)$, and therefore profit functions and first-order conditions are isomorphic under the new and extended model.

Turning to election outcomes, we imagine there is a small cost of voting so that a fund abstains if it has no information and no advice (the Online Appendix presents a formal extension of the model with costly voting). Consider the following example: Funds are divided into two groups, uninformed funds, $\gamma_L = 0$, and partially informed funds, $\gamma_H \in$

$(0,1)$, with uninformed funds the majority. Let $\theta_m^H \in (0,1)$ be the median fund in the partially informed subset, and let $\theta_m \in (0,1)$ be the median fund in the population as before. Without an advice market, for a mass γ_H of proposals, informed funds know their own v_{ij} and vote according to their preferences, and uninformed funds abstain. The electoral outcome coincides with the vote of the subgroup's median fund θ_m^H. For a mass $1-\gamma_H$ of proposals, all funds are uninformed and abstain, and the proposals are rejected, yielding an expected policy payoff of zero (recall that this is equivalent to randomly selecting an outcome, since all funds are indifferent between approval and rejection). The informed representation in this market is then

$$\gamma_H \left[\theta_m V_s(\theta_m^H) + (1-\theta_m) V_r(\theta_m^H)\right]. \tag{29}$$

If a monopolist advisory firm is added to this market, it chooses α^* as before. Because uninformed funds constitute a majority and all of them follow the monopolist's advice, the election outcome is determined by the monopolist's advice. Therefore, the level of informed representation is $\theta_m V_s(\alpha^*) + (1-\theta_m) V_r(\alpha^*)$. The monopolist weakly increases informed representation if and only if

$$\theta_m V_s(\alpha^*) + (1-\theta_m) V_r(\alpha^*) \quad \geq \quad \gamma_H \left[\theta_m V_s(\theta_m^H) + (1-\theta_m) V_r(\theta_m^H)\right]. \tag{30}$$

Inequality (30) brings out the key tradeoffs. The left-hand side is strictly positive for all $\alpha^* \in [0,1]$, and the right-hand side is strictly positive for all $\gamma_H \in (0,1)$ and $\theta_m^H \in (0,1)$. Therefore, the inequality holds for sufficiently low γ_H. Even if the subgroup of informed voters have a preference θ_m^H very close to the overall median θ_m, and the monopolist offers a policy α^* very far from the median, it is better to have a monopolist if the level of internal information is low. In this case, the monopolist's biased policy deteriorates the election outcome for a small fraction γ_H of proposals, but improves the outcome for the remaining proposals. Conversely, if θ_M^H is better for the median voter than α^*, then the monopolist reduces informed representation if γ_H is sufficiently high.

The potential distortion arising from existence of an advice market is different than the traditional crowding-out effect. The crowding-out effect is that the presence of advisory firms causes funds to acquire less internal information, or to disregard their internal information, resulting in votes that are based on less information. In our model, the amount of internal

information is fixed by assumption. When there is no advice market and voting is entirely based on internal information, elections are determined by the median of the self-informed funds. When an advice market exists, the median fund may change, and become one that purchases advice. If advice is slanted toward the preferences of funds with high-expressive benefits, then corporate elections become slanted in that direction.

VIII. Empirical and Policy Implications

A. Slant in Advice

One contribution of our model is a theory of advice that can explain the apparent ideological slant in proxy advisor recommendations (Bolton et al. 2020). Existing explanations of advice bias emphasize ideological preferences of advisory firm managers, or conflicts of interest that lead advisory firms to favor issuers that purchase their other services (Li 2018; Ma and Xiong 2021; Business Roundtable 2025). The implications of our theory differ in several respects.

First, our theory predicts a slant in advice even if the managers of advisory firms are not ideological and have no conflicts of interest. And our theory predicts the direction of the slant: equations (15), (22), and (23) imply that a proxy advisor's recommendations are slanted toward the preferences of customers with the highest expressive benefits. We conjecture that these customers include SRI funds that use their voting record as a marketing tool and public pension funds that are under scrutiny by state officials.

The model also implies that if the covariance between fund intensity and policy preference is positive, an increase in average expressive intensity $E[\lambda]$ reduces bias — see (16) and Corollary 1. If the expressive value of voting increases with external scrutiny, as seems plausible, then heightened scrutiny reduces the slant in advice. To the extent that investors, journalists, and politicians have paid more attention to fund voting in recent years, the model predicts less slant. Similarly, the SEC's recent enhanced vote disclosure rule (SEC 2022) that makes "proxy voting records more usable and easier to analyze, improving investors' ability to monitor how their funds vote and compare different funds' voting records" would reduce slant, according to the model.

Our model also links slant to competition. An immediate prediction is that the slant of different proxy advisors does not converge – which appears to be the case (Bolton et al. 2020)

– and conflicting recommendations are a feature of equilibrium. Another implication concerns the effect of new entry. According to the conflict-of-interest theory, bias falls after new entry because self-serving recommendations are more likely to be detected with an increased supply of information (Li 2018). In contrast, our model predicts that new entry will cause existing firms to slant their recommendations away from the new entrant. Consistent with our prediction, Li (2018) finds that ISS recommendations tilted more in favor of left-leaning shareholder activists after Glass Lewis entered with centrist recommendations in 2003. One could also imagine testing the effect of entry or exit of Egan Jones, ProxyGovernance, Segal Marco Advisors, or Strive Asset Management.

Our model also has implications for disagreement in recommendations among competing proxy advisors. In Figure 4, duopolists issue opposing recommendations in the non-shaded regions, which occur in quadrants II and IV. Consider corporate governance proposals compared to environmental and social (E&S) proposals. We conjecture that governance proposals load primarily on the r dimension and have $s \approx 0$, meaning that they mainly fall along the r-axis, and therefore are unlikely to be in the disagreement regions. E&S proposals, on the other hand, contain significant nonzero r and s components, making them more likely to appear in the disagreement regions. The empirical prediction is that ISS and Glass Lewis are more likely to disagree on E&S than corporate governance issues. Table 1 reports the frequency of disagreement between their recommendations across a range of governance and E&S issues for the period 2008-2021. While not a formal test, disagreement is generally higher on E&S issues (46 percent) than governance issues (16 percent), as predicted. This pattern, which to our knowledge is new to the literature, lends some support to the model and illustrates how it might be applied to new empirical issues.

B. Robo-voting and Fund Preferences

In our baseline model, funds that pay for advice always follow the recommendations when voting. The practice of fully delegating a fund's vote to a proxy advisor is called "robo-voting" and is common. Matsusaka and Shu (2024), for example, show that in 2021, one-third of mutual funds voted according to their advisor's recommendation on 99 percent or more election items. Robo-voting is facilitated by the use of voting platforms that pre-

Table 1: Opposing Recommendations from ISS and Glass Lewis

Proposal type	Number	% ISS-GL Disagree
Environment and Social (total)	**780**	**46%**
Animal rights	61	3%
Board diversity	47	26%
Climate change	91	40%
Equal employment opportunity	139	58%
Gender pay gap	77	26%
Greenhouse gases	169	62%
Human rights	52	56%
Sustainability	144	53%
Governance (total)	**3,474**	**16%**
Board classification	695	0%
Board size	31	0%
Cumulative voting	149	33%
Majority vote	328	1%
Independent chair	527	43%
Poison pill	279	11%
Proxy access	354	36%
Special meeting	502	19%
Supermajority	609	3%
All Environmental, Social, and Governance	**4,254**	**22%**

Source: ISS recommendations from ISS Voting Analytics; Glass Lewis recommendations from a Public Records Law request. The sample contains shareholder proposals from regular annual meetings in which both firms issued recommendations, from 2008 to 2021.

populate an advisory firm's recommendations on the thousands of issues facing a fund, so that the fund can simply "press a button" to cast its votes.

One justification for robo-voting is that it is too costly for funds to collect and process the information required to cast votes that accurately reflect their preferences on thousands of items (Matsusaka and Shu 2024). Our model provides another rationale: casting a vote aligned with proxy advice can reflect a fund's preference better than uninformed voting, even if the proxy advice is known to be slanted. In our extension with self-information, funds follow their own information on the occasions when they collect it, but continue to follow proxy advice on all other issues.

The prevalence of robo-voting (or near robo-voting) creates challenges for estimating fund preferences. Spatial models, such as those used in Bolton et al. (2020), compress a large

number of votes cast by a fund into a lower-dimensional "preference" parameter. Ideally, these methods would recover a fund's structural preferences, but to the extent that funds vote based on slanted advice – which our model predicts is common – their votes are not accurate indicators of their preferences. Rather, their votes represent a compounding of their true preferences and the slant of advice. If advice is slanted toward the preferences of a particular subset of funds – say SRI funds – the ideal-point estimates will produce preferences that are more aligned with those of SRI funds than the true preferences. Estimates of fund preferences are likely to be more accurate as the number of proxy advisors increases and funds are able to purchase advice that is closer to their preferences.

C. Antitrust and Concentration

The American proxy advice industry, essentially a duopoly, increasingly attracts antitrust scrutiny. U. S. House committees in 2013, 2023, and 2025 investigated if ISS and Glass Lewis were violating antitrust laws (U.S. House 2013; U. S. House Committee on the Judiciary 2023, 2025). The committees expressed concern that the companies' market power allowed them to collude with activists to impose policies on companies, and to extort monopoly prices in their related consulting businesses.

Our model suggests that a lack of competition does pose a risk that the preferences of a small group of activists will be imposed on issuers contrary to the preferences of a majority of shareholders. However, this doesn't happen through collusion between advisory firms and activists, but through competitive product differentiation. A prohibition on collusion would not address the problem because it stems from the dynamics of profit maximization. A natural solution, in terms of our model, would be to increase the number of advisory firms. With a sufficiently large number of firms, corporate elections reflect the informed preferences of a majority of funds, and are not slanted toward the preferences of an activist minority. Competition also reduces the prices charged by advisory firms, allaying the traditional concern with market power. The lack of competition in the existing market suggests the presence of barriers to entry – we conjecture that these might stem from fixed costs of collecting information and establishing a vote execution platform.[15]

[15]The CEO of a small proxy advisory firm Egan-Jones testified that "there's little question that there's

It is interesting to note that an increase in the average expressive benefit from voting $E[\lambda]$ may have a pro-competitive effect. If average expressive benefit increases, funds are willing to pay more for advice. To the extent that entry is deterred by substantial fixed costs, the higher willingness to pay would encourage new entry. A collateral effect of new regulations such as the SEC's enhanced disclosure rules, which make voting records more accessible to the public, then might be to enhance competition.

D. Mandatory Voting

One of the most consequential shareholder democracy regulations of the 21st century was the SEC's decision to (in effect) require all mutual funds to vote, which came on the heels of a Labor Department regulation that required pension funds to vote.[16] While the regulations were framed in terms of a fund's fiduciary duty, presumably they were also intended to improve the quality of corporate elections – if all funds participate, one might expect the full array of preferences to be registered, and if votes are also required to be informed, the dispersed information of all shareholders would be incorporated into outcomes.

Our model offers a perspective on the effects of mandatory voting. The basic calculus of voting pushes most funds in the direction of abstaining: because their votes are unlikely to be pivotal, their gain from investing in information and voting exceeds their private benefit. Suppose there is a cost of voting, so that if not required to vote, funds that care little about their voting records (sufficiently low λ) do not purchase proxy advice or vote. Also assume, as in the extension above, that some funds have internal information γ, so that advisory firms offer policies that target funds with high $\tilde{\lambda} \equiv (1-\gamma)\lambda$. Election outcomes are then

---

massive barriers [to] anybody entering the business," specifically singling out voting platforms: "from our perspective, it's critical to get on that platform. But you know what? We can't get on that platform. We simply can't." (U. S. Securities and Exchange Commission 2018, pages 204-205).

[16]SEC: Proxy Voting by Investment Advisors, 68 Federal Register 6585, February 7, 2003. Department of Labor: Letter from Alan D. Lebowitz, Deputy Assistant Secretary, Pension and Welfare Benefits Administration of the U.S. Department of Labor, to Helmuth Fandl, Chair of the Retirement Board, Avon Products, Inc., February 23, 1988. Investment advisors, strictly speaking, are not required to vote their proxies but 90 percent of them choose to do so (SEC Staff Legal Bulletin No. 20, June 30, 2014; Broadridge + PwC, 2019). It is widely understood that funds can satisfy their fiduciary responsibilities by following the recommendations of a proxy advisor.

determined by three factors: the internal information available to informed funds that vote, the slant of advice provided by advisory firms, and the number of funds that choose to vote.

In this environment, a regulatory mandate to vote can improve informed representation, but can also harm it. To see how mandatory voting can be harmful (the non-obvious case), suppose that the median preference of internally informed funds is similar to the overall median preference, the number of uninformed funds that vote according to external advice is small, and there are many funds with a low λ that do not purchase advice and abstain from voting. In this case, election outcomes are very close to the overall median preference and informed representation is high. Although relatively few funds vote, the majority are informed, and their preferences are representative of the population.

If funds are required to vote, and their votes must be based on some information, the large group of funds with a low λ that previously abstained now purchases advice and vote according to the advice. Since this new group of customers has a low λ, advisory firms do not change their policies much. Consequently, election outcomes are no longer determined by the internally informed funds with preferences close to the overall median, but by the advice policy of advisory firms. To the extent that the advice over weights the preferences of funds with high λ, forcing funds to vote might decrease informed representation.

More generally, our argument suggests that "forcing" funds with a low expressive value of voting to vote may not be the best policy. These funds will simply follow the recommendations from slanted advisors that put little weight on their preferences.

Appendix A. Proofs

Proof of Lemma 1: We prove the lemma by proving a series of claims.

Claim 1) $V_s(\alpha) \geq 0$ and, for all $\alpha \in (0,1)$, we have $V_s'(\alpha) > 0$.

To prove the claim, take the derivative of (4)

$$\begin{aligned} V_s'(\alpha) &= \int_{-\infty}^{\infty} \left[\frac{r}{\alpha^2} \frac{(1-\alpha)r}{\alpha} f\left(-\frac{(1-\alpha)r}{\alpha}, r \right) + \frac{r}{\alpha^2} \frac{(1-\alpha)r}{\alpha} f\left(-\frac{(1-\alpha)r}{\alpha}, r \right) \right] dr \\ &= (1-\alpha)\frac{2}{\alpha^3} \int_{-\infty}^{\infty} \left[r^2 f\left(-\frac{(1-\alpha)r}{\alpha}, r \right) \right] dr \\ &= (1-\alpha)K(\alpha) > 0, \end{aligned} \tag{A.1}$$

where

$$K(\alpha) \equiv \frac{2}{\alpha^3}\int_{-\infty}^{\infty}\left[r^2 f\left(-\frac{(1-\alpha)r}{\alpha}, r\right)\right]dr. \tag{A.2}$$

The term $K(\alpha)$ is strictly positive because the integral is over the pdf of a line that goes through the origin. We assumed that F has full support on a set and $(0,0)$ belongs to the interior of this set. Therefore, f is strictly positive around the origin.

When $\alpha = 0$, the firm recommends approval if and only if $r > 0$, therefore $V_s(0) = \Pr(r > 0)E[s|r > 0] - \Pr(r \leq 0)E[s|r \leq 0]$. We assumed that $E[s|r > 0] \geq 0$, which together with assumption $E[s] = 0$ implies that $E[s|r \leq 0] \leq 0$. Therefore, $V_s(0) \geq 0$; since this is an increasing function, $V_s(\alpha) \geq 0$, completing this claim.

Claim 2) $V_r(\alpha) \geq 0$ and, for all $\alpha \in (0,1)$, $V_r'(\alpha) < 0$.

To prove the claim, take the derivative of (5)

$$\begin{aligned} V_r'(\alpha) &= \int_{-\infty}^{\infty}\left[-\frac{s}{(1-\alpha)^2}\frac{\alpha s}{(1-\alpha)}f\left(s, -\frac{\alpha s}{(1-\alpha)}\right) - \frac{s}{(1-\alpha)^2}\frac{\alpha s}{(1-\alpha)}f\left(s, -\frac{\alpha s}{(1-\alpha)}\right)\right]ds \\ &= -2\frac{\alpha}{(1-\alpha)^3}\int_{-\infty}^{\infty}\left[s^2 f\left(s, -\frac{\alpha s}{(1-\alpha)}\right)\right]ds. \end{aligned}$$

Perform a change in variables $s = -(1-\alpha)r/\alpha$:

$$\begin{aligned} V_r'(\alpha) &= -2\frac{\alpha}{(1-\alpha)^3}\int_{\infty\alpha/(1-\alpha)}^{-\infty\alpha/(1-\alpha)}\left[\left(\frac{-(1-\alpha)r}{\alpha}\right)^2 f\left(\frac{-(1-\alpha)r}{\alpha}, r\right)\right]d\frac{-(1-\alpha)r}{\alpha} \\ &= -2\frac{\alpha}{(1-\alpha)^3}\int_{\infty}^{-\infty}\left[\frac{-(1-\alpha)}{\alpha}\frac{(1-\alpha)^2 r^2}{\alpha^2}f\left(\frac{-(1-\alpha)r}{\alpha}, r\right)\right]dr \\ &= -2\frac{\alpha}{\alpha^3}\int_{-\infty}^{\infty}\left[r^2 f\left(\frac{-(1-\alpha)r}{\alpha}, r\right)\right]dr = -2[0.5\alpha K(\alpha)] = -\alpha K(\alpha). \end{aligned} \tag{A.3}$$

This implies $V_r'(\alpha) < 0$. When $\alpha = 1$, the firm recommends approval if and only if $s > 0$, therefore $V_r(1) = \Pr(s > 0)E[r|s > 0] - \Pr(s \leq 0)E[r|s \leq 0]$. We assumed that $E[r|s > 0] \geq 0$, which together with assumption $E[r] = 0$ implies that $E[r|s \leq 0] \leq 0$. Therefore, $V_r(1) \geq 0$; since this is a decreasing function, $V_r(\alpha) \geq 0$, completing this claim.

Claim 3) Fund (λ, θ) is willing to pay $U_{\lambda\theta}(\alpha)$ for advice α. Given the prior $E[s] = E[r] = 0$, every fund receives an expected utility of zero if it does not purchase information. Claims 1 and 2 imply that if the fund purchases the advice and follows the voting recommendation, then it receives a weakly positive payoff $U_{\lambda\theta}(\alpha) = \lambda\{\theta V_s(\alpha) + (1-\theta)V_r(\alpha)\} \geq 0$. Therefore, the value of information is $U_{\lambda\theta}(\alpha) \geq 0$.

Claim 4) Fund (λ, θ) benefits from voting according to the recommendation provided by advice policy α. This result follows immediately from the previous result that the value of information is positive. Because $E[s] = E[r] = 0$ and following the recommendation yields a positive payoff $U_{\lambda\theta}(\alpha) \geq 0$, voting against the recommendation must yield a weakly negative payoff — strictly negative if $U_{\lambda\theta}(\alpha) > 0$.

Claim 5) $U'_{\lambda\theta}(\alpha) > 0$ if $\alpha < \theta$, and $U'_{\lambda\theta}(\alpha) < 0$ if $\alpha > \theta$.

Compute the marginal value of an advice and simplify by using (A.1) and (A.3):

$$\begin{aligned} U'_{\lambda\theta}(\alpha) &= \lambda\left[\theta V'_s(\alpha) + (1-\theta)V'_r(\alpha)\right] \\ &= \lambda\left[\theta(1-\alpha)K(\alpha) - (1-\theta)\alpha K(\alpha)\right] \\ &= \lambda\left[\theta - \alpha\right]K(\alpha). \end{aligned} \tag{A.4}$$

Therefore, $U'_{\lambda\theta}(\alpha) > 0$ if $\alpha < \theta$, and $U'_{\lambda\theta}(\alpha) < 0$ if $\alpha > \theta$. This implies that $U_{\lambda\theta}(\alpha)$ is maximized at $\alpha = \theta$, concluding the proof. □

Proof of Proposition 1: In the text. □

Proof of Proposition 2: We start by proving the following claim: There exists policies $\alpha_1^* < \alpha_2^*$ that maximize total expressive voting utility $\mathcal{U}$. If advice policies $\alpha_1^* < \alpha_2^*$ maximize $\mathcal{U}$, then conditions (22), (23), and (24) hold.

To see this, note that the total expressive utility function $\mathcal{U}(\alpha_1, \alpha_2)$ is continuous on the compact interval $(\alpha_1, \alpha_2) \in [0,1]^2$, therefore a maximum exists. Let (α_1^*, α_2^*) be a maximum. It cannot be the case that $\alpha_1^* = \alpha_2^*$ is a maximum because changing one of the policies would strictly increase $\mathcal{U}$. Therefore, it must be the case that $\alpha_1^* \neq \alpha_2^*$. By symmetry, if some pair $\alpha_1^* > \alpha_2^*$ maximizes $\mathcal{U}$, then the pair $\alpha_1' = \alpha_2^*$ and $\alpha_2' = \alpha_1^*$ also maximizes the function, completing the first part of the proof.

Fix any pair $\alpha_1^* < \alpha_2^*$ that maximizes $\mathcal{U}$, which implies that $\alpha_2^* > 0$. Given α_2^*, it must be the case that

$$\alpha_1^* \in \arg\max_{\alpha_1 \in [0, \alpha_2^*]} \mathcal{U}(\alpha_1, \alpha_2^*). \tag{A.5}$$

For $\alpha_1 < \alpha_2^*$, use (19) and the definition of the indifferent fund $\tilde{\theta}$ to rewrite

$$\begin{aligned}
& \mathcal{U}(\alpha_1, \alpha_2^*) \\
= \; & E[\lambda] \int_0^1 \max\{\theta V_s(\alpha_1) + (1-\theta)V_r(\alpha_1), \theta V_s(\alpha_2^*) + (1-\theta)V_r(\alpha_2^*)\}\hat{g}(\theta)d\theta \\
= \; & E[\lambda] \left\{ \int_0^{\tilde{\theta}} \left[\theta V_s(\alpha_1) + (1-\theta)V_r(\alpha_1)\right] \hat{g}(\theta)d\theta + \int_{\tilde{\theta}}^1 \left[\theta V_s(\alpha_2^*) + (1-\theta)V_r(\alpha_2^*)\right] \hat{g}(\theta)d\theta \right\}.
\end{aligned}$$

Taking into account that $E[\lambda] > 0$ is constant and $\tilde{\theta}$ is a function of α_1, the partial derivative with respect to α_1 has the same sign as

$$\begin{aligned}
\frac{\partial \mathcal{U}(\alpha_1, \alpha_2^*)}{\partial \alpha_1} \quad \propto \quad & \frac{\partial \tilde{\theta}}{\partial \alpha_1} \hat{g}(\tilde{\theta}) \left[\tilde{\theta} V_s(\alpha_1) + (1-\tilde{\theta})V_r(\alpha_1) - \tilde{\theta} V_s(\alpha_2^*) - (1-\tilde{\theta})V_r(\alpha_2^*)\right] \\
& + \int_0^{\tilde{\theta}} \left[\theta V_s'(\alpha_1) + (1-\theta)V_r'(\alpha_1)\right] \hat{g}(\theta)d\theta.
\end{aligned}$$

The term in brackets in the first line is zero because, by definition, a fund with preference $\tilde{\theta}$ is indifferent between advice policies α_1 and α_2^* — see (18). Therefore, the first-order condition simplifies to making the term in the second line (the integral) equal to zero. This condition is equivalent to the first-order condition that maximizes the expressive voting utility of a fund with a preference equal to the average θ in this segment. This is analogous to the monopolist in (13), with the difference that we are changing the support of the preference distribution.

We can also see this by substituting (A.1) and (A.3)

$$\begin{aligned}
\frac{\partial \mathcal{U}(\alpha_1, \alpha_2^*)}{\partial \alpha_1} \quad \propto \quad & \int_0^{\tilde{\theta}} \left[\theta(1-\alpha_1)K(\alpha_1) - (1-\theta)\alpha_1 K(\alpha_1)\right] \hat{g}(\theta)d\theta \\
= \quad & K(\alpha_1) \int_0^{\tilde{\theta}} \left[\theta - \alpha_1\right] \hat{g}(\theta)d\theta \\
= \quad & K(\alpha_1) Pr[\theta \leq \tilde{\theta}] \left[\hat{E}[\theta | \theta \leq \tilde{\theta}] - \alpha_1\right].
\end{aligned}$$

This derivative is similar to (A.4), and implies that the first-order condition must satisfy $\alpha_1^* = \hat{E}[\theta | \theta \leq \tilde{\theta}]$, which implies that (22) must hold.

Fixed α_1^*, the same logic applied to α_2 implies that (23) must hold. Finally, condition (24) defines the preference $\tilde{\theta}$ of the funds which are indifferent between the two policies, concluding the proof.

Part (i) We start by writing the profit function of Firm 1, given policies (α_1, α_2). Funds with preference (λ, θ) such that $U_{\lambda\theta}(\alpha_1) > U_{\lambda\theta}(\alpha_2)$ will purchase from Firm 1 and pay the

price $U_{\lambda\theta}(\alpha_1) - U_{\lambda\theta}(\alpha_2)$. Funds with preference such that $U_{\lambda\theta}(\alpha_1) < U_{\lambda\theta}(\alpha_2)$ will purchase from Firm 2 and pay the price $U_{\lambda\theta}(\alpha_2) - U_{\lambda\theta}(\alpha_1)$. Bertrand competition implies that the firms cannot profit from the funds that are indifferent between the two advice policies.[17] Therefore, the profit of Firm 1 is

$$\begin{aligned}
\text{Profit}_1(\alpha_1, \alpha_2) &= \int_0^1 \int_0^{\overline{\lambda}} \max\left\{U_{\lambda\theta}(\alpha_1) - U_{\lambda\theta}(\alpha_2), 0\right\} g(\lambda, \theta) d\lambda d\theta \\
&= \int_0^1 \int_0^{\overline{\lambda}} \max\left\{U_{\lambda\theta}(\alpha_1), U_{\lambda\theta}(\alpha_2)\right\} g(\lambda, \theta) d\lambda d\theta - \int_0^1 \int_0^{\overline{\lambda}} U_{\lambda\theta}(\alpha_2) g(\lambda, \theta) d\lambda d\theta \\
&= \mathcal{U}(\alpha_1, \alpha_2) - \int_0^1 \int_0^{\overline{\lambda}} U_{\lambda\theta}(\alpha_2) g(\lambda, \theta) d\lambda d\theta.
\end{aligned}$$

Therefore, to maximize its own profit given α_2, the best response of Firm 1 is to choose the policy α_1 that maximizes the total expressive voting $\mathcal{U}(\alpha_1, \alpha_2)$. This result is analogous to the main result in Lederer and Hurter (1986). Similarly, the profit of Firm 2 is

$$\text{Profit}_2(\alpha_1, \alpha_2) = \mathcal{U}(\alpha_1, \alpha_2) - \int_0^1 \int_0^{\overline{\lambda}} U_{\lambda\theta}(\alpha_1) g(\lambda, \theta) d\lambda d\theta.$$

To maximize its own profit given α_1, the best response of Firm 2 is to choose the policy α_2 that maximizes the total expressive voting $\mathcal{U}(\alpha_1, \alpha_2)$.

Consequently, if policies (α_1^*, α_2^*) maximize $\mathcal{U}$, then inequalities (20) and (21) most hold; therefore, our results imply that (α_1^*, α_2^*) form a competitive equilibrium. From our previous claim, we know that there exists policies that maximize $\mathcal{U}$, therefore a competitive equilibrium exists, concluding this part of the proof.

Part (ii) Suppose policies $\alpha_1^* < \alpha_2^*$ form a competitive equilibrium. Equilibrium implies that (20) holds, which in turn also implies that (A.5) holds since $\alpha_1^* < \alpha_2^*$. We have shown that (A.5) implies $\alpha_1^* = \hat{E}[\theta | \theta \leq \tilde{\theta}]$. Similar logic applies to Firm 2: inequality (21) together with $\alpha_2^* > \alpha_1^*$ imply that $\alpha_2^* = \hat{E}[\theta | \theta \geq \tilde{\theta}]$. These two optimality conditions together with the definition of the indifferent fund $\tilde{\theta}$ imply that (22), (23), and (24) hold. Moreover, because $G(\lambda, \theta)$ has full support, we have that $\hat{E}[\theta | \theta \leq \tilde{\theta}] < \hat{E}[\theta] < \hat{E}[\theta | \theta \geq \tilde{\theta}]$ for any $\tilde{\theta} \in (0, 1)$. Therefore, $\alpha_1^* < \alpha^* < \alpha_2^*$, concluding this part of the proof.

Uniqueness: We now define sufficient conditions for equilibrium uniqueness. Combine conditions (22), (23), and (24) into a single equation. Rewrite condition (24) and then

[17] There is a measure zero of indifferent funds for any $\alpha_1 \neq \alpha_2$ because of our assumption that the preference distribution has no atoms.

substitute $\alpha_1 = \hat{E}[\theta \leq \tilde{\theta}]$ and $\alpha_2 = \hat{E}[\theta \geq \tilde{\theta}]$:

$$\tilde{\theta}\left[V_s(\hat{E}[\theta \leq \tilde{\theta}]) - V_s(\hat{E}[\theta \geq \tilde{\theta}])\right] + (1-\tilde{\theta})\left[V_r(\hat{E}[\theta \leq \tilde{\theta}]) - V_r(\hat{E}[\theta \geq \tilde{\theta}])\right] \quad = \quad 0. \quad \text{(A.6)}$$

Define the LHS as the function $\Gamma(\tilde{\theta})$. A necessary condition for a competitive equilibrium is that the cutoff fund $\tilde{\theta}$ separating the two markets is such that $\Gamma(\tilde{\theta}) = 0$. Note that

$$\begin{aligned}
\Gamma(1) &= V_s(\hat{E}[\theta]) - V_s(1) < 0, \\
\Gamma(0) &= V_r(0) - V_r(\hat{E}[\theta]) > 0,
\end{aligned}$$

where the strict inequalities hold because V_s is strictly increasing and V_r is strictly decreasing — see Lemma 1. Since Γ is continuous, it crosses zero at least once. If it only crosses zero once, then the equilibrium must be unique. Moreover, if the equilibrium is unique, then it must maximize $\mathcal{U}$, since there always exist policies that maximize total expressive voting utility and form a competitive equilibrium. For example, if s and r are independently drawn from a standard normal distribution and $\hat{g}(\theta)$ is a uniform distribution, then Γ is strictly decreasing and the equilibrium is unique. □

Proof of Proposition 3: The proof follows very closely the arguments of the proof of Proposition 2, therefore we do not repeat some previous details.

Part (i) Write the profit function of Firm k, given policies (α_k, α_{-k}). Let $\max\{U_{\lambda\theta}(\alpha_{-k})\}$ denote the maximum $U_{\lambda\theta}(\cdot)$ from the set of policies offered by the firms excluding Firm k. Funds with preference (λ, θ) such that $U_{\lambda\theta}(\alpha_k) > \max\{U_{\lambda\theta}(\alpha_{-k})\}$ will purchase from Firm k and pay the price $U_{\lambda\theta}(\alpha_k) - \max\{U_{\lambda\theta}(\alpha_{-k})\}$. Therefore, the profit of Firm k is

$$\begin{aligned}
\text{Profit}_k(\alpha_k|\alpha_{-k}) &= \int_0^1 \int_0^{\overline{\lambda}} \max\left\{U_{\lambda\theta}(\alpha_k) - \max\{U_{\lambda\theta}(\alpha_{-k})\}, 0\right\} g(\lambda,\theta) d\lambda d\theta \\
&= \int_0^1 \int_0^{\overline{\lambda}} \max\left\{U_{\lambda\theta}(\alpha_k), \max\{U_{\lambda\theta}(\alpha_{-k})\}\right\} g(\lambda,\theta) d\lambda d\theta \\
&\quad - \int_0^1 \int_0^{\overline{\lambda}} \max\{U_{\lambda\theta}(\alpha_{-k})\} g(\lambda,\theta) d\lambda d\theta \\
&= \mathcal{U}(\alpha_k|\alpha_{-k}) - \int_0^1 \int_0^{\overline{\lambda}} \max\{U_{\lambda\theta}(\alpha_{-k})\} g(\lambda,\theta) d\lambda d\theta.
\end{aligned}$$

Therefore, to maximize its own profit given α_{-k}, the best response of Firm k is to choose the policy α_k that maximizes the total expressive voting $\mathcal{U}(\alpha_k|\alpha_{-k})$.

Consequently, if a policy vector maximizes total expressive voting utility, then it constitutes a competitive equilibrium. Because the policy vector belongs to a compact set and $\mathcal{U}$ is continuous, there exists a vector that maximizes $\mathcal{U}$, concluding this part of the proof.

Part (ii) The fact that the first-order condition (28) and indifference condition (27) are necessary conditions follows from the same arguments as Proposition 2, therefore we omit the details.

Part (iii) The result follows for two reasons. First, if the median fund θ_m prefers to purchase the advice α_m^*, then all funds will a lower preference ($\theta \leq \theta_m$) will purchase from firms with a weakly lower policy $\alpha_k^* \leq \alpha_m^*$. Conversely, all funds will a higher preference ($\theta \geq \theta_m$) will purchase from firms with a weakly higher policy $\alpha_k^* \geq \alpha_m^*$. Second, whenever fund α_m^* votes to approve or reject a proposal, either all firms with a lower policy $\alpha_k^* \leq \alpha_m^*$ or all firms with a higher policy $\alpha_k^* \geq \alpha_m^*$ recommend the same vote (by the same logic of Footnote 8). Therefore, a majority of funds will be following the same recommendation as α_m^*. From α_m^* we can compute the equilibrium voting outcome q_m^* and $\Psi(q_m^*)$. □

References

[1] Alonso, Ricardo and Odilon Câmara, "Persuading Voters," *American Economic Review*, 2016, 106(11), 3590-3605.

[2] American Accountability Foundation, "Proxy Wars: Glass Lewis," 2023, available at: https://americanaccountabilityfoundation.com/wp-content/uploads/2023/07/Proxy-Wars-Glass-Lewis.pdf

[3] Bergemann, Dirk, and Alessandro Bonatti. "Markets for information: An introduction." *Annual Review of Economics*, 2019, 11(1), 85-107.

[4] Bernhardt, D., Krasa, S., Polborn, M., "Political polarization and the electoral effects of media bias." *Journal of Public Economics*, 2008, 92, 1092-1104.

[5] Bolton, Patrick, Tao Li, Enrichetta Ravina, and Howard Rosenthal, "Investor Ideology," *Journal of Financial Economics*, August 2020, Vol. 137(2), 320-352.

[6] Broadridge + PwC, "2019 Proxy Season Review," ProxyPulse, 2019, available at: https://www.broadridge.com/_assets/pdf/broadridge-proxypulse-2019-review.pdf.

[7] Buechel, B., L. Mechtenberg, and A. F. Wagner, "When Two Experts Are Better Than One: The Example of Shareholder Voting," *ECGI Finance Working Paper* No. 832/2022, 2024.

[8] Business Roundtable, "The Need for Bold Proxy Process Reforms," April 2025, available at https://www.businessroundtable.org/business-roundtable-calls-for-bold-reform-of-the-corporate-proxy-process-in-new-white-paper

[9] Chan, Jimmy and Wing Suen, "A Spatial Theory of News Consumption and Electoral Competition," *The Review of Economic Studies*, 2008, 75, 699-728,

[10] Chan, Jimmy and Wing Suen, "Media as watchdogs: The role of news media in electoral competition," *European Economic Review*, 2009, 53, 799-814.

[11] Duggan, J. and C. Martinelli, "A Spatial Theory of Media Slant and Voter Choice," *Review of Economic Studies*, 2011 (78), 640-666.

[12] Fiorina, Morris P., "The Voting Decision: Instrumental and Expressive Aspects," *Journal of Politics*, May 1976, Vol. 38(2), 390-415.

[13] Garicano, Luis, "Hierarchies and the Organization of Knowledge in Production," *Journal of Political Economy*, 2000, 108(5), 874–904.

[14] Gentzkow, Matthew and Jesse M. Shapiro, "Media Bias and Reputation", *Journal of Political Economy*, 2006, 114(2), pp. 280-316.

[15] Hayne, Christie and Marshall Vance, "Information Intermediary or De Facto Standard Setter? Field Evidence on the Indirect and Direct Influence of Proxy Advisors," *Journal of Accounting Research*, September 2019, Vol. 57(4), 969-1011.

[16] Hirshleifer, Jack, "The private and social value of information and the reward to inventive activity," *American Economic Review*, 1971, 61(4), 561–574.

[17] Hu, Edwin, Nadya Malenko, and Jonathon Zytnick, "Custom Proxy Voting Advice," *NBER Working Paper* 32559, 2024.

[18] Kerber, Ross, "Proxy Adviser ISS Expands Offerings for 'ESG Skeptic' Clients." *Reuters*, March 4, 2024, available at https://www.reuters.com/sustainability/boards-policy-regulation/proxy-adviser-iss-expands-offerings-esg-skeptic-clients-2024-03-04/.

[19] Lederer, Phillip J., and Arthur P. Hurter Jr., "Competition of Firms: Discriminatory Pricing and Location," *Econometrica*, 1986, 54(3), 623-640.

[20] Li, Tao, "Outsourcing Corporate Governance: Conflicts of Interest Within the Proxy Advisory Industry," *Management Science*, 2018, 64 (6), pp. 2951-2971.

[21] Lupia, Arthur, "Shortcuts versus encyclopedias: Information and voting behavior in California insurance reform elections," *The American Political Science Review*, 1994, 88(1): 63-76.

[22] Ma, Shichao and Yan Xiong, "Information Bias in the Proxy Advisory Market," *Review of Corporate Finance Studies*, March 2021, Vol. 10(1), 82-135.

[23] Malenko, Andrey and Nadya Malenko, "Proxy Advisory Firms: The Economics of Selling Information to Voters," *Journal of Finance*, 2019, Vol. 74(5), 2441-2490.

[24] Malenko, Andrey, Nadya Malenko, and Chester S. Spatt, "Creating Controversy in Proxy Voting Advice," *Journal of Finance*, 2025.

[25] Matsusaka, John and Chong Shu, "A Theory of Proxy Advice when Investors Have Social Goals," *Working Paper*, 2021.

[26] Matsusaka, John and Chong Shu, "Robo-Voting: Does Delegated Proxy Voting Pose a Challenge for Shareholder Democracy?," *Seattle University Law Review*, 2024, Vol. 47(2), 605-634.

[27] May, Kenneth O., "A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decisions," *Econometrica*, 1952, Vol. 20(4), 680-684.

[28] McCahery, Joseph A., Zacharias Sautner, and Laura T. Starks. "Behind the scenes: The corporate governance preferences of institutional investors". *The Journal of Finance*, 2016, 71: 2905–32.

[29] Mullainathan, Sendhil, and Andrei Shleifer. "The market for news." *American Economic Review*, 2005, 95(4), 1031-1053.

[30] Perego, Jacopo and Sevgi Yuksel, "Media Competition and Social Disagreement," *Econometrica*, January 2022, Vol. 90(1), 223-265.

[31] Prat A, Strömberg D., "The Political Economy of Mass Media." In: Acemoglu D, Arellano M, Dekel E, eds. Advances in Economics and Econometrics: Tenth World Congress. Econometric Society Monographs. Cambridge University Press; 2013:135-187.

[32] Riedl, Arno and Paul Smeets, "Why Do Investors Hold Socially Responsible Mutual Funds?," *Journal of Finance*, December 2017, Vol. 72(6), 2505-2550.

[33] Shu, Chong, "The Proxy Advisory Industry: Influencing and Being Influenced," *Journal*

of Financial Economics, 2024, 154, 103810.

[34] Smolin, Alex, "Disclosure and pricing of attributes," *RAND Journal of Economics*, 2023, 54(4), pp. 570–597.

[35] Strömberg, David, "Mass Media Competition, Political Competition, and Public Policy," *Review of Economic Studies*, 2004, 71, 265-284.

[36] U. S. House Committee on the Judiciary, *Letter to Mr. Kevin Cameron, Glass, Lewis & Co.*, August 1, 2023, available at: https://judiciary.house.gov/sites/evo-subsites/republicans-judiciary.house.gov/files/evo-media-document/letters.pdf.

[37] U. S. House Committee on the Judiciary, *Letter to Mr. Gary Retelny, Institutional Shareholder Services*, March 31, 2025, available through Wise (2025).

[38] U. S. House of Representatives, Examining the Market Power and Impact of Proxy Advisory Firms: Hearing before the Subcommittee on Capital Markets and Government Sponsored Enterprises of the Committee on Financial Services, Serial No. 113-27, June 5, 2013.

[39] U. S. Securities and Exchange Commission, "Roundtable on the Proxy Process," November 2018, available at https://www.sec.gov/files/proxy-round-table-transcript-111518.pdf

[40] U. S. Securities and Exchange Commission, "SEC Adopts Rules to Enhance Proxy Voting Disclosure by Registered Investment Funds and Require Disclosure of 'Say-on-Pay' Votes for Institutional Investment Managers," SEC Press Release 2022-198, November 2, 2022, available at: https://www.sec.gov/newsroom/press-releases/2022-198.

[41] Vanguard, Investment Stewardship: 2023 Annual Report, available at: https://corporate.vanguard.com/content/corporatesite/us/en/corp/how-we-advocate/investment-stewardship/reports-and-policies.html.

[42] White, Alexandra, "Glass Lewis to end benchmark voting recommendations on proxy issues," *Financial Times*, October 14, 2025, available at https://www.ft.com/content/7860f94c-7b1c-49c9-9f4c-a561ee39ee31.

[43] Wise, Justin, "House Committee Launches Antitrust Probe of Proxy Advisers," *Bloomberg Law*, March 31, 2025, available at: https://news.bloomberglaw.com/antitrust/house-committee-launches-antitrust-probe-of-proxy-advisers.

Appendix B. Online Appendix

This Online Appendix presents some extensions and generalizes the main equilibrium characterization of *Shareholder Democracy and the Market for Voting Advice*, by Odilon Câmara, John Matsusaka, and Chong Shu, CMS henceforth.

B.1. Extensions

B.1.1. Buying Advice from More Than One Advisor

CMS have assumed that a fund purchases only a single advice policy. Here we allow funds to purchase multiple advice policies, and show that their demand for additional policies is limited. Suppose that firms offer policies $\{\alpha_1, \ldots, \alpha_N\}$, and consider a fund with preference θ. What is the value of obtaining a second advice policy?

Remark 4. *If fund θ receives recommendations based on advice policy $\alpha \geq \theta$, then the marginal value of acquiring recommendations based on advice policy $\alpha' > \alpha$ is zero. If fund θ receives recommendations from advice policy $\alpha \leq \theta$, then the marginal value of advice policy $\alpha' < \alpha$ is zero.*

To see the intuition, suppose that a fund receives recommendations from some policy $\alpha > \theta$, which assigns too much weight to the social dimension from the fund's perspective. If the fund receives additional recommendations from a more extreme policy $\alpha' > \alpha$, the new recommendations do not change its voting behavior. Therefore, the new recommendations provide no valuable of information.

The remark implies that if fund θ receives advice from its ideal policy $\alpha = \theta$, then the additional value of any other advice policy is zero. Less obviously, it also implies that if a fund θ purchases more than one policy, it purchases exactly two of them, the one immediately to the left and the one immediately to the right of θ. A full solution of this extension of the model requires information on the shape of the distributions F and G to compute equilibrium choices and election outcomes, which we do not pursue here.

The substantive takeaway is that in a world in which multiple firms offer multiple policies, the utility of each fund (and hence its willingness to pay) is maximized by acquiring the closest policy to the left of its ideal point, and the closest policy to the right. Therefore,

as in our benchmark model, firms still want to offer policies targeting the preferences of the funds purchasing their advice, taking into account how much they care about voting.

B.1.2. Instrumental Voting

Instead of considering expressive voters, we could have adapted the assumptions of Perego and Yuksel (2022) to introduce instrumental voters in a tractable manner. Suppose there is a finite number of funds, all purely instrumental voters, and policies are implemented with a probability equal to the percentage of total shares voting to approve the proposal. A fund of size τ_i will increase (decrease) the probability of approval by τ_i if it votes yes (no). Voters are pivotal in all elections (their vote always change the probability of approval) but they do not learn about the state from the event of being pivotal.

We can reinterpret the payoff functions to match our benchmark model. Suppose all funds equally care about the financial return, but they differ on how much they care about the social return. Fund i receives a policy payoff $\beta_i s_j + r_j$ if proposal j is approved by shareholders, and $-(\beta_i s_j + r_j)$ if it is rejected, where $\beta_i > 0$ is how much the fund cares about the social return. Voting $a_i \in \{-1, +1\}$ causes an expected payoff change of

$$a_i \tau_i (\beta_i s_j + r_j),$$

because the vote changes the approval probability by $a_i \tau_i$. Define $\lambda_i = \tau_i(1 + \beta_i)$ and $\theta_i = \beta_i/(1 + \beta_i)$, and rewrite the payoff change as

$$a_i \lambda_i (\theta_i s_j + (1 - \theta_i) r_j),$$

which is the equation that we are using for our expressive voters.

We have chosen to focus on the expressive voting interpretation because we think this is particularly relevant in this scenario where voting records are public and funds may care differently about these records. Moreover, the welfare analysis needs to take into account the additional randomness of probabilistic voting in election outcomes.

B.1.3. Paying a Premium for Close Elections

In our model, funds vote for expressive reasons, meaning that they derive utility from voting their interests regardless of whether their vote is pivotal. In contrast, in the instrumental

or strategic voting model, voters derive no utility from non-pivotal votes, and therefore are indifferent about how they vote in non-pivotal states. An "intermediate" approach is to assume that voters care about their votes in all states, but care more when their vote is pivotal. To study this case, let $A(s,r) \in [0,1]$ be the share of votes cast in favor of a proposal (s,r) by all shareholders. Define the instrumental premium as a differentiable function $\Lambda(A)$ with the following characteristics: $\Lambda(0) = \Lambda(1) = 0$; it is symmetric around $A = 0.5$; and it strictly increases from $A = 0$ to $A = 0.5$. The investor's total voting utility is then

$$a_{ij}\lambda_i v_{ij} + \Lambda(A(s,r))a_{ij}v_{ij}.$$

The term $a_{ij}\lambda_i v_{ij}$ represents expressive utility as in the benchmark model, while $\Lambda(A(s,r))a_{ij}v_{ij}$ captures the instrumental premium: the fund is willing to pay more to correctly register its preferences in close elections.

In the case of a monopolist, this premium does not affect our results. Since the monopolist provides the same advice to all funds, the electoral outcome is always $A(s,r) = 0$ or $A(s,r) = 1$. Hence, only the expressive voting utility matters and the monopolist targets the average voter.

In the duopoly case, whereas previously the value of purchasing a policy α did not depend on the other policies sold in the market, now the value of an advice policy sold by one firm depends on the policy of the other firm and the purchase decisions of other voters. The instrumental premium does not change voter's preference order over polices. The duopoly equilibrium is then characterized by a pair of policies and a cutoff voter, $\alpha_1 < \tilde{\theta} < \alpha_2$, where all voters $\theta < \tilde{\theta}$ purchase from firm 1, voters $\theta > \tilde{\theta}$ purchase from Firm 2, and voter $\tilde{\theta}$ is indifferent between the two advice policies.

The instrumental premium is zero if both advisors make the same recommendation; it is positive only if the advisors issue opposing recommendations. When recommendations disagree, the instrumental premium is the highest when the median voter θ_m is the cutoff voter $\tilde{\theta}$ because, in this case, votes are equally split — half of the votes are for and half are against the proposal. If the cutoff voter $\tilde{\theta}$ is far from the median voter, then one firm controls more votes than the other and split elections are far from 50/50. Such lopsided elections decrease the instrumental premium. Therefore, the instrumental premium creates an incentive for firms to move $\tilde{\theta}$ in the direction of the median voter.

Perhaps surprisingly, this incentive to move toward the median could be detrimental to the median voter and therefore reduce informed representation. Suppose that in the benchmark model $\alpha_1 < \tilde{\theta} < \theta_m < \alpha_2$, so that the median voter and a majority of voters purchase advice α_2. Since α_2 serves a majority of the market, all voting outcomes follow its recommendation. Now suppose there is a small instrumental premium. Both firms now have an incentive to marginally move their policies α_1 and α_2 to the right, which moves $\tilde{\theta}$ to the right, closer to the median voter, making elections more competitive. The median voter and a majority of voters would continue to purchase advice from Firm 2, which decides the elections, but the new policy α_2 is further away from the median, reducing informed representation.

Malenko et al. (2025) shows how a monopolist facing instrumental voters can increase its profit by providing a free signal that leads investors to believe that elections will be close. Our results uncover similar incentives in competitive markets: firms profit from generating close elections, and they do so by slanting their advice policy and issuing opposing recommendations.

B.1.4. Fully Informative Reports

Our running assumption is that there are limits on the information that advisors can produce and investors can decode. To reflect this limit, we assumed that advice takes the form of binary recommendations. Here we consider an extension in which advisory firms can obtain and transmit complete information about r and s for a subset of proposals.

We modify the model from Section VII.B to assume that each advisory firm k is characterized by an exogenous parameter $\gamma_k \in [0, 1)$ that represents its "capacity" to generate fully informative reports. A firm with capacity γ_k learns the true returns (r, s) of a measure γ_k of proposals. As in Section VII.B, if a firm with capacity γ_k learns the returns of a proposal j, then all firms with a higher capacity $\gamma_{k'} \geq \gamma_k$ also learn these returns for proposal j. We assume that funds can interpret the full reports. For the remaining $1 - \gamma_k$ proposals, the firm generates a binary recommendation according to the chosen policy α_k, as before.

Abusing notation, the expected expressive voting utility (6) from purchasing advice from

firm k becomes

$$U_{\lambda\theta}(\alpha_k, \gamma_k) = \gamma_k \lambda \Big[\theta V_s(\theta) + (1-\theta) V_r(\theta)\Big] + (1-\gamma_k)\lambda \Big[\theta V_s(\alpha_k) + (1-\theta) V_r(\alpha_k)\Big] \quad \text{(B.1)}$$

It is as if the fully informative reports provide perfect recommendations ($\alpha = \theta$) for a measure γ_k of proposals, while the remaining $1 - \gamma_k$ recommendations follow the slanted policy α_k. Advisory firms are now both horizontally differentiated (the location of their policy choice α_k) and vertically differentiated (their full information capacity γ_k).

Our main qualitative results continue to hold. In equilibrium, each fund purchases from the firm providing the highest expected utility, and the price paid is the difference between this highest utility and the second best option available. Firms segment the market: each firm chooses a policy targeting the (weighted) average preference of its customers.

In the case of a monopolist, the firm's profit is a strictly increasing function of γ. A higher capacity allows the firm to provide more valuable information to funds and charge higher prices. In the case of competition, increasing the capacity of only one firm increases the profit of this firm, but increasing the capacity of all firms at the same time can decrease profits. To see this, suppose that all firms have the same capacity γ. As we discussed in Section V.B, in a competitive market, a firm can only earn a profit from issuing recommendations that are different than the second best option available to an investor. Since all firms offer full information on a measure γ of proposals (and they offer full information on an identical set of proposals, by assumption), they earn no profit from offering these reports (there is no product differentiation here). Their profit comes entirely from selling the heterogeneous binary advice policies for the remaining $1 - \gamma$ proposals. Hence, increasing γ for all firms at the same time decreases profits but enhances shareholder democracy.

B.2. Generalizations

We start by introducing some useful notation. For any pair of vectors $a, b \in \mathbb{R}^T$, where $T \geq 2$ is finite, let $\langle a, b \rangle \equiv \sum_{t=1}^{T} a_t b_t$ be the inner product. Define the probability simplex over T elements

$$\Delta(T) \equiv \left\{ \delta = (\delta_1, \ldots, \delta_T) \in [0,1]^T \,\middle|\, \sum_{t=1}^{T} \delta_t = 1,\ \delta_t \geq 0 \text{ for all } t \right\}.$$

We next describe which assumptions we are generalizing.

1. **Proposals:** CMS assumes that each proposal j has two return dimensions: $s_j, r_j \in \mathbb{R}$. Now suppose that there is a finite number $T \geq 2$ of return dimensions, where t is a generic dimension. Abusing notation, let $r_{jt} \in \mathbb{R}$ be the return from dimension t, with the vector $\boldsymbol{r_j} = (r_{j1}, \ldots, r_{jT}) \in \mathbb{R}^T$ representing the returns of some proposal j. Following CMS, we continue to assume that the return distribution is given by cdf $F(\boldsymbol{r})$ and pdf $f(\boldsymbol{r})$, with no atoms, full support on some subset $\mathcal{R} \subset \mathbb{R}^T$, and the zero vector $\mathbf{0}$ belongs to the interior of $\mathcal{R}$. To simplify presentation, CMS assumes $E[s] = E[r] = 0$, $E[s|r \geq 0] \geq 0$, and $E[r|s \geq 0] \geq 0$. We do not impose these assumptions here. Define $\hat{\boldsymbol{r}} \equiv E[\boldsymbol{r}]$ as the expected return vector given the prior belief F. The main consequence of this generalization will be that some funds might prefer to not purchase certain types of recommendations.

2. **Funds:** CMS assumes that a fund i attaches weight θ_i to the social return dimension and weight $(1-\theta_i)$ to the financial return. With T return dimensions, we abuse notation and generalize preferences to be represented by the vector $\boldsymbol{\theta_i} = (\theta_{i1}, \ldots, \theta_{iT}) \in \Delta(T)$. We continue to assume that funds also differ on preference parameter $\lambda_i \in (0, 1]$.

   We can then rewrite (1) and (3): the value of proposal $\boldsymbol{r}$ to fund $\boldsymbol{\theta}$ becomes

   $$v_{ij} \quad = \quad \langle \boldsymbol{\theta}, \boldsymbol{r} \rangle, \tag{B.2}$$

   and the expressive utility becomes

   $$u(\lambda, \boldsymbol{\theta}, \boldsymbol{r}, a) \quad = \quad a\lambda \langle \boldsymbol{\theta}, \boldsymbol{r} \rangle. \tag{B.3}$$

   As before, the fund would like to approve the proposal ($a = 1$) if the weighted return of the proposal is positive, and reject ($a = -1$) if negative. Because different funds assign different weights to the different dimensions, funds might disagree on whether a proposal should be approved. We continue to assume that preference parameters are distribute according to some cdf $G(\lambda, \boldsymbol{\theta})$ and continuous pdf $g(\boldsymbol{\theta})$.

3. **Firms:** We continue to assume that there is a finite number N of firms and each firm can only offer one advice policy. As before, our main assumption is that each fund can only offer a binary advice, which is naturally interpreted as an approval or rejection recommendation. To simplify exposition, CMS constrained the firm to

choose an advice policy from a particular subset of experiments indexed by $\alpha \in [0,1]$. We now drop this constraint and allow the firm to choose *any* binary experiment. An experiment has a message space $m \in \{m^+, m^-\}$, where m^+ is a recommendation to approve the proposal and m^- a recommendation to reject. The experiment is represented by a measurable function $M : \mathcal{R} \to [0,1]$, where $M(\boldsymbol{r})$ represents the probability of sending an approval message m^+ conditional on the proposal having return $\boldsymbol{r}$. Let $\boldsymbol{r}^+ \equiv E[\boldsymbol{r}|m = m^+]$ and $\boldsymbol{r}^- \equiv E[\boldsymbol{r}|m = m^-]$ be the posterior beliefs after observing the realized recommendation from policy M.

4. **Market:** The market operates as in CMS. Firms simultaneously choose their policies, $\{M_1, \ldots, M_N\}$. Firms observe the chosen policies and make discriminatory price offers to the different funds. Funds make their purchase decisions (a fund can buy at most one advice, or choose to not buy an advice). Firms issue voting recommendations to their clients, elections are held, and payoffs are realized.

B.2.1. Value of Information

The value of advice M for fund $(\lambda, \boldsymbol{\theta})$ is the expected value of voting according to the recommendations generated by M. With probability $\Pr(m^+)$ the firm recommends approval, which results in a posterior belief $\boldsymbol{r}^+$ about the expected return of the proposal. Voting approval ($a = 1$) results in an expected utility $\lambda\langle\boldsymbol{\theta}, \boldsymbol{r}^+\rangle$. With probability $\Pr(m^-)$ the firm recommends rejection, which results in a posterior belief $\boldsymbol{r}^-$ about the expected return of the proposal. Voting rejection ($a = -1$) results in an expected utility $-\lambda\langle\boldsymbol{\theta}, \boldsymbol{r}^-\rangle$. We can rewrite (6) as

$$U_{\lambda\boldsymbol{\theta}}(M) \;=\; \lambda\left[\Pr(m^+)\langle\boldsymbol{\theta}, \boldsymbol{r}^+\rangle - \Pr(m^-)\langle\boldsymbol{\theta}, \boldsymbol{r}^-\rangle\right]. \tag{B.4}$$

We next compute the value of not acquiring an advice. Given the prior belief $\hat{\boldsymbol{r}}$, an uninformed fund that does not purchase an advice will always vote approve if $\langle\boldsymbol{\theta}, \hat{\boldsymbol{r}}\rangle > 0$, always vote reject if $\langle\boldsymbol{\theta}, \hat{\boldsymbol{r}}\rangle < 0$, and will be indifferent if $\langle\boldsymbol{\theta}, \hat{\boldsymbol{r}}\rangle = 0$. Therefore, in all three cases, the value of not acquiring an advice is given by the absolute value $\lambda|\langle\boldsymbol{\theta}, \hat{\boldsymbol{r}}\rangle|$. We will use M_0 to represent the option of not acquiring information, which yields expected utility

$$U_{\lambda\boldsymbol{\theta}}(M_0) = \lambda|\langle\boldsymbol{\theta}, \hat{\boldsymbol{r}}\rangle|. \tag{B.5}$$

B.2.2. Price Competition

We use backward induction and first consider the subgame following a set of policy choices $\mathcal{M} = \{M_0, M_1, \ldots, M_N\}$, with M_0 the option of not acquiring information. Consider firms competing for fund $(\lambda, \boldsymbol{\theta})$. In equilibrium, the fund will always select the choice $M \in \mathcal{M}$ that yields the highest expected payoff. If there is a tie for first place, then the fund will equally randomize between the best options.

Consider a firm k choosing a price to offer to fund $(\lambda, \boldsymbol{\theta})$. Define

$$\bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k}) = \max\{U_{\lambda\boldsymbol{\theta}}(M_{k'})\}_{k' \neq k, k' \in \{0,1,\ldots,N\}}$$

as the highest payoff that can be obtained by fund $(\lambda, \boldsymbol{\theta})$ from the policies in the set $\mathcal{M}_{-k}$. If $U_{\lambda\boldsymbol{\theta}}(M_k) > \bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k})$, then firm k will offer a price equal to the difference between the utility that it offers and the second best option available, $U_{\lambda\boldsymbol{\theta}}(M_k) - \bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k})$. Otherwise, the firm offers a price equal to the marginal cost (which is zero).

B.2.3. Policy Choices

Given the effects of price competition, we can now describe the best response of each firm during the policy choice stage. Consider the choice of firm k. Fix the choices of the other firms and the outside option, $\mathcal{M}_{-k}$. The profit of firm k is

$$\text{Profit}_k(M_k|\mathcal{M}_{-k}) \equiv \int_0^1 \int_{\boldsymbol{\theta}\in\Delta(T)} \max\left\{U_{\lambda\boldsymbol{\theta}}(M_k) - \bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k}), 0\right\} g(\lambda, \boldsymbol{\theta}) d\boldsymbol{\theta} d\lambda. \quad \text{(B.6)}$$

Rewriting (25), the total expressive utility in the market is

$$\mathcal{U}(M_k|\mathcal{M}_{-k}) \equiv \int_0^1 \int_{\boldsymbol{\theta}\in\Delta(T)} \max\left\{U_{\lambda\boldsymbol{\theta}}(M_0), U_{\lambda\boldsymbol{\theta}}(M_1), \ldots, U_{\lambda\boldsymbol{\theta}}(M_N)\right\} g(\lambda, \boldsymbol{\theta}) d\boldsymbol{\theta} d\lambda. \quad \text{(B.7)}$$

Lemma B. 1. *Given $\mathcal{M}_{-k}$, experiment M_k maximizes firm's k profit if and only if M_k maximizes the total expressive utility. Therefore, policies $\mathcal{M}^* = \{M_1^*, \ldots, M_N^*\}$ form a competitive equilibrium if and only if*

$$\mathcal{U}(M_k^*|\mathcal{M}_{-k}^*) \geq \mathcal{U}(M_k|\mathcal{M}_{-k}^*) \text{ for all } M_k \text{ and all } k. \quad \text{(B.8)}$$

Proof. The proof follows the results in CMS, which uses the same logic as the main result

in Lederer and Hurter (1986). Rewrite (B.6)

$$\begin{aligned}\text{Profit}_k(M_k|\mathcal{M}_{-k}) &= \int_0^1 \int_{\boldsymbol{\theta}\in\Delta(T)} \max\left\{U_{\lambda\boldsymbol{\theta}}(M_k), \bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k})\right\} g(\lambda,\boldsymbol{\theta}) d\boldsymbol{\theta} d\lambda \\ &\quad - \int_0^1 \int_{\boldsymbol{\theta}\in\Delta(T)} \bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k}) g(\lambda,\boldsymbol{\theta}) d\boldsymbol{\theta} d\lambda \\ &= \mathcal{U}(M_k|\mathcal{M}_{-k}) - \int_0^1 \int_{\boldsymbol{\theta}\in\Delta(T)} \bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k}) g(\lambda,\boldsymbol{\theta}) d\boldsymbol{\theta} d\lambda. \qquad \text{(B.9)}\end{aligned}$$

Since the second term is not a function of M_k, profit maximization is equivalent to choosing policy M_k that maximizes total expressive utility, given the other policies $\mathcal{M}_{-k}$. □

As in (12), construct the following auxiliary distribution $\hat{g}(\boldsymbol{\theta})$, which represents the distribution of vector $\boldsymbol{\theta}$ weighted by λ:

$$\hat{g}(\boldsymbol{\theta}) \equiv \frac{\int_0^1 \lambda g(\lambda,\boldsymbol{\theta}) d\lambda}{E[\lambda]}. \qquad \text{(B.10)}$$

We also need to characterize the set of consumers who purchase advice from each firm. Fix $\mathcal{M}_{-k}$, and consider experiment M_k. Note that if M_k results in the same posterior beliefs as another experiment in $\mathcal{M}_{-k}$ (no product differentiation), then firm k would have zero profits. Therefore, without loss of generality, consider the case in which M_k does not result in the same posterior beliefs as some other experiment in $\mathcal{M}_{-k}$. Define the set of fund preferences $\boldsymbol{\theta}$ that purchase advice from firm k as

$$C(M_k|\mathcal{M}_{-k}) \equiv \left\{\boldsymbol{\theta}\in\Delta(T) | U_{\lambda\boldsymbol{\theta}}(M_k) \geq \bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k})\right\}. \qquad \text{(B.11)}$$

Define $\hat{E}[\boldsymbol{\theta}|\boldsymbol{\theta}\in C(M_k|\mathcal{M}_{-k})]$ as the expected preference $\boldsymbol{\theta}$ of the consumers who purchase advice from firm k, where we use $\hat{g}(\boldsymbol{\theta})$ to assign more weight to funds with a higher λ.

Define a class of experiments which we call "targeted policies." Let the vector $\boldsymbol{\alpha}\in\Delta(T)$ represent weights attached to the different dimensions, and consider a binary experiment $M^{\boldsymbol{\alpha}}$ that recommends approval of a proposal $\boldsymbol{r}$ if and only if this proposal has a positive value when evaluated with this weight, $\langle\boldsymbol{\alpha}, r\rangle > 0$. If $\boldsymbol{\alpha} = \boldsymbol{\theta}$, then this recommendation allows fund $(\lambda,\boldsymbol{\theta})$ to make the same voting decisions that it would make if it had full information about all proposals; hence, $M^{\boldsymbol{\alpha}}$ maximizes $U_{\lambda\boldsymbol{\theta}}(M)$. Therefore, we say that policy $M^{\boldsymbol{\alpha}}$ is a policy "targeted" at fund $\boldsymbol{\theta} = \boldsymbol{\alpha}$. The next lemma shows that the optimal experiment is a targeted policy.

Lemma B.2. *Fix $\mathcal{M}_{-k}$ and suppose that experiment M_k maximizes firm's k profit. Then this experiment is a target policy $M^{\boldsymbol{\alpha}}$, where policy $\boldsymbol{\alpha}$ targets the (weighted) average consumer of firm k,*

$$\boldsymbol{\alpha} \quad = \quad \hat{E}[\boldsymbol{\theta}|\boldsymbol{\theta} \in C(M_k|\mathcal{M}_{-k})]. \tag{B.12}$$

Proof. Fix $\mathcal{M}_{-k}$ (which may or may not be targeted experiments) and suppose that experiment M_k maximizes firm's k profit. By contradiction, suppose that M_k is not a targeted policy. That is, M_k generates a pair of posterior beliefs $\boldsymbol{r}^+$ and $\boldsymbol{r}^-$ that is not generated by any targeted policy $M^{\boldsymbol{\alpha}}$, $\boldsymbol{\alpha} \in \Delta(T)$. We will construct an alternative experiment that strictly increases the profit of Firm k, which yields the contradiction.

First, suppose that the set $C(M_k|\mathcal{M}_{-k})$ is empty or has no interior point. This implies that no consumer is purchasing the advice, or consumers who are purchasing are paying a price of zero (they are indifferent between the advice from k and another option). In this case, since the number N of firms is finite, there always exists a targeted experiment $M^{\boldsymbol{\alpha}}$ target at some consumer $\boldsymbol{\theta}$ which results in this consumer and all funds sufficiently close to $\boldsymbol{\theta}$ choosing to purchase this advice from k at a strictly positive price. This strictly increases profit, a contradiction.

Second, suppose that the set $C(M_k|\mathcal{M}_{-k})$ has an interior point. Note that the boundaries of C are either consumers who are indifferent between M_k and another option, or are consumers at the boundary of the simplex $\Delta(T)$ who strictly prefer M_k over the alternatives. Recall that consumers that are indifferent between M_k and another alternative are paying a price of zero, hence they do not generate profit for firm k.

Since λ is a level shift of preferences (it does not change how funds rank different experiments), we abuse notation and define

$$\begin{aligned} U_{\boldsymbol{\theta}}(M) \quad &\equiv \quad \frac{U_{\lambda\boldsymbol{\theta}}(M)}{\lambda}, \\ \bar{U}_{\boldsymbol{\theta}}(\mathcal{M}_{-k}) \quad &\equiv \quad \frac{\bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k})}{\lambda}. \end{aligned}$$

Rewrite (B.7)

$$\mathcal{U}(M_k|\mathcal{M}_{-k}) \equiv E[\lambda] \left\{ \int_{\boldsymbol{\theta} \in C(M_k|\mathcal{M}_{-k})} U_{\boldsymbol{\theta}}(M_k)\hat{g}(\boldsymbol{\theta})d\boldsymbol{\theta} + \int_{\boldsymbol{\theta} \notin C(M_k|\mathcal{M}_{-k})} \bar{U}_{\boldsymbol{\theta}}(\mathcal{M}_{-k})\hat{g}(\boldsymbol{\theta})d\boldsymbol{\theta} \right\} \tag{B.13}$$

The firm must choose the advice M_k that maximizes this function. Construct experiment $\tilde{M}_\epsilon$, indexed by parameter $\epsilon > 0$, as follows. With probability ϵ, experiment $\tilde{M}_\epsilon$ generates a recommendation according to the targeted policy $M^{\boldsymbol{\alpha}^*}$, where $\boldsymbol{\alpha}^* = \hat{E}[\boldsymbol{\theta}|\boldsymbol{\theta} \in C(M_k|\mathcal{M}_{-k})]$. With probability $1-\epsilon$, experiment $\tilde{M}_\epsilon$ generates a recommendation according to the original policy M_k. This implies that

$$\begin{aligned} U_{\boldsymbol{\theta}}(\tilde{M}_\epsilon) &= \epsilon \left[U_{\boldsymbol{\theta}}(M^{\boldsymbol{\alpha}^*})\right] + (1-\epsilon)\left[U_{\boldsymbol{\theta}}(M_k)\right], \\ \frac{\partial U_{\boldsymbol{\theta}}(\tilde{M}_\epsilon)}{\partial \epsilon} &= U_{\boldsymbol{\theta}}(M^{\boldsymbol{\alpha}^*}) - U_{\boldsymbol{\theta}}(M_k). \end{aligned}$$

We will show that if ϵ is sufficiently small, then $\tilde{M}_\epsilon$ is strictly more profitable than M_k.

From (B.13), take the derivative of $\mathcal{U}(\tilde{M}_\epsilon|\mathcal{M}_{-k})$ with respect to ϵ evaluated at $\epsilon = 0$:

$$\begin{aligned} \left.\frac{\partial \mathcal{U}(M_k|\mathcal{M}_{-k})}{\partial \epsilon}\right|_{\epsilon=0} &\equiv E[\lambda]\left\{\int_{\boldsymbol{\theta}\in C(M_k|\mathcal{M}_{-k})} \frac{\partial U_{\boldsymbol{\theta}}(\tilde{M}_\epsilon)}{\partial \epsilon}\hat{g}(\boldsymbol{\theta})d\boldsymbol{\theta}\right\} \\ &= E[\lambda]\left\{\int_{\boldsymbol{\theta}\in C(M_k|\mathcal{M}_{-k})} \left[U_{\boldsymbol{\theta}}(M^{\alpha^*}) - U_{\boldsymbol{\theta}}(M_k)\right]\hat{g}(\boldsymbol{\theta})d\boldsymbol{\theta}\right\} \\ &\propto U_{\hat{E}[\boldsymbol{\theta}|\boldsymbol{\theta}\in C(M_k|\mathcal{M}_{-k})]}(M^{\alpha^*}) - U_{\hat{E}[\boldsymbol{\theta}|\boldsymbol{\theta}\in C(M_k|\mathcal{M}_{-k})]}(M_k) > 0. \end{aligned}$$

When taking this derivate, we can ignore the fact that a marginal change in ϵ changes the set of consumers C who purchase advice from firm k because this has a zero effect on profit. Marginal consumers (the ones just indifferent between purchasing from firm k or not) are paying a price equal to marginal cost (which is zero) and generate zero profit. Hence, gaining or losing marginal consumers does not change profits. In the proofs of CMS, this zero-profit effect was captured by the marginal change in $\tilde{\theta}$, which was the identity of the fund indifferent between two different policies.

By construction, $\boldsymbol{\alpha}^* = \hat{E}[\boldsymbol{\theta}|\boldsymbol{\theta} \in C(M_k|\mathcal{M}_{-k})]$ implies that $M^{\boldsymbol{\alpha}^*}$ provides the highest expected utility to the average fund $\hat{E}[\boldsymbol{\theta}|\boldsymbol{\theta} \in C(M_k|\mathcal{M}_{-k})]$. Moreover, we are considering the case in which M_k results in different posterior beliefs than $M^{\boldsymbol{\alpha}^*}$, hence it yields a strictly lower payoff for fund $\hat{E}[\boldsymbol{\theta}|\boldsymbol{\theta} \in C(M_k|\mathcal{M}_{-k})]$. Therefore, the last inequality holds, which is a contradiction to M_k being an optimal policy, concluding the proof. □

The following proposition summarizes our results.

Proposition B. 1. *There exists a competitive equilibrium in which firms choose policies that maximize the expected expressive voting utility. In every competitive equilibrium, firms*

optimally use targeted policies $M^{\alpha_k^*}$, where policy α_k^* targets the average preference of the funds purchasing advice from firm k, with preferences weighted by λ.

Proof. From Lemma B.2, policies that maximize total expressive voting utility and policies that are best responses for firms must be targeted policies. Consider a social planner choosing N targeted policies $\{\alpha_1, \ldots, \alpha_N\} \in \Delta(T)^N$ to maximize the total expressive voting utility $\mathcal{U}(M^{\alpha_1}, \ldots, M^{\alpha_N})$. Since the domain is compact and the function is continuous, there exists a set of targeted policies that maximizes $\mathcal{U}$. From Lemma B.1, this set of policies constitutes a competitive equilibrium. □

B.2.4. Customized Advice

We now characterize the equilibrium when each firm can offer multiple advice policies. Suppose that there are N firms, and each firm k can offer a finite number Z_k of policies. Let M_{kz} denote one of the policies offered by firm k, with $\mathcal{M}_k = \{M_{k1}, \ldots, M_{kZ_k}\}$ the set of of policies offered by firm k, and $\mathcal{M}_{-k}$ be the set of policies offered by all other firms, including the option of not acquiring information, M_0. $\mathcal{M}$ represent all polices.

Using backward induction, consider the subgame after $\mathcal{M}$ has been chosen. Consider the price offer of firm k to fund $(\lambda, \boldsymbol{\theta})$. Let $\bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_k) = \max\{U_{\lambda\boldsymbol{\theta}}(M_{kt})\}_{M_{kt} \in \mathcal{M}_k}$ be the highest expected utility that fund $(\lambda, \boldsymbol{\theta})$ can obtain if it selects the best available advice from firm k. Similarly define $\bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k})$ as the highest expected utility that fund $(\lambda, \boldsymbol{\theta})$ can obtain if it selects the best available advice from the other firms (or the outside option M_0). Firm k optimally charges a price equal to $\max\{\bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_k) - \bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k}), 0\}$. The next Lemma generalizes Lemma B.1.

Lemma B.3. *Given $\mathcal{M}_{-k}$, the set of experiments $\mathcal{M}_k$ maximizes firm's k profit if and only if $\mathcal{M}_k$ maximizes the total expressive utility. Therefore, policies $\mathcal{M}^*$ form a competitive equilibrium if and only if*

$$\mathcal{U}(\mathcal{M}_k^* | \mathcal{M}_{-k}^*) \;\geq\; \mathcal{U}(\mathcal{M}_k | \mathcal{M}_{-k}^*) \text{ for all } \mathcal{M}_k \text{ and all } k. \tag{B.14}$$

Proof. We essentially follow the proof of Lemma B.1 and rewrite the firm's profit as

$$\begin{aligned}
\text{Profit}_k(\mathcal{M}_k|\mathcal{M}_{-k}) &= \int_0^1 \int_{\boldsymbol{\theta}\in\Delta(T)} \max\left\{\bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_k) - \bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k}), 0\right\} g(\lambda,\boldsymbol{\theta}) d\boldsymbol{\theta} d\lambda \\
&= \int_0^1 \int_{\boldsymbol{\theta}\in\Delta(T)} \max\left\{\bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_k), \bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k})\right\} g(\lambda,\boldsymbol{\theta}) d\boldsymbol{\theta} d\lambda \\
&\quad - \int_0^1 \int_{\boldsymbol{\theta}\in\Delta(T)} \bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k}) g(\lambda,\boldsymbol{\theta}) d\boldsymbol{\theta} d\lambda \\
&= \mathcal{U}(\mathcal{M}_k|\mathcal{M}_{-k}) - \int_0^1 \int_{\boldsymbol{\theta}\in\Delta(T)} \bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k}) g(\lambda,\boldsymbol{\theta}) d\boldsymbol{\theta} d\lambda. \qquad \text{(B.15)}
\end{aligned}$$

Since the second term is not a function of $\mathcal{M}_k$, profit maximization is equivalent to choosing a set of policies $\mathcal{M}_k$ that maximizes total expressive utility, given the other policies $\mathcal{M}_{-k}$. □

Given that each firm wants to choose policies that maximize the total expressive voting utility, it is straightforward to generalize our previous results. The best response of each firm must be to choose targeted policies, with each policy targeting the (weighted) average consumer that buys this policy. The set of targeted policies is compact and the total expressive utility is continuous, therefore a maximum exists and these policies form a competitive equilibrium.

B.2.5. Voting Cost

We now extend the model to consider the case in which each fund faces a cost to maintain a system to cast votes.

A fund that does not purchase an advice needs to pay a cost $c_0 \geq 0$ to keep the infrastructure necessary to cast votes. Recall that there is a measure one of proposals to be voted on, and the expected benefit of casting an uninformed vote is $\lambda|\langle\boldsymbol{\theta}, \hat{\boldsymbol{r}}\rangle|$ — see (B.5). Therefore, a fund that does not purchase advice will pay for the voting system and vote if $\lambda|\langle\boldsymbol{\theta}, \hat{\boldsymbol{r}}\rangle| \geq c_0$, and will not vote and attain a payoff of zero otherwise. Abusing notation, we can rewrite (B.5) and define the *net* value of the outside option M_0 as

$$U_{\lambda\boldsymbol{\theta}}(M_0) = \max\{\lambda|\langle\boldsymbol{\theta}, \hat{\boldsymbol{r}}\rangle| - c_0, 0\}. \qquad \text{(B.16)}$$

We now consider firms that sell a bundled service: when purchasing advice from firm k, the firm will offer the advice and a vote execution service. Firm k incurs a cost $c_k \geq 0$

per fund to offer this service, but it also adjusts the price it charges, given the market competition. Note that we are allowing different firms to have different costs c_k. We assume that $c_k \leq c_0$, so that it is never optimal for funds to purchase an advice but not purchase the voting service. Abusing notation, we can use (B.4) to define the net value of the information provided by firm k,

$$U_{\lambda\boldsymbol{\theta}}(M) \quad = \quad \lambda \left[\Pr(m^+) \langle \boldsymbol{\theta}, \boldsymbol{r}^+ \rangle - \Pr(m^-) \langle \boldsymbol{\theta}, \boldsymbol{r}^- \rangle \right] - c_k. \tag{B.17}$$

Using backward induction, consider the subgame following policy choices $\mathcal{M}$. In the benchmark model, allocation of funds to firms was efficient, in the sense that each fund would always select the option that offers the highest value of information, and pay a price equal to the difference between the highest value and the value of the second highest option available. Now, allocations will be efficient in the sense that each fund will select the option that offers the highest *net value* of information (net of the voting cost c_k) — the price paid for advice will be the voting cost c_k plus the difference between the net value of the highest option available and the second highest net value available.

That is, consider the price that firm k will offer to fund $(\lambda, \boldsymbol{\theta})$. Let $\bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k})$ be the maximum net value that the fund can obtain from all the other options. If $U_{\lambda\boldsymbol{\theta}}(M_k) \leq \bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k})$, then the best the fund can do is to offer a price equal to its marginal cost c_k. The firm will serve this fund with some probability if the fund is indifferent, or not serve the fund. The firm earns zero profit in both cases. If $U_{\lambda\boldsymbol{\theta}}(M_k) > \bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k})$, then the firm will optimally offer the price $c_k + U_{\lambda\boldsymbol{\theta}}(M_k) - \bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k})$, and will serve this fund at the cost c_k. Therefore, the firm's profit from this fund is $U_{\lambda\boldsymbol{\theta}}(M_k) - \bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k})$.

Thus, profit equation (B.6) continues to hold, and we reinterpret equation (B.7) as the total expressive utility net of the voting cost. Consequently, all results and proofs in this Online Appendix continue to hold. In particular, since the cost c_k is the same for all consumers purchasing advice from firm k, the optimal signal must be a targeted policy, where the optimal $\boldsymbol{\alpha}$ targets the average consumer of firm k, weighted by λ. Moreover, there is a set of targeted policies $\{\boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_N\}$ that maximizes the total expressive voting utility net of the voting cost, and these policies constitute a competitive equilibrium.

The only substantive change is that we need to adjust (B.18), which defined the set of consumers purchasing advice from firm k. In the original definition, all consumers with the

same preference $\boldsymbol{\theta}$ would choose the same firm, independently of λ. Now choices over firms (and the outside option) depend on λ: the set of consumers purchasing for firm k is

$$C(M_k|\mathcal{M}_{-k}) \equiv \left\{(\lambda,\boldsymbol{\theta}) \in (0,1] \times \Delta(T)|U_{\lambda\boldsymbol{\theta}}(M_k) \geq \bar{U}_{\lambda\boldsymbol{\theta}}(\mathcal{M}_{-k})\right\}. \tag{B.18}$$

This generalization allows us to uncover interesting new results. Funds with a low λ will focus on the voting execution service cost and will prefer to purchase advice from the firm with the lowest c_k. Funds with a high λ will prefer the firm offering the advice policy closest to their preferences.

To see this, consider proposals with two dimensions (s,r) as in CMS. Consider a duopoly market where firms 1 and 2 offer target policies $\alpha_1 < \alpha_2$, and suppose voting is mandatory. If both firms have the same cost of providing voting services, $c_1 = c_2$, then there is a cutoff type $\tilde{\theta} \in (\alpha_1, \alpha_2)$ such that all funds $\theta < \tilde{\theta}$ strictly prefer firm 1 while all funds $\theta > \tilde{\theta}$ prefer firm 2, independently of λ. Now suppose that firm 2 has a lower cost, $c_1 > c_2$. In this case, a fund will prefer firm 2 if

$$\begin{aligned} \lambda U_\theta(M^{\alpha_2}) - c_2 &> \lambda U_\theta(M^{\alpha_1}) - c_1 \\ \Rightarrow U_\theta(M^{\alpha_2}) - U_\theta(M^{\alpha_1}) &> \frac{-(c_1 - c_2)}{\lambda}. \end{aligned}$$

The inequality holds for all funds $\theta > \tilde{\theta}$, so they continue to purchase from firm 2. However, since the RHS is strictly negative, now all funds $\theta < \tilde{\theta}$ with a sufficiently low λ will prefer to purchase the service from firm 2. Funds with a low λ are more sensitive to voting costs. In some sense they are more price elastic: they are eager to switch from a policy α_1 that they prefer to a worse policy α_2 if the price is sufficiently lower.

Moreover, potential entrants might not be able to enter this market if their cost to provide voting execution services is too high. In our duopoly example, if c_2 is low and c_1 is sufficiently high, then all consumers (even consumers with $\lambda = 1$) would prefer to purchase advice from firm 2, who would become a monopolist.

Finally, if voting is not mandatory and voting cost is strictly positive, then all funds with a sufficiently small λ would not purchase advice and would abstain.

References

[1] Câmara, Odilon, John Matsusaka, Chong Shu, "Shareholder Democracy and the Market for Voting Advice," *working paper*, 2025.

[2] Lederer, Phillip J., and Arthur P. Hurter Jr., "Competition of Firms: Discriminatory Pricing and Location," *Econometrica*, 1986, 54(3), 623-640.

[3] Malenko, Andrey, Nadya Malenko, and Chester S. Spatt, "Creating Controversy in Proxy Voting Advice," *Journal of Finance*, 2025.

[4] Perego, Jacopo and Sevgi Yuksel, "Media Competition and Social Disagreement," *Econometrica*, January 2022, Vol. 90(1), 223-265.