# 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*

September 12, 2024

**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 through prices and their advice policies. Shareholders have heterogeneous goals, differing in the weight they place on financial returns versus nonfinancial ("social") returns, such as reductions in carbon emissions. We assume that investors vote for expressive reasons and may differ in how much they care about voting correctly. In equilibrium, advising firms tailor their advice to reflect the preference of their average investor, which can result in election outcomes being skewed away from those that would prevail if investors had full information. We derive conditions under which advisory firms skew their advice and therefore the election outcome in favor of a minority of investors who have a strong preference for nonfinancial returns. We also study how increasing competition affects equilibrium outcomes.

*This paper incorporates some ideas from the working paper, "A Theory of Proxy Advice when Investors Have Social Goals" (2021), which is now retired. 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 Hong Kong Baptist University, Nanyang Technological University, National University of Singapore, and USC Marshall.

# 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 provided 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 the economic forces that shape proxy advice and voting in corporate elections. An important feature of our model – and one that distinguishes it from previous research – is that investors have heterogeneous goals: even if they have full information about a proposal, they may disagree about whether to approve it. Specifically, we assume that investors 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 types of preference heterogeneity, such as disagreement on short-run versus long-run returns. We also assume that investors vote for expressive reasons and differ in how much they care about voting correctly. For example,

[1] As of 2019, institutional investors held 70 percent of corporate equity, compared to 30 percent held by retail investors (Broadridge + PwC, 2019). For an overview of institutions, laws, and issues, see Edelman et al. (2014) and Gallagher (2014).

[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. Larcker et al. (2013) and U. S. House of Representatives (2013, 2024) criticize proxy advice for the quality of the underlying research.

the voting record of a public pension fund might be scrutinized by labor unions, media, and politicians, while a small passive mutual fund might be entirely unconcerned about its voting record.

In the model, investors vote on many proposals, each characterized by a level of financial and social returns, which 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. Each advisory firm chooses an advice policy that consists of the relative weight that it places on the expected social versus financial return when evaluating proposals. Using its advice policy, an advisory firm determines whether each proposal generates a positive or a negative expected return, and then recommends approval or rejection of the proposal accordingly. The prices of advice are the outcome of discriminatory price competition.

For an investor, the ideal advice policy would be the one that assigns the same weights to social and financial returns as the investor's own preference. Such a policy maximizes the value of the information provided by the recommendation. The investor is willing to pay less for recommendations produced by advice policies that use weights farther from the ideal weights.

We begin by considering a market monopolized by a single advisory firm. In interviews, proxy advisors emphasize that their "guidelines come from an aggregation of institutional investor preferences [for which] issues are not viewed uniformly across investors" (Hayne and Vance, 2019, p. 975). We characterize how this aggregation occurs: the proxy advisor maximizes profit by choosing a recommendation policy that targets the customer with the average (rather than the median) preference, weighted by how much investors care about their voting reputation. This has two implications. First, advice is slanted toward the preferences of funds with the highest expressive benefits from voting. If, as we argue is likely, socially responsible investment (SRI) funds have the highest expressive benefits, the model predicts that voting recommendations are slanted toward their preferences, and as a result, their interests are overweighted in corporate elections. Second, because the recommendation policy preferred by the average voter is generally different than the policy preferred by the median voter, proxy advice does not result in the conventional democratic outcome of majority rule.

If the proxy advisor were to offer advice targeted at the median voter, then the outcome of elections would be the same as if shareholders had free access to fully informative signals about the proposals. We argue that biasing outcomes away from the median fund and toward the average fund is normatively undesirable: it allows the preference of less than a majority to determine the outcome; it differs from the outcome that would prevail under full information; and the distortions are driven by expressive benefits – how much a fund benefits from expressing its views – that lack a compelling normative value. As our normative benchmark, therefore, we introduce the concept of "informed representation," which is the election outcome that would prevail if all voters were fully informed and the majority (median voter) ruled.

We then extend the model to explore the role of competition in the advice market. We show that an increase in the number of advisory firms leads them to offer differentiated products, advice policies with different weights on the return dimensions. Advice policy differentiation increases the total expressive voting utility of investors. However, differentiation may end up offering the median investor advice that is less aligned with its preferences, undermining the informed representation of corporate elections. We identify conditions under which competition improves versus deteriorates the informed representation of corporate elections.

We also extend the model to study a case of considerable practical importance: the regulatory requirement that funds vote in every election. In practice, U.S. regulations do not permit pension and mutual funds to abstain from voting. The Department of Labor ruled in 1988 that pensions had a fiduciary duty to vote under ERISA, and the SEC issued a no-action letter in 2003 stating that mutual funds had a fiduciary duty to vote. The SEC later elaborated that funds "could demonstrate that the vote was not a product of a conflict of interest if it voted client securities in accordance with a pre-determined policy, based upon the recommendations of an independent third party."[3] This was broadly interpreted

[3]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. SEC: Proxy Voting by Investment Advisors, 68 Federal Register 6585, February 7, 2003. Investment advisors, strictly speaking, are not required to vote their proxies and enforcement actions are rare, yet 90 percent of them choose to do so (SEC Staff Legal Bulletin No. 20, June

to mean that funds are presumed to have acted as proper fiduciaries if they follow the recommendations of a proxy advisor when voting. As a result, most funds vote, and if they do not conduct their own independent research, they utilize recommendations from a proxy advisor. We use the model to explore how this mandate affects the quality of advice, its slant, and the outcomes of corporate elections.

To study the effect of mandatory voting, we assume that a small fraction of funds are endowed with private information, while the remaining funds are uninformed and must acquire advice if they wish to cast informed votes. We show that requiring the uninformed funds to vote can distort corporate elections. The intuition is that without a voting mandate, the uninformed funds abstain, so that election outcomes are determined by the median of the informed funds. If uninformed funds are forced to participate and base their votes on proxy advice, then the median voter becomes one of the uninformed funds. As a result, the election outcome is determined by the recommendation of an advisory firm, which is slanted toward the preferences of investors with the highest expressive voting benefit.

As noted, an important assumption in our analysis is that investors have heterogeneous preferences over corporate policies. When voters have common values, the Condorcet Jury Theorem tells us that voting is very likely to achieve the efficient outcome even with a small almost-uninformed electorate. However, without common values, the law of large numbers does not necessarily lead to better representation, and the traditional focus on efficiency as a normative principle loses force. For this reason, we develop a new normative principle, as mentioned above, that we call "informed representation." With this benchmark, we can rank different market structures according to how close their electoral outcomes are to the democratic full-information benchmark. We show that this ranking coincides with the preference ranking of the median fund.

These normative issues connect our analysis to an ongoing discussion about the appropriate objective function for firms in a world where some investors care about more than financial returns (Zingales et al., 2020). When investors care only about monetary returns, standard neoclassical principles imply that firms should maximize value or profit, as famously

30, 2014; Broadridge + PwC, 2019). Securities and Exchange Commission, Final Rule: Proxy Voting by Investment Advisers, 17 CFR Part 275, available at: https://www.sec.gov/rules/final/ia-2106.htm.

argued by Milton Friedman (Friedman, 1970; Fama and Miller, 1972). However, when some investors care about nonpecuniary issues such as human rights, there is no compelling theoretical basis for adopting value maximization as the appropriate goal. Hart and Zingales (2017) argue that shareholder welfare maximization would be more appropriate. This idea is appealing, but requires managers to be able to determine shareholder preferences. Voting seems like a natural way to gauge preferences, but our analysis shows that election outcomes are not guaranteed to accurately mirror voter preferences if votes are based on recommendations from profit-maximizing advisory firms. Indeed, under some conditions advice is slanted toward a minority group of biased investors.

Theoretical research on the proxy advice industry is in its early stages. Malenko and Malenko (2019) develop a strategic voting model in which funds can buy advice from a single advisory firm that provides an unbiased but noisy recommendation, and also acquire information through independent research. They explore how proxy advice can crowd out funds' own acquisition of information. Buechel et al. (2024) consider a similar strategic voting model with a single advisory firm in which investors also receive advice from the board. They show that having two sources of advice may increase the incentive of funds to acquire their own information if the sources offer conflicting advice. Ma and Xiong (2021) also study a strategic voting model with a single advisory firm. In their model, the firm may offer biased advice when investors have a wrong prior belief or are not value-maximizers. Malenko et al. (2021) develop a strategic voting model with a single advisory firm that provides a free public recommendation and sells a more detailed private recommendation. They show that in equilibrium the proxy advisory firm's public recommendations are biased against the a priori most likely election outcome in order to make elections more competitive and boost demand for their private recommendations.

All of these models assume a monopoly advisory firm, abstracting away from the market structure issues that we explore. We believe our study is the first to investigate the disciplining role of competition in the proxy voting advice market, and the interaction between market structure, the quality of advice, and the representativeness of corporate elections. These models also abstract away from conflicts *between* investors: they all assume that investors have homogeneous preferences, in the sense that they consider a binary state space (a proposal is either good or bad) and all voters would agree on how to vote if they knew the realized state. In contrast, our analysis features a rich state space and heterogenous voters, giving rise to a slant in advice that mirrors the findings from empirical studies. In this respect, perhaps the closest study to ours is Levit and Tsoy (forthcoming), which studies a monopolist advisor giving advice to two agents with different preferences in a cheap talk communication game; they show how an advisor may adopt one-size-fits-all recommendations in order to obscure its biases. While their advisor has preferences over voting outcomes and offers free advice, our model features profit-maximizing advisors who compete to sell information. In the equilibrium of our model, the advice may be slanted towards the social dimension not because the advisor has an intrinsic preference bias, but because it is the profit-maximizing choice.

Moving away from corporate governance, our model shares some economic features with models of information markets, particularly Perego and Yuksel (2022), in which competitive firms sell information to voters. In their paper, the advisor (media outlet) provides a point estimate of the value of the voting options, such as the choice between two candidates for governor; while we assume that the advisor sells a simple binary recommendation, which is the form of many voting recommendations in practice.

## II. Model

### *A. Motivating Example*

To motivate our assumptions, consider the familiar example of a juror in a trial. The juror wants to vote correctly: "guilty" if the defendant is guilty and "not guilty" if the defendant is innocent. The juror receives a positive payoff from voting correctly and a negative payoff from voting incorrectly, the size of which depends on the stakes of the case. If the case is about a serious crime with a heavy sanction, then the juror receives a large positive payoff from voting correctly, and a large negative payoff from voting incorrectly. If the case is a minor crime, the juror receives a small gain or loss for casting the right or wrong vote. There may also be heterogeneity across jurors. Some jurors might care more about voting correctly than others, and they may have different views of what should be considered a violation of the law.

*B. Investors*

Our voting model incorporates all of the above features. Instead of jurors voting guilty or not guilty, we consider investment funds voting "yes" or "no" on proposals. An investment fund or investor $i$ is a shareholder and votes to "approve" (yes) or "reject" (no) a proposal $j$. The vote is an action $a_{ij} \in \{-1, 1\}$ , where $a_{ij} = 1$ represents approval and $a_{ij} = -1$ is rejection.

In line with the juror example, our model assumes that each fund cares about voting correctly (it cares about its voting record), captured by variable $v_{ij} \in \mathbb{R}$, which represents the value of making a correct voting decision. The sign of $v_{ij}$ indicates how fund $i$ should vote: a positive $v_{ij}$ means that the vote should be yes, and a negative $v_{ij}$ means that the vote should be no. The absolute value of $v_{ij}$ indicates the importance of voting correctly on this proposal. Formally, the fund's payoff is $v_{ij}$ if it votes to approve and $-v_{ij}$ if it votes to reject:

$$u_i(v_{ij}, a_{ij}) = a_{ij} v_{ij}. \tag{1}$$

There is a measure one of proposals to be voted on. In practice, the value of each proposal depends on a complex list of variables. To keep the model tractable and focus on the tradeoff between social return and financial return, we assume that the proposal's value takes the form:

$$v_{ij} = \lambda_i \left[\theta_i s_j + (1 - \theta_i) r_j\right]. \tag{2}$$

Variable $r_j \in \mathbb{R}$ captures the proposal's financial return. Variable $s_j \in \mathbb{R}$ captures the proposal's social return. Note that 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 of this, even if funds know all of the facts related to the proposal, they may differ on whether they favor approval or rejection. Preference parameter $\lambda_i \in (0, 1]$ is the overall importance that the fund attaches to voting correctly. For example, a fund targeting active green investors can advertise a strong environmental voting record to attract a larger fund flow, in which case voting correctly is very important. In contrast, an index fund targeting investors who do not pay attention to voting records is less concerned about its voting records. Each fund knows its own preference parameters, but they do not 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. Given $F$, the variances are $\sigma_s^2$ and $\sigma_r^2$, with covariance $cov_{sr}$ and Pearson correlation coefficient $\rho_{sr}$.

There is a measure one of funds, with preference parameters distributed according to the cdf $G(\lambda,\theta)$ and continuous pdf $g(\lambda,\theta)$.[4] The variances are $\sigma_\lambda^2$ and $\sigma_\theta^2$, with covariance $cov_{\lambda\theta}$ and Pearson correlation coefficient $\rho_{\lambda\theta}$. The unconditional expectations $E[\lambda]$ and $E[\theta]$ are strictly between zero and one.

Abusing notation, the (ex post) payoff of fund $(\lambda,\theta)$ voting on proposal $(s,r)$ is

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

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

Note that investors receive voting payoffs based on the vote they cast, not the election outcome: this is a model of "expressive" voting rather than "instrumental voting," a common approach in the political science literature.[5] We adopt expressive voting for several reasons. First, a typical fund's chance of casting a decisive vote in a typical election is negligible. The vast majority of corporate elections are one-sided — in only 2.7 percent of elections is the margin of victory less than 5 percent, and the typical margin is 67 percent.[6] Second, with expressive voting, demand for information is not driven by the probability of casting a pivotal vote, but as a way to attract fund flow (SRI funds) or accommodate pressure from external constituencies (public pension funds). 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,"

[4] We assume that this distribution is continuous on $\lambda$ to simplify notation. Our qualitative results continue to hold if $\lambda$ is a discrete variable. The key assumption is that the marginal distribution on $\theta$ is continuous, to guarantee that the adviser's profit (defined later) is a differentiable function of its policy choice.

[5] Downs (1957) identified the logical problem with instrumental voting theories. Fiorina (1976) is an early statement of the expressive voting theory; Brennan and Lomasky (1993) is another notable work on the idea.

[6] These are our calculations using data from 489,657 elections during 2003-2018.

immediately below which it provides its proxy votes for the last 13 years.[7]

### C. Proxy Advisory Firms

We first consider a benchmark version of our model in which: (i) there is a single firm selling advice, (ii) the firm can only sell one type of advice, and (iii) funds are not able to conduct their own private research to learn about the proposals. We later consider a more general version of the model to show how key insights from the benchmark model continue to hold when we drop these assumptions, and how the general model provides additional insights regarding how competition affects shareholder democracy.

In the benchmark model, the advisory firm (monopolist) creates a test defined by an advice policy $\alpha \in [0,1]$. We imagine that the advisory firm hires a team of specialists and designs protocols such that, for each proposal, it will recommend voting yes if $\alpha s+(1-\alpha)r > 0$ and voting no if $\alpha s + (1-\alpha)r \leq 0$. In order to credibly and accurately implement a specific policy $\alpha$, the firm needs to specialize in the particular weight it attaches to social values. To simplify exposition, we abstract from fixed and marginal costs of providing advice, and instead constrain each firm to produce at most one type of advice policy $\alpha$. We consider multiple policies per firm in Section VII.A.

It is important to emphasize that the advisory firm is not providing an estimate of $s$ and $r$ for each proposal $j$. The firm only provides a binary yes/no recommendation. We have in mind that the firm may not even collect the raw information on $r$ and $s$; rather its researchers search for critical pieces of information that reveal if the proposal lies in the accept or reject region ("checking boxes").

[7] 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. The model does not assume that the payoff from voting depends on whether it actually changes the issuer's policy, which in practice would be quite costly for investors to determine.

*D. Monopoly Prices*

The advisory firm negotiates a price directly with each fund; it does not post a single list price.[8] For most of the analysis, we assume that the advisory firm makes a take-it-or-leave-it price offer to each fund; however, our analysis goes through in its entirety if we instead assume that the advisory firm captures a fraction $\eta \in (0, 1]$ of the surplus it generates for each fund, where $\eta$ is exogenously fixed. In the general case, if the advisory firm implements a policy $\alpha$ that yields expected utility $U_{\lambda\theta}(\alpha)$ to a fund with preference $(\lambda, \theta)$, then the advisory firm charges a price $\eta U_{\lambda\theta}(\alpha)$. Note that $\eta = 1$ represents the case of a take-it-or-leave-it offer, or first-degree price discrimination. After observing the price, each fund then chooses whether to purchase the advice from the firm, or decline to purchase and receive a payoff of zero.

## III. The Value of Advice

*A. Basic Properties*

To simplify presentation and more clearly convey the value of information, we assume that the expected values of $s$ and $r$ are zero. This means 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.

Now suppose that a fund obtains recommendations from an advisory firm, generated according to some advice policy $\alpha \in (0, 1)$ that prescribes a vote yes if $s > \frac{-(1-\alpha)}{\alpha} r$ and a vote no otherwise. By following this 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 $\alpha$ 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

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

[8] Based on Public Records Act requests, Shu (2024) reports that the annual fees paid by the largest public pension funds to their proxy advisors ranged from \$50,000 to \$500,000.

The value of advice $\alpha$ for fund $(\lambda, \theta)$ is the expected value of voting according to the recommendations generated by $\alpha$, which we express 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 $\alpha$. We assume that $E[r|s > 0] \geq 0$ and $E[s|r > 0] \geq 0$ in order to guarantee a positive value of advice (see the proof of Lemma 1); a sufficient condition for this to hold is that returns $s$ and $r$ are independent; it may also hold with negatively correlated returns if there is a positive non-linear relationship between the variables.[9] The following lemma summarizes some useful properties of the value of information.

**Lemma 1.** *Fund $(\lambda, \theta)$ is willing to pay $U_{\lambda\theta}(\alpha)$ to purchase advice from an advisor offering policy $\alpha$. This value is maximized at $\alpha = \theta$, and strictly decreases as $\alpha$ moves away from $\theta$. Formally, for all $\alpha \in (0, 1)$ 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 fundamental 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. At one extreme, policy $\alpha = 1$ implies $V_s(1) = E[|s|]$ (the policy provides perfect advice on the social dimension) and $V_r(1) = \Pr(s > 0)E[r|s > 0] - \Pr(s \leq 0)E[r|s \leq 0]$ (the minimum amount of information about the financial dimension). At the other extreme, $\alpha = 1$ implies $V_s(0) = \Pr(r > 0)E[s|r > 0] - \Pr(r \leq 0)E[s|r \leq 0]$ (the minimum amount of information about the social dimension) and $V_r(0) = E[|r|]$ (perfect advice on 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 $\alpha$, which brings it closer to the fund's optimal policy $\theta$, and conversely when $\alpha > \theta$.

[9] For example, consider the following distribution: with probability 0.3, both $s$ and $r$ are independently drawn from a uniform $[0, 1]$ distribution; with probability 0.3, both $s$ and $r$ are independently drawn from a uniform $[-1, 0]$ distribution; with probability 0.4, $s$ and $r$ are drawn from a bivariate normal distribution with mean zero, variance one, and covariance -0.4. This implies $E[s|r > 0] = E[r|s > 0] \approx 0.03$ and $cov_{sr} = -0.01$.

## B. *Example*

We next provide an example to build 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 $\alpha = 0.5$, which means equal weight on both dimensions. Figure 1 depicts a set of random proposals drawn from $F$. Each point represents a realized proposal $j$, with the financial return $r_j$ on the horizontal axis and the social return $s_j$ on the vertical 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 $\alpha$ is represented by a clockwise rotation of the dashed line. The dashed line overlays the vertical axis when the policy assigns no weight to the social dimension ($\alpha = 0$) and overlays the horizontal axis when the policy assigns no weight to the financial dimension ($\alpha = 1$).

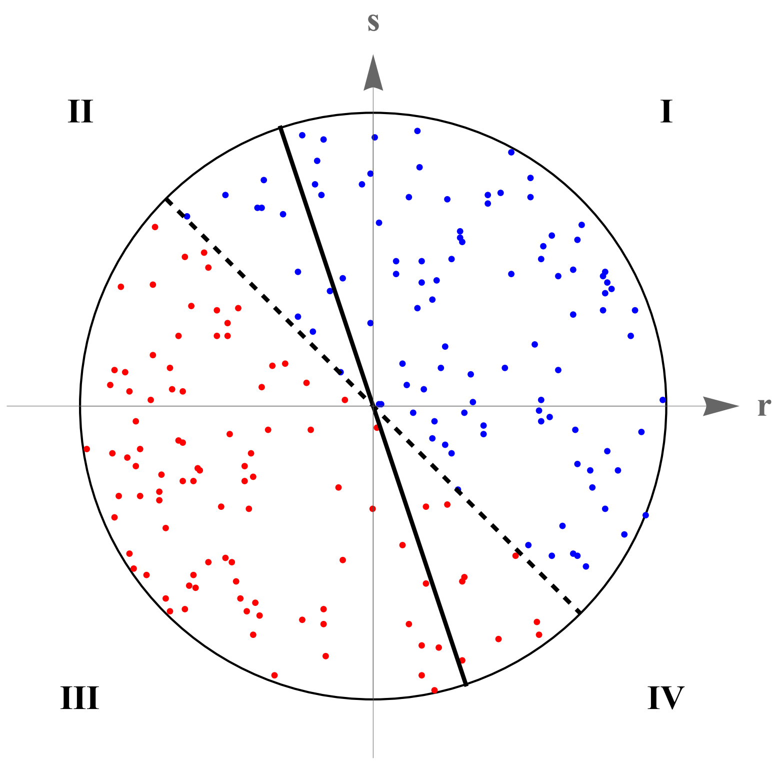


**Figure 1: Proposals uniformly distributed on a disk.** Each dot represents a proposal. The dashed line represents the advice policy $\alpha = 0.5$. The solid line represents a fund with $\theta = 0.25$.

Consider a fund with preference $\theta = 0.25$. The solid line in Figure 1 has a slope $-(1-\theta)/\theta = -3$ and passes through the origin. This line represents the fund's indifference curve — along this line the fund is indifferent between accepting and rejecting proposals. The fund would like to vote yes on proposals above the solid line and vote no on proposals below the line, but without knowledge about the realized values of proposals it must rely on the advisor's recommendations.

Independent of $\alpha$ and $\theta$, 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. All proposals in the south-west quadrant III deliver non-positive values in both dimensions, hence every policy $\alpha \in [0, 1]$ recommends rejection and every fund $\theta \in [0, 1]$ prefers 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. A higher policy $\alpha$ means that more of those proposals receive an approval recommendation. Conflict arises when $\alpha \neq \theta$. In Figure 1, we assume $\alpha > \theta$, implying that the advisor's policy overweights the social dimension relative to the fund's ideal policy. 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. 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 $\alpha$, bringing it closer to $\theta$. 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 $\alpha$.

To illustrate the payoff functions, consider two funds that have the same $\lambda$, 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 $U_{\lambda\theta}(\alpha)$ 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.

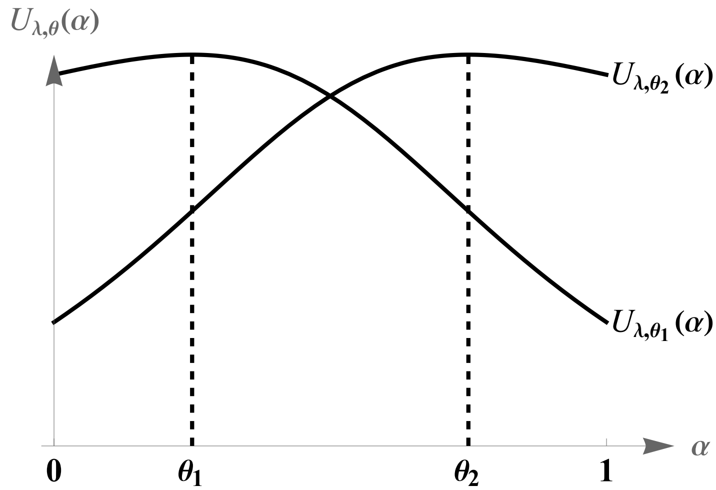


**Figure 2: Utility as a function policy $\boldsymbol{\alpha}$.** Proposals are uniformly distributed on a disk, and the two funds have preferences $\theta_1 = 0.25$ and $\theta_2 = 0.75$, with the same $\lambda$.

### *C. Advice and the Effectiveness of 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. Our ranking criterion focuses on election outcomes – whether proposals are approved or rejected – not on the individual vote decisions except insofar as they determine the outcomes. We assume that the optimal election outcomes are those that would occur if all funds had perfect information and election outcomes were determined by majority vote; and develop a method to rank alternative outcomes in descending order according to their distance – in a sense to be made precise – from the optimal outcomes. An alternative utilitarian approach would be to rank outcomes based on aggregate expected expressive voting utility. We discuss the tradeoffs of our approach versus others after defining our concept.

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 $\tau_i$ 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 $\theta$, within this group order funds from smallest to largest according to their ownership share.

Election outcomes are determined by majority voting, so that proposal $j$ is approved if a majority of funds vote yes, where votes are weighted by the fraction of shares owned by a fund, and rejected otherwise. That is, if a subset of funds $A$ vote to approve a certain proposal $j$ and a subset $R$ vote to reject it, then the proposal is approved if $\int_{i\in A} \tau_i di > \int_{i\in R} \tau_i di$, and rejected otherwise. 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 the fund's preference parameter. Note that the median fund is not the median $\theta_m$ from 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 $\theta_m$ is decisive. That is, if the median fund prefers to approve a proposal $j$, then a majority of funds prefer to approve proposal $j$, and conversely.*[10]

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. We next argue that the median's preference is also a sufficient statistic to define how a majority of funds rank different election outcomes under partial information.

Consider a market structure $k$ in which different funds may have different amounts of information. This market structure encompasses the number of firms in the market, the equilibrium strategies of all players, the information available to different 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 voting funds (weighted by their size), and $q_k(s,r) = -1$ if it is rejected. For fund $(\lambda, \theta)$, the ex ante expected payoff from election outcomes $q_k$ is

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

[10] If the median fund prefers to approve a project, 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 $\theta < \theta_m$, $\theta s + (1-\theta) r > 0$. Similar logic applies when the median fund prefers to reject a project.

where the expectation is taken over all possible values of $(s, r)$ according to $F$. Note the difference between $U_{\lambda\theta}(\alpha)$ in (6) and $\hat{U}_{\lambda\theta}(q_k)$ in (7). The payoff $U_{\lambda\theta}(\alpha)$ captures a fund's expressive voting payoff from voting according to advice policy $\alpha$. The payoff $\hat{U}_{\lambda\theta}(q_k)$ in (7) captures a fund's ex post utility as a result of the election outcomes $q_k$. When choosing an advisory firm, a fund maximizes $U_{\lambda\theta}(\alpha)$ because it votes for expressive reasons.

Comparing across two outcome functions, fund $(\lambda, \theta)$ weakly prefers $q_A$ to $q_B$ if it delivers a higher expected policy payoff, that is, if

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

Hence, from this policy outcome perspective, we say that fund $(\lambda, \theta)$ weakly prefers market $q_A$ whenever (8) holds.

**Remark 2.** *Consider two different market structures, with election outcomes $q_A$ and $q_B$. If the median fund weakly prefers $q_A$ over $q_B$, then a majority of funds weakly prefer $q_A$ over $q_B$.*[11]

If we imagine asking funds to choose between $q_A$ and $q_B$, the outcome favored by the median fund would attract majority support. The median fund's preference is thus a sufficient statistic to rank the popularity of different $q$'s among the funds. Because the information available to voters determines the election outcomes $q_k$, 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 method for ranking different 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.** *Market structures can be ranked by the median voter's preference $\Psi(\cdot)$. Specifically, if $\Psi(q_A) > \Psi(q_B)$, then a majority of funds receive a higher expected policy payoff*

[11] 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.

*under $q_A$ than $q_B$. In this 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 $k$ where all funds have full information about proposals is

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

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 $q_{\text{ No Information}}(s, r) = 0$ for all projects, and $\Psi(q_{\text{ No Information}}) = 0$.

Note that although a market $k$ with $\Psi(q_k) = \Psi_{max}$ produces electoral outcomes equal to the outcomes obtained in the full-information majority-rule benchmark, the election outcomes are not necessarily optimal from a social perspective that incorporates the downstream utilities associated with $s$ and $r$. Our ranking concept here adopts a democratic governance point of view in which majoritarian outcomes are privileged. We take informed majoritarian outcomes as the benchmark, and explore which advisory market structures bring election outcomes closer to that standard. The model is silent on whether the policies adopted in those votes maximize shareholder value or achieve some other social goal such as cleaning the environment.

### D. *Majoritarian versus Utilitarian Welfare Comparisons*

This study focuses on the positive question of how competition in the proxy advice market affects the nature of advice and the outcomes of corporate elections. But we are also interested in the normative question of how proxy advice affects the welfare of market participants. Unfortunately, there is not a consensus on how to evaluate welfare in social choice situations.

One general approach is majority rule: given two outcomes, the one favored by the majority is preferred. 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). Our criterion of informed representation is essentially majority rule: we rank outcomes based on how likely they are to produce the majority voting outcome that would prevail under full information.

An alternative approach is classic utilitarianism: the outcome that maximizes the sum of utilities is preferred. This assumes that interpersonal utility comparisons are possible and implies that those with more intense preferences receive more weight in the social welfare function. Majority rule ignores preference intensity.

Both approaches have their appeal, and although we frame the paper around our majoritarian concept of informed representation, we also discuss the implications from a utilitarian perspective. One reason we emphasize majority rule is that it mirrors democratic practices: almost all democracies use some version of majority rule in which one person has one vote, independent of intensity of preferences. This holds for candidate elections, referendum elections, legislatures, and jury decisions. Shareholder voting gives one vote per share but does not adjust for intensity of preferences. Another advantage of majority rule is that it avoids the possibility of a small group of extremists dictating decisions to everyone else. For example, in a corporate election in which a small minority of shareholders have near infinite utility from abating carbon emissions because they believe that greenhouse gases pose an immediate existential risk for future generations of humanity, utility maximization would allow that group to dictate the decision even if a substantial majority of shareholders disagreed. Not taking into account preference intensity may seem like a defect of democracy from a utilitarian perspective, but it could be a design feature not a bug.

Giving welfare weight to preference intensity is particularly a concern in our model because we work with expressive preferences. Our model assumes that funds have "fundamental" preferences over $r$ and $s$, but on top of this, they receive utility from expressing those preferences given by $\lambda$. Maximizing utility in terms of the underlying utility from $r$ and $s$ has an appeal, but it's harder to justify why the expressive value $\lambda$ of those preferences should have a normative weight. If SRI funds have a high value from voting on green issues because that gains them fund flow, does that mean that it is socially optimal to give their opinions on, say, decarbonization more weight?

Having said this, we want to emphasize that we are not arguing that our concept is necessarily superior to utilitarianism or other alternatives. We think informed representation is a potentially useful lens through which to view election outcomes, but other approaches are useful too. By calling attention to the underlying complexity of welfare comparisons in corporate elections, we hope to stimulate further thinking on the issue, which is underexplored to date.

## IV. Monopoly

Consider a market with a single proxy advisory firm. We next characterize the policy $\alpha$ that maximizes the firm's profit. In this section we work with the general case of $\eta \in (0, 1]$ to illustrate that the model can accommodate it, before reverting to $\eta = 1$ in subsequent sections. If the firm offers a policy $\alpha$, then each fund $(\lambda, \theta)$ is willing to pay at most $U_{\lambda\theta}(\alpha)$ for its recommendations. The firm then charges each fund a price $\eta U_{\lambda\theta}(\alpha)$. Given the distribution $G$ of funds, the firm's profit from offering policy $\alpha$ is

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

Note that the monopolist maximizes profit by choosing the policy that maximizes the total expected expressive utility of the funds. This reason is intuitive: 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

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

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

$$\hat{g}(\theta) \;\equiv\; \frac{\int_0^1 \lambda g(\lambda, \theta) d\lambda}{E[\lambda]}. \tag{12}$$

Note that $\hat{g}(\theta) \geq 0$ and $\int_0^1 \hat{g}(\theta) d\theta = 1$. This new distribution assigns higher weight to values of $\theta$ associated with higher values of $\lambda$. The profit function can then be expressed as

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

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

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

By definition, $E[\lambda\theta] = E[\lambda]E[\theta] + cov_{\lambda\theta}$, which implies the following result:

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

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

*which produces Profit($\alpha^*$). This policy maximizes total expected expressive utility of funds. However, this policy maximizes informed representation if and only if*

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

To understand the intuition, consider first the case in which $cov_{\lambda\theta} = 0$, so that the advice firm selects the policy $\alpha^*$ that targets the average fund $E[\theta]$. This differs from most political economy models of elections, where candidates target the median voter. Here, the advice firm 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 deviating its policy $\alpha$ from the median fund's preferred policy to the average fund's preferred policy, the firm lowers the surplus of a majority of funds in order to create a larger surplus for a minority of funds with preferences further from the median. Figure 3 depicts an example in which the pdf of $\theta$ is right skewed and $\theta_m < E[\theta]$.

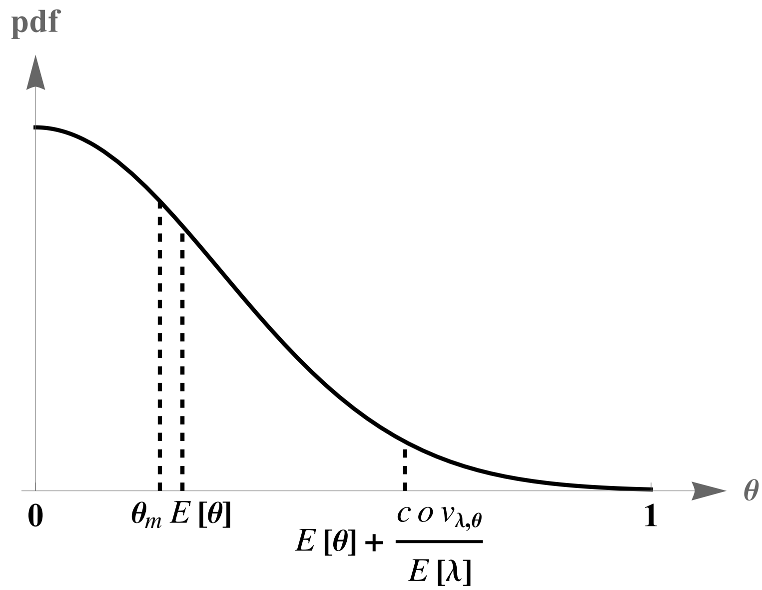


**Figure 3: Right-skewed pdf of $\theta$.** The median fund has preference $\theta_m$, but the monopolist chooses policy $\alpha^* = E[\theta]$ if preference parameters are independent, and $\alpha^* = E[\theta] + \frac{cov_{\lambda\theta}}{E[\lambda]}$ if $cov_{\lambda\theta} > 0$.

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.

In the case of $cov_{\lambda\theta} \neq 0$, because $cov_{\lambda\theta} = \rho_{\lambda\theta}\sigma_\lambda\sigma_\theta$, the profit-maximizing policy can be written as

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

In the example of Figure 3, a positive correlation between $\theta$ and $\lambda$ 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 $\alpha^*$ 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 ($\sigma_\theta$ is large), when there is greater heterogeneity in expressive benefits ($\sigma_\lambda$ is large), and when expressive benefits are low on average ($E[\lambda]$ is small).*

The corollary offers one way to interpret some features of actual shareholder activism today. Arguably, SRI funds are more concerned with their voting record while traditional funds focus primarily on financial returns (SRI funds post their voting records on their websites and use their voting record as part of their marketing strategy). In terms of the model, this is $\rho_{\lambda\theta} > 0$. Furthermore, there is variation across funds in the weights they place on social versus financial returns, and in their expressive benefits from voting. All of these features would lead a monopolist to choose an advice policy that is tilted toward the preferences of the SRI funds.

To characterize the information loss in such a market, note that all funds vote according to advice $\alpha^*$. 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 at all. Nevertheless, there is still some information loss unless $\alpha^* = \theta_m$, and the loss grows as $\alpha^*$ becomes more distant from $\theta_m$. 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 $\lambda$, 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. The timing of the game is the following. Both firms simultaneously choose advice policies $\alpha_1$ and $\alpha_2$. After observing each other's policies, the firms simultaneously offer a customized price to each fund, a form of Bertrand competition. From here on we assume that advisory firms make take-it-or-leave-it offers to funds. Once the firms have offered prices, funds choose whether to purchase advice, and if so, from which advisory firm. We assume that a fund can purchase advice from only one firm. We focus on subgame perfect equilibria with pure strategies, which we refer to it as an equilibrium.

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. We characterize the equilibrium by examining two cases. First, suppose that the firms were to offer the same policy $\alpha_1 = \alpha_2$. Bertrand competition would yield a price equal to the marginal cost, which is zero by assumption, resulting in a profit of zero. If instead the firms offer different policies, the product differentiation results in consumers paying a price above marginal cost, strictly increasing profits. Therefore, there is no equilibrium in which $\alpha_1 = \alpha_2$. Second, suppose that the firms offer different policies with $\alpha_1 < \alpha_2$. Note that 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. We can solve for the value of $\tilde{\theta}$:

$$\begin{aligned}
U_{\lambda\tilde{\theta}}(\alpha_1) &= U_{\lambda\tilde{\theta}}(\alpha_2) \\
\lambda\left\{\tilde{\theta}V_s(\alpha_1) + (1-\tilde{\theta})V_r(\alpha_1)\right\} &= \lambda\left\{\tilde{\theta}V_s(\alpha_2) + (1-\tilde{\theta})V_r(\alpha_2)\right\} \\
\tilde{\theta}\left[V_s(\alpha_1) - V_s(\alpha_2)\right] + (1-\tilde{\theta})\left[V_r(\alpha_1) - V_r(\alpha_2)\right] &= 0 \qquad (18)\\
\tilde{\theta} &= \frac{-\left[V_r(\alpha_1) - V_r(\alpha_2)\right]}{\left[V_s(\alpha_1) - V_s(\alpha_2)\right] - \left[V_r(\alpha_1) - V_r(\alpha_2)\right]}.
\end{aligned}$$

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 each fund to retain some of the surplus: if $U_{\lambda\theta}(\alpha_1) > U_{\lambda\theta}(\alpha_2)$, then fund $(\lambda, \theta)$ purchases advice from Firm 1, receives a gross surplus $U_{\lambda\theta}(\alpha_1)$, pays a price $U_{\lambda\theta}(\alpha_1) - U_{\lambda\theta}(\alpha_2)$, and receives a positive net surplus $U_{\lambda\theta}(\alpha_2)$. The case of $U_{\lambda\theta}(\alpha_1) < U_{\lambda\theta}(\alpha_2)$ is symmetric.

Having described the price and purchase decisions in the subgame after firms choose advice policies, it remains to characterize the policy choices themselves. Before doing that, it is useful to establish an intermediate result regarding the policies that maximize the total expressive voting utility. Because each fund purchases the advice policy that yields it the highest expressive voting utility, the total expressive voting utility in the market is

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

where $\hat{g}(\theta)$ is defined by (12). We wish to characterize the two policies that maximize $\mathcal{U}$. The solution is not identical policies $\alpha_1=\alpha_2$ because at that point $\mathcal{U}$ could be increased by changing one of the policies. Consider $\alpha_1<\alpha_2$. The next lemma shows that $\mathcal{U}$ is maximized by advice policies that maximize the average expressive utility of the funds that purchase that advice, where the average is computed using the weighted distribution $\hat{g}(\theta)$.

**Lemma 2.** *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*

$$\alpha_1^* = \hat{E}[\theta|\theta\le\tilde{\theta}] = E[\theta|\theta\le\tilde{\theta}] + \frac{cov_{\lambda\theta}(\theta\le\tilde{\theta})}{E[\lambda|\theta\le\tilde{\theta}]}, \tag{20}$$

$$\alpha_2^* = \hat{E}[\theta|\theta\ge\tilde{\theta}] = E[\theta|\theta\ge\tilde{\theta}] + \frac{cov_{\lambda\theta}(\theta\ge\tilde{\theta})}{E[\lambda|\theta\ge\tilde{\theta}]}, \tag{21}$$

$$\tilde{\theta} = \frac{-[V_r(\alpha_1^*)-V_r(\alpha_2^*)]}{[V_s(\alpha_1^*)-V_s(\alpha_2^*)]-[V_r(\alpha_1^*)-V_r(\alpha_2^*)]}. \tag{22}$$

With this established, we can present the main result of this section.

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

*(i) Advice policies $(\alpha_1^*,\alpha_2^*)$ form a competitive equilibrium if and only if*

$$\mathcal{U}(\alpha_1^*,\alpha_2^*)\ge\mathcal{U}(\alpha_1,\alpha_2^*) \quad \text{for all } \alpha_1\in[0,1], \text{and} \tag{23}$$

$$\mathcal{U}(\alpha_1^*,\alpha_2^*)\ge\mathcal{U}(\alpha_1^*,\alpha_2) \quad \text{for all } \alpha_2\in[0,1]. \tag{24}$$

*This implies that if the advice policies $(\alpha_1^*,\alpha_2^*)$ maximize total expressive voting utility $\mathcal{U}$, then they also form a competitive equilibrium. Therefore, a competitive equilibrium exists.*

*(ii) If advice policies $\alpha_1^*<\alpha_2^*$ form a competitive equilibrium, then conditions (20), (21), and (22) hold. Moreover, $\alpha_1^*<\alpha^*<\alpha_2^*$, where $\alpha^*$ is the optimal policy of a monopolist.*

*(iii) For $\tilde{\theta} \in [0,1]$, define the function:*

$$\Gamma(\tilde{\theta}) \equiv \tilde{\theta}\left[V_s(\hat{E}[\theta \leq \tilde{\theta}]) - V_s(\hat{E}[\theta \geq \tilde{\theta}])\right] + (1-\tilde{\theta})\left[V_r(\hat{E}[\theta \leq \tilde{\theta}]) - V_r(\hat{E}[\theta \geq \tilde{\theta}])\right]. \quad (25)$$

*If $\Gamma(\tilde{\theta})$ crosses zero only once, then the competitive equilibrium is unique and the equilibrium policies maximize the total expressive voting utility.*

To understand part (i), fix an advice policy $\alpha_2^*$. 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 $\alpha_1$ 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 preference parameter $\theta \in [0, \tilde{\theta}]$, hence its optimal advice policy is $\alpha_1^* = \hat{E}[\theta|\theta \leq \tilde{\theta}]$ from Lemma 2; 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 fund of 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 $\theta \geq \tilde{\theta}$).

For part (iii), we can rearrange (22) , which came from (18), and then substitute $\alpha_1^* = \hat{E}[\theta \leq \tilde{\theta}]$ and $\alpha_2^* = \hat{E}[\theta \geq \tilde{\theta}]$, yielding:

$$\tilde{\theta}\left[V_s(\hat{E}[\theta \leq \tilde{\theta}]) - V_s(\hat{E}[\theta \geq \tilde{\theta}])\right] + (1-\tilde{\theta})\left[V_r(\hat{E}[\theta \leq \tilde{\theta}]) - V_r(\hat{E}[\theta \geq \tilde{\theta}])\right] = 0. \quad (26)$$

Defining the left-hand side 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 $\Gamma(\tilde{\theta}) = 0$. The proof of the proposition shows that $\Gamma(0) > 0$ and $\Gamma(1) < 0$. Since $\Gamma$ is continuous, it crosses zero at least once. If it crosses zero only once, then the equilibrium is unique. An example would be if $s$ and $r$ are independently drawn from a standard normal distribution and $\hat{g}(\theta)$ is a uniform distribution, in which case $\Gamma$ is a strictly decreasing function. When the competitive equilibrium is not unique, there may be an equilibrium in which total voting utility is not maximized. Proposition 2 states that each firm chooses a policy that maximizes the total voting utility given the policy choice of its competitor. Therefore, both firms could be trapped at a local maximum that is not a global maximum.

Our results in Proposition 2 are closely related to the insights from Lederer and Hurter (1986). They consider a spatial competition model with price discrimination. There is a distribution of consumers on a plane, each consumer wants to purchase a single unit of a good, and they have the same willingness to pay. Firms simultaneously choose a location and then make price-discriminating offers to consumers. Firms face transportation costs that depend on their location and the location of their consumers. Lederer and Hurter (1986) shows that in a competitive equilibrium, given the locations of the other firms, each firm chooses a location that minimizes the total transportation cost. Therefore, the locations that minimize the overall transportation cost are a competitive equilibrium. In our model, instead of firms choosing spatial locations and facing transportation costs, we have firms choosing advice policies that endogenously generate different values to different consumers.

### *B. Competition and the Proxy Advice Market*

For the remainder of the paper, we assume that the competitive equilibrium is unique (hence it maximizes the total expressive voting utility).[12] An immediate consequence of Proposition 2 is 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 competition increases expressive voting utility.

An important question is how competition in the advice market affects the amount of informed representation in corporate elections. Are outcomes better informed and more representative with a duopoly or a monopoly? It turns out that both cases are possible. Intuitively, entry can be harmful in a market in which the monopolist offers a policy very close to the preferred policy of the median fund. This is because duopolists would differentiate their advice policies, offering $\alpha_1^*$ and $\alpha_2^*$ that are farther from $\theta_m$. Figure 4a provides an example. The median voter is then forced to choose either $\alpha_1^*$ or $\alpha_2^*$, which are farther from $\theta_m$. Figure 4b shows the other case: with competition, the median fund can choose policy

[12] Alternatively, we could assume that the firms coordinate in the equilibrium that maximizes total expressive voting utility.

$\alpha_1^*$, which is better (closer to $\theta_m$) than the policy $\alpha^*$ that would be offered by a monopolist. In this case, competition increases informed representation.

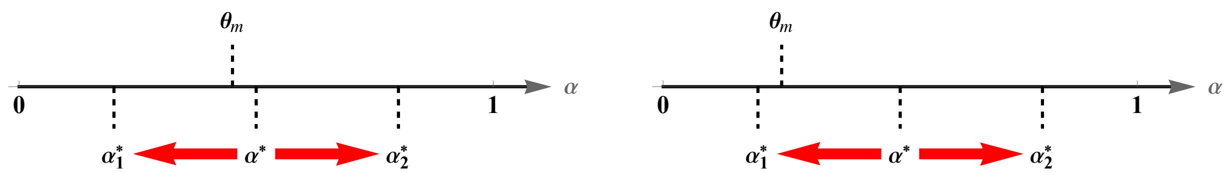


**(a)** Competition Decreases Informed Representation

**(b)** Competition Increases Informed Representation

**Figure 4: Comparison of policy choices by monopolist and duopolists.** The monopolist's policy is $\alpha^*$; the duopolists' policies are $\alpha_1^*$ and $\alpha_2^*$; and the median fund's preference is $\theta_m$.

Formally, fix a market and consider two scenarios: a monopoly and a duopoly. Let $\alpha^*$ be the equilibrium policy of the monopolist, with corresponding voting outcome $q^*$ and informed representation $\Psi(q^*)$ described by (17). Let $\{\alpha_1^*, \alpha_2^*\}$ be the equilibrium policies of the duopoly market, with $\alpha_1^* < \alpha_2^*$. In the duopoly case, let $\tilde{\theta}$ be the fund just indifferent between purchasing policies $\alpha_1^*$ and $\alpha_2^*$, defined in (22). If $\theta_m \leq \tilde{\theta}$, then the median fund and all funds $\theta < \theta_m$ purchase advice from Firm 1, forming a simple majority of funds. Therefore, the market voting outcome coincides with the recommendations of Firm 1. Similarly, if $\theta_m > \tilde{\theta}$, then the median fund and all funds $\theta \geq \theta_m$ purchase advice from Firm 2, and the market voting outcome coincides with the recommendations of Firm 2. In equilibrium, the advising firm hired by the median fund dictates the market outcome. Define $\alpha_m^*$ as the policy of the median-supported advisory firm, and $q_m^*$ the resulting voting outcome.

**Corollary 2.** *Informed representation is higher under duopoly than monopoly if and only if $\Psi(q^*) \leq \Psi(q_m^*)$. Informed representation is higher under duopoly than monopoly if $\theta_m$ is sufficiently far from $\alpha^*$, and lower if $\theta_m$ is sufficiently close to $\alpha^*$.*

The corollary follows from the fact that $\alpha_1^* < \alpha^* < \alpha_2^*$ (see Proposition 2), which is driven by the advisory firms' product differentiation. If the median fund's preference $\theta_m$

is sufficiently close to the monopolist's policy $\alpha^*$, then the median fund strictly prefers the monopolist's policy $\alpha^*$ over both of the duopoly policies $\alpha_1^*$ and $\alpha_2^*$, and competition reduces informed representation. Conversely, if $\theta_m$ is sufficiently far from $\alpha^*$, 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.

Our focus is primarily on the connection between market structure and the slant of advice, but the model also has implications for prices. Intuitively, moving from monopoly to duopoly has two potentially offsetting effects on prices. 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 $\alpha^*$, 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 the monopolist's advice policy, they may experience higher or lower prices depending on parameters.

Note that whenever $\theta_m$ is sufficiently close to $\alpha^*$, the median fund is unable to make a better decision in a duopoly even if it could observe the recommendation of both firms. That is, even if we consider the joint information implied by $\alpha_1^*$ and $\alpha_2^*$ together the duopoly firms provide less useful information than the monopolist. To see this, consider the example in Figure 5, where proposals are distributed on a disk. A monopolist would recommend approval of all policies to the northeast of the dashed line $\alpha^*$, and rejection of proposals below the line. In a duopoly, each firm would give recommendations according to $\alpha_1^*$ and $\alpha_2^*$. All three policies recommend approval of proposals in the blue-shaded northeast sector, and rejection of proposals in the red-shaded southwest sector. Duopoly and monopoly produce the same outcomes in these sectors. The difference appears in the other sectors. In the northwest sector between $\alpha_1^*$ and $\alpha_2^*$, Firm 1 recommends rejection and Firm 2 recommends approval, while the opposite is true in the southeast sector between $\alpha_1^*$ and $\alpha_2^*$.

Suppose the median fund were able to simultaneously observe the recommendations of both firms. Whenever Firm 1 recommends rejection and Firm 2 recommends approval, the median fund learns that the proposal must be in the northwest sector between $\alpha_1^*$ and $\alpha_2^*$ and updates its belief accordingly. The fund would have to make its voting decision based on this posterior belief about the expected (positive) social return $s$ and the expected

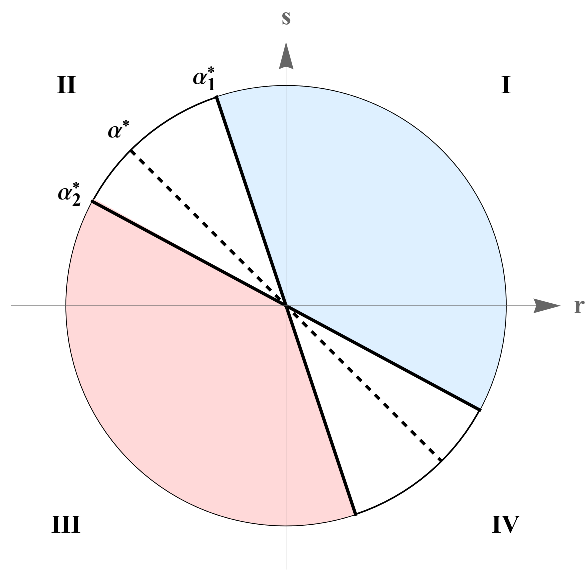


**Figure 5: Comparison of policy choices of monopolist and duopolists.** This example shows the monopolist's policy $\alpha^*$ and the duopolists' policies $\alpha_1^*$ and $\alpha_2^*$, when proposals are distributed on a disk.

(negative) financial return $r$ is in this area. With this limited information, the fund would either approve all proposals in this area or reject all proposals in this area. In contrast, the monopolist's advice tells the fund precisely whether the proposal is above or below the dashed line. Hence, if the fund's preference $\theta_m$ is sufficiently close to $\alpha^*$, the information provided by the monopolist is strictly more valuable than the joint information provided by the duopolists. Similar logic applies to the remaining sector.

We conclude this section with what could be seen as a paradox. Consider the policies represented in Figure 4a. Suppose that the distribution of $\theta$ is such that most of the funds are concentrated around the two extreme points $\{0, 1\}$. When we move from the monopolist $\alpha^*$ to the duopoly $\{\alpha_1^*, \alpha_2^*\}$, the vast majority of funds are able to receive advice from a firm 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 are making voter choices more aligned with their preferences, as a group the funds are making a worse collective decision, in the sense that informed representation is deteriorating: 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 might help an 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 desire to maximize profit, which can lead to a detrimental product differentiation.

## VI. Many Advisory Firms

We can generalize Proposition 2 to competition between $N$ firms. To this end, we extend the duopoly model by supposing there is a finite number $N \geq 2$ of firms. The timing of the game is the same as before. All firms simultaneously choose advice policies $\{\alpha_1, \cdots, \alpha_N\}$. After observing the chosen policies, the firms simultaneously offer a take-it-or-leave-it price to each fund. Funds then choose which advice to purchase, if any. We continue to assume that each fund can purchase advice from at most one firm. We focus on subgame perfect equilibria in which players use pure strategies, henceforth, simply called an equilibrium. If there are multiple equilibria, then we select the equilibrium that maximizes the total expressive voting utility (we will prove that such equilibrium exists).

In equilibrium, as before, two firms never offer the same policy, and for the same reason: if they did, each firm could strictly increase its (zero) profit by choosing a policy different from the other firms. Therefore, the firms segment the market. Consider the policy profile in which firms are ordered by their policies, $\alpha_1 < \cdots < \alpha_N$. Between each pair of adjacent firms $\alpha_k$ and $\alpha_{k+1}$, there is a cutoff customer $\tilde{\theta}_{k,k+1} \in (\alpha_k, \alpha_{k+1})$ that is indifferent between the advice provided by the two firms,

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

and define $\tilde{\theta}_{0,1} = 0$ and $\tilde{\theta}_{N,N+1} = 1$. Each firm $k \in \{1, \cdots ., N\}$ serves the segment of funds with preferences $\theta \in [\tilde{\theta}_{k-1,k}, \tilde{\theta}_{k,k+1}]$. As with duopoly, the first-order condition for all firms $k$ requires the optimal policy $\alpha_k$ to target the average consumer in the market segment, which implies:

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

The total expressive utility in the market, expressed in equation (19) in the duopoly case, here becomes:

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

Because the policy vector belongs to a compact set and $\mathcal{U}$ is continuous, there exists a vector of policies that maximizes $\mathcal{U}$. We use the notation $\mathcal{U}(\alpha_k|\alpha_{-k})$ to represent the total utility as a function of the policy $\alpha_k$, given a fixed vector $\alpha_{-k}$ of policies of the other firms. With this established, we can present the general result.

**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 each firm $k$,*

$$\mathcal{U}(\alpha_k^*|\alpha_{-k}^*) \geq \mathcal{U}(\alpha_k|\alpha_{-k}^*) \qquad \text{for all } \alpha_k \in [0, 1]. \quad (29)$$

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

*(ii) If advice policies $\alpha_1^* < \ldots < \alpha_N^*$ form a competitive equilibrium, then conditions (27) and (28) hold.*

*(iii) Let $\alpha_m^*$ be the policy of the advisory firm that sells to the median voter.*[13] *Then the advisory firm $\alpha_m^*$ is decisive for election outcomes. That is, in equilibrium, each proposal is approved if and only if it receives an approval recommendation by the advisory firm with policy $\alpha_m^*$.*

Figure 6 illustrates a market with four advisory firms. In this example, the partitions have the same size for simplicity. The actual size of each market segment depends on the distributions $F$ and $G$.

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

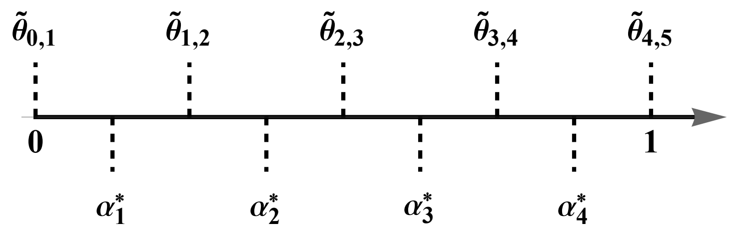


**Figure 6: Example of a market with four advisory firms.** Firm $k$ serves consumers with preferences $\theta \in [\tilde{\theta}_{k-1,k}, \tilde{\theta}_{k,k+1}]$ and offers policy $\alpha_k^* = \hat{E}[\theta | \tilde{\theta}_{k-1,k} \leq \theta \leq \tilde{\theta}_{k,k+1}]$.

From a utilitarian point of view, increasing the number of firms in the market monotonically increases the total expressive voting utility. Because funds' willingness to pay for advice is derived from their expressive voting utility, firms customize their policies to maximize total voting utility. More firms means more policies to choose from, and therefore higher expected utility in aggregate.

One remaining question of interest is how shareholder democracy performs as the number of firms increases. In the limit, we reach the "perfectly competitive" outcome: a continuum of advisory firms, with each firm $k$ serving one fund type $\theta_i$, providing this type's favorite advice $\alpha_k^* = \theta_i$. Perfect competition maximizes informative representation, our metric of the effectiveness shareholder democracy:

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

Together Corollaries 2 and 3 paint a nuanced picture of the effects of competition in the proxy advice market on shareholder democracy. Marginal increases in competition (such as moving from one to two firms) are sometimes detrimental, but perfect competition maximizes shareholder democracy.

## VII. Customized Advice and Entry

### *A. Customized Advice*

One of our main findings is that unless there are many advisory firms, outcomes of corporate elections may not be informationally representative, and may be slanted toward the preferences of SRI funds. A natural question is to what extent this unfortunate outcome could be mitigated 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.[14] Here we extend the model to allow customization and explore its implications for shareholder democracy.

Consider the following extension of our model. There is a finite number $N$ of advisory firms and each firm $k$ offers a finite number of policies, denoted by $m_k$. As before, we continue to assume that the number of firms and the number of policies per firm are fixed exogenously. Let $M \equiv \sum_{k=1}^{N} m_k$ be the number of policies in the market; let $\alpha^k = (\alpha_1^k, \cdots, \alpha_{m_k}^k)$ be the menu (vector) of policies offered by firm $k$, where $\alpha_p^k$ is a generic policy in this set; and let $\alpha^{-k}$ be the menu of policies offered by all firms not including $k$. All firms simultaneously select their menu of policies. After observing the set of all policies in the market, each firm quotes a price for each of its policies to each fund. Each fund then chooses which policy to purchase. We continue to assume that a fund can only purchase a single policy.[15]

We start with some intuition about equilibrium prices. When firm $k$ considers which prices to offer a fund, the only relevant policy in the set $\alpha^k$ is the policy $\alpha_p^k$ that maximizes the fund's utility. This is because a firm can never earn more profit by selling a given fund a policy that the fund prefers less than another policy offered by the firm. Therefore, when

[14] In a survey, 30 percent of investors stated proxy advisors' advice is too standardized (McCahery et al., 2016).

[15] Our model of customized advice differs from Hu et al. (2024) in that their model (see their Appendix C) does not explicitly consider competition between advisory firms; each investor can pay an exogenously defined price to receive benchmark (non-customized) voting recommendations, or pay an exogenously defined additional price to receive perfectly customized voting recommendations. In contrast, our model considers profit maximizing advisory firms that strategically choose prices and are limited on the degree of customization that they can offer – each firm $k$ competes by selecting $m_k$ different types of policies to offer.

computing the equilibrium, we can treat each firm as if it has only one policy for each fund as before.

Formally, fix the set of policies offered and consider the price that firm $k$ offers to fund $(\lambda, \theta)$. Let

$$\tilde{U}_{\lambda\theta}(\alpha^k) \equiv \max_{\alpha_p^k \in \alpha^k} U_{\lambda\theta}(\alpha_p^k)$$

be the maximum voting utility that firm $k$ can provide to fund $(\lambda, \theta)$, and let

$$\tilde{U}_{\lambda\theta}(\alpha^{-k}) \equiv \max_{\alpha_p^{-k} \in \alpha^{-k}} U_{\lambda\theta}(\alpha_p^{-k})$$

be the maximum voting utility that competitors can provide this fund. If $\tilde{U}_{\lambda\theta}(\alpha^k) < \tilde{U}_{\lambda\theta}(\alpha^{-k})$, then firm k offers a price equal to marginal cost (which is zero) and loses the fund to a competitor. If $\tilde{U}_{\lambda\theta}(\alpha^k) > \tilde{U}_{\lambda\theta}(\alpha^{-k})$, then firm k wins the price competition with an offer price of $\tilde{U}_{\lambda\theta}(\alpha^k) - \tilde{U}_{\lambda\theta}(\alpha^{-k})$.

Our model allows firm $k$ to offer fund $(\lambda, \theta)$ a different price for each policy in its menu, but it does not require this to happen. If firm $k$ offers an identical (flat) price $\tilde{U}_{\lambda\theta}(\alpha^k) - \tilde{U}_{\lambda\theta}(\alpha^{-k})$ for each policy it offers to fund $(\lambda, \theta)$, and allows the fund to choose from the menu of offerings, the fund always selects the policy that yields the highest utility,

It is straightforward to restate Proposition 3 in terms of these prices. Given the policies $\alpha^{-k}$ offered by the other firms, firm $k$ chooses the policy menu $\alpha^k$ that maximizes the total expressive voting utility because this maximizes the extractable surplus. If the vector of policies $\boldsymbol{\alpha^*} = (\alpha_1^*, \cdots, \alpha_M^*)$ maximizes total expected voting utility, then there is an equilibrium in which the $M$ policies offered by the firms equal $\boldsymbol{\alpha^*}$. The first-order conditions continue to apply: each equilibrium policy $\alpha_p^k$ 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 $\alpha_p^k$ purchased by the median fund is decisive for election outcomes. Because Proposition 3 can be restated in this way, it implies that the equilibrium implications and key insights remain the same under this extension.

There is a close connection between the electoral outcomes of the market studied in the previous sections, with $N$ firms offering one policy each ("single-policy market"), and the market studied in this section with $N'$ firms offering multiple policies that sum to $M = N$ total policies ("multiple-policy market"). If the policy vector $\boldsymbol{\alpha^*} = (\alpha_1^*, \cdots, \alpha_M^*)$ maximizes the total expressive voting utility, then both types of market have an equilibrium where these are the policies offered, and the resulting electoral outcomes are the same.

However, all funds pay a weakly lower price in the single-policy market than the multiple-policy market. 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.

Also, the prices in a multiple-policy market depend on how the policies are distributed across the advisory firms. For example, funds pay more if adjacent policies are offered by the same firm, and they pay less if adjacent policies are offered by different firms. To see this, consider a market with two firms in which each firm offers two policies, and the policies $(0.2, 0.4, 0.6, 0.8)$.

**Example 1 (adjacent policies):** Suppose Firm 1 offers policies $\alpha^1 = (0.2, 0.4)$ and Firm 2 offers $\alpha^2 = (0.6, 0.8)$. Consider a fund with a preference $\theta$ very close to 0.2, so that this policy yields the highest utility. Firm 1 then charges a price $U_{\lambda\theta}(0.2) - U_{\lambda\theta}(0.6)$ and allows the fund to choose a policy in its menu, which will be $\alpha = 0.2$. Note that Firm 1 does not need to worry about the fund switching from policy $\alpha = 0.2$ to $\alpha = 0.4$ because Firm 1 is charging a price for $\alpha = 0.4$ that is high relative to the utility it delivers. Firm 1 only needs to worry about the fund switching over to the best policy offered by the competition, which is $\alpha = 0.6$ in this case.

**Example 2 (nonadjacent policies):** Suppose Firm 1 offers policies $\alpha^1 = (0.2, 06)$ and Firm 2 offers $\alpha^2 = (0.4, 0.8)$. Again consider a fund with a preference $\theta$ very close to 0.2, so that this policy yields the highest utility. Firm 1 then charges a price $U_{\lambda\theta}(0.2) - U_{\lambda\theta}(0.4)$ and allows the fund to choose a policy in its menu, which will be $\alpha = 0.2$. Now Firm 1 needs to worry about the fund switching to policy $\alpha = 0.4$, which is the best policy offered by the competitor. For any fund $\theta$ close to 0.2, policy $\alpha = 0.4$ yields a higher utility than policy $\alpha = 0.6$. Therefore, Firm 1 charges a lower price to fund $\theta$ in Example 2 than in Example 1.

More generally, suppose that policy $\alpha_p^k$ is the best policy for some fund $\theta$. Then the second-best policy for $\theta$ is either the policy immediately to the left or to the right of $\alpha_p^k$. If both these policies are controlled by a competitor (the nonadjacent policies case), then firm $k$ faces a stronger price competition. If one or both these adjacent policies are controlled by firm $k$ (the adjacent case), then the firm faces a weaker price competition.

It follows that all funds retain more of their surpluses in Example 2 than Example 1, and the advisory firms earn a higher total profit in Example 1 than Example 2. Therefore, it matters for prices and profits whether firms have bunched policies that partition the market (as in Example 1) or alternating policies (as in Example 2).

Note that holding constant the number of policies, electoral outcomes are the same, whether offered by single-policy firms or multiple policy firms. And if offered by multiple-policy firms, the electoral outcomes are the same regardless of how the policies are distributed across the firms. This implies that informed representation is no better or worse if firms offer customized advice – what matters is not the menu of advice offered by individual firms, but the total array of advice available in the market.

### *B. Entry*

One of the key findings of our analysis is that corporate elections are not reliably informationally representative, and specifically, that their outcomes are skewed toward the preferences of funds with high expressive benefits from voting, perhaps such as SRI funds. However, as shown above, as the number of advisory firms or the number of customized policies goes to infinity, the slant in election outcomes goes to zero. This raises the question whether entry can resolve the problems associated with voting advice.

This boils down to a question of whether there is sufficiently entry to remove the distortion. We do not explicitly analyze entry in our model, but one could imagine an extension in which entry is allowed and firms have a fixed cost. This would produce an equilibrium number of firms, with the number inversely related to the fixed cost. In practice, the American advice market is essentially a duopoly consisting of ISS and Glass Lewis. Because there are no legal or obvious institutional barriers to entry, this suggests that the fixed costs are substantial. One significant source of fixed costs is investment in proxy execution services. Both ISS and Glass Lewis typically sell their advice services bundled together with vote execution services. For a fund that has to execute tens of thousands of votes each year, the value of vote execution services is likely to be immense, and could easily exceed the value of advice itself. It is possible the proxy advice market has not yet reached its long run equilibrium and new firms will enter – the industry is not that old – but its structure appears to have been fairly stable going back at least to 2010 (Shu 2024) so there is currently no reason to expect that new entry will significantly reduce the potential slant.

A more plausible way that slant may be reduced is through "entry" in the form of customized advice: the industry may remain a duopoly but the firms may proliferate the number of advice products they offer. Theoretically, we could also extend the model to allow fixed costs in providing customized services. The existence of such fixed costs is plausible since an advisory firm must set up a distinct "policy experiment" for each advice policy. Practically speaking, this means designing rubrics and hiring and training researchers to detect information that is pertinent to the specific form of customization (for example, religious groups may desire the recommendations they receive to take into account the ethical behavior of corporations and the social impact of their actions).[16] The effect of these considerations in practice is ultimately an empirical question, but the amount of customization in the market today suggests that we are not yet at the perfectly competitive equilibrium.

### *C. Multiple Purchases*

Another way funds might reduce distortion in election outcomes is by purchasing advice from more than one advisory firm. Our assumption that a fund purchases only a single advice policy simplifies the characterization of optimal policies and election outcomes, and allows us to work with very general distributions $F$ and $G$. Here we extend the model to allow funds to purchase more than one advice policy. Suppose that firms offer policies $\{\alpha_1, \ldots, \alpha_N\}$, and consider a fund with preference $\theta$. What is the value of obtaining a second advice policy?

**Remark 4.** *If fund $\theta$ 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 $\theta$*

[16] If an advisory firm conducted research that revealed $r$ and $s$ then it would be conceptually easy to offer additional advice at zero marginal cost, but in our model the firm only acquires information on which side of the policy dividing line the proposal lies.

*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.

Trivially, the remark implies that if fund $\theta$ receives advice from its ideal policy $\alpha = \theta$, then the additional value of any other advice policy is zero. It also follows that if fund $\theta$ 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 $\theta$. Therefore, a fund either purchases one advice policy (the advice that yields the highest surplus $U_{\lambda\theta}(\alpha)$), or it purchases two advice policies: the highest available $\alpha < \theta$ and lowest available $\alpha > \theta$.

This brings us back to the situation we discussed in Section V.B, regarding access to two signals. 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.

## VIII. 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 sufficiently large funds conduct their own research on election items, and some funds may acquire election-related information as a by-product of conducting investment research. This raises the interesting question of how the emergence of firms selling advice affects the incentives of firms to vote based on their own internal information: does it "crowd out" the use of internal information – or does it increase the value of internal information through complementarities? And how does this affect the amount of information underlying corporate elections, and thus the effectiveness of shareholder democracy?

Here we extend the model to include firms with internal information, and explore the impact of an advice market. Malenko and Malenko (2019) develop a model in which funds have identical preferences, and therefore corporate elections are common value decisions; in their context, a market for advice crowds out internal provision of information, and results in less informative elections. Our analysis abstracts away from these forces by assuming that funds are exogenously endowed with internal information; this allows us to focus on a different force that has not been previously examined: how existence of an advice market can shift the position of the median voter, and possibly reduce the informed representation of elections.

### *A. Internal Production of Voting Information*

In this extension each fund is characterized by three parameters, $(\lambda, \theta, \gamma)$. The new parameter $\gamma \in [0, 1]$ represents the fund's internally produced information. Specifically, fund $i$ with internal information $\gamma$ exogenously learns the value $v_{ij}$ of a measure $\gamma$ of proposals (recall that there is a unit mass of proposals). Funds are distributed according to $G(\lambda, \theta, \gamma)$.

In our benchmark model, because unadvised firms have no information ($\gamma = 0$), funds that do not purchase advice have a payoff of zero. In the extension here, self-informed funds ($\gamma > 0$) have a strictly positive payoff if they do not purchase advice. This limits the amount of surplus that advisory firms can extracted from them.

We assume that each fund is independently informed about a random subset of proposals that is uncorrelated with fund characteristics and proposal characteristics. This is a potentially substantive assumption, as we discuss below, but we maintain it here to isolate the key issues. For the mass $\gamma$ of proposals about which a fund is fully informed, it votes yes only if a proposal delivers a striclty positive payoff, yielding an average payoff of $\lambda[\theta V_s(\theta) + (1-\theta)V_r(\theta)]$. For the remaining mass $1-\gamma$ of proposals, the fund is uninformed unless it purchases advice, and therefore receives a zero payoff. The expected utility of a fund that does not purchase advice is then:

$$\gamma\lambda \left[\theta V_s(\theta) + (1-\theta)V_r(\theta)\right]. \tag{30}$$

If the fund purchases advice with policy $\alpha$, then for the measure $1-\gamma$ of proposals without internal information, the fund votes according to the recommendation given by policy $\alpha$. (Advice does not affect its vote on proposals for which it has internal information, since internal information is presumed to be fully accurate.) The expected utility of a fund that purchases advice $\alpha$ (gross of its price) is then:

$$\gamma\lambda\left[\theta V_s(\theta)+(1-\theta)V_r(\theta)\right]+(1-\gamma)\lambda\left[\theta V_s(\alpha)+(1-\theta)V_r(\alpha)\right]. \tag{31}$$

The fund's willingness to pay for advice $\alpha$ is the difference between (31) and (30):

$$\bar{U}_{\lambda\theta\gamma}(\alpha) \;=\; (1-\gamma)\lambda\left[\theta V_s(\alpha)+(1-\theta)V_r(\alpha)\right]=(1-\gamma)U_{\lambda\theta}(\alpha).$$

Therefore, the reservation price of a fund with preference $(\lambda,\theta)$ and information $\gamma$ equals the original reservation price $U_{\lambda\theta}(\alpha)$ of a fund that is not self-informed times $(1-\gamma)$. Internal information reduces the fund's willingness to pay for advice, but does not change how the fund ranks different recommendation policies. From the viewpoint of the advisory firms, a market in which funds have internal information $\gamma$ and preference intensity $\lambda$ is isomorphic to a market in which funds have no internal information but a lower preference intensity $\tilde{\lambda}\equiv(1-\gamma)\lambda$. We can use this observation to restate the original distribution $G(\lambda,\theta,\gamma)$ as a new distribution $\tilde{G}(\tilde{\lambda},\theta)$. Because, as this shows, profit functions and first order conditions are isomorphic under the new and extended model, our results characterizing optimal policies $\alpha_k^*$ continue to hold as before.

### B. *Shareholder Democracy with Internal Information Production*

In this section we explain how the presence of internal information production changes how funds vote, which proposals are approved, and the level of informed representation. Shareholder democracy no longer depends only on the policies offered by advisory firms, but also on the amount of internally produced information and the identity of the funds that produce the information.

In order to study this case, we need to close the model with an assumption about whether a fund votes if it has no internal information and no advice. This case does not arise in the benchmark model because all funds acquire information and therefore vote. Here we assume that a fund with neither internal nor external information abstains from voting (we can imagine there is a small cost of voting, as below). This assumption is not critical: our results would hold if instead we assumed that uninformed funds flipped a coin when deciding how to vote.

We also assume, primarily to simplify the analysis, that if a fund with internal information $\gamma$ learns the value of a proposal $j$, then all funds with internal information $\gamma' \geq \gamma$ also learn about the value of proposal $j$. This is equivalent to saying that each proposal $j$ has a third dimension $z \sim U[0,1]$ that might be called "visibility." For each proposal $(s, r, z)$, all funds with internal information $\gamma \geq (1-z)$ are able to compute their value $v_{ij}$ of the proposal, while other funds are not. Therefore, the set of proposals with values known by a fund with a lower $\gamma$ is a subset of the proposals with values known by a fund with a higher $\gamma$. Funds are then ordered by $\gamma$ on what they know about the proposals. This is similar in spirit to the idea behind Garicano (2000), where workers are ordered by their ability to solve problems of different levels of complexity. Here, funds are ordered by the set of proposals that they are able to value. Proposals with a low visibility $z$ can only be valued by the funds with a high $\gamma$, while proposals with a high visibility $z$ can be valued even by funds with a low $\gamma$.

It is useful to start by considering a limit case.

**Remark 5.** *Suppose that $\gamma$ is distributed independently of $\lambda$, $\theta$, and $\tau$, and suppose that there are some funds with perfect knowledge, $g(\lambda, \theta, \gamma = 1) > 0$. Then a market without advisory firms achieves the highest level of informed representation, $\Psi_{max}$.*

The conclusion follows because, for each proposal $j$, there is a non-empty and unbiased group of funds that are able to value the proposal and cast votes that perfectly reflect their preferences. Because this group is unbiased, the median fund in each group is the same as the median fund $\theta_m$ in the general population. Therefore, the voting outcome of each small subgroup is representative of the larger population. This implies that existence of an advice market has the potential to improve informed representation only if funds have limited internal information (no fund has $\gamma = 1$) or the distribution of internal information across funds is biased away from the median fund $\theta_m$.

To see how introduction of an advice market can reduce informed representation, 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 this subset, and let $\theta_m \in (0,1)$ be the median fund in the population as before. Without an advice market, for a mass $\gamma_H$ of proposals, the informed funds know their own $v_{ij}$ and vote according to their preferences, and the uniformed funds abstain. The electoral outcome coincides with the vote of the subgroup's median fund $\theta_m^H$. 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

$$\gamma_H \left[\theta_m V_s(\theta_m^H) + (1-\theta_m) V_r(\theta_m^H)\right]. \tag{32}$$

If a monopolist advisory firm is added to this market, it chooses

$$\alpha^* = E[\theta] + \frac{cov_{(1-\gamma)\lambda,\theta}}{E[(1-\gamma)\lambda]}$$

as before. Because uninformed funds are a majority and all follow the monopolist's advice, the voting 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

$$\theta_m V_s(\alpha^*) + (1-\theta_m) V_r(\alpha^*) \quad \geq \quad \gamma_H \left[\theta_m V_s(\theta_m^H) + (1-\theta_m) V_r(\theta_m^H)\right]. \tag{33}$$

Inequality (33) allows us to study the 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 $\gamma_H$. Even if the subgroup of informed voters have a preference $\theta_m^H$ very close to the overall median $\theta_m$, and the monopolist offers a policy $\alpha^*$ 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 $\gamma_H$ of proposals, but improves the outcome for the remaining proposals. Conversely, if $\theta_M^H$ is better for the median voter than $\alpha^*$, then the monopolist strictly reduces informed representation if $\gamma_H$ is sufficiently high.

The intuition is similar if more than one advisory firm is added to the market. The advice market provides valuable new information to uninformed funds, but if the uninformed funds are sufficiently numerous, advice-based voting swamps voting based on internal information, and makes advisory firms the kingmakers. If a sufficiently large number of advisory firms is added, enough to make the market perfectly competitive, then the informed representation would be maximized.

The potential distortion arising from existence of an advice market in our model is different than the traditional crowding out effect. In the standard crowding out logic, the presence of advisory firms causes funds to acquire less internal information, or to disregard their internal information. Consequently, each fund becomes less informed or votes based on less information. That is not the case here, where the amount of internal information has been fixed by assumption. When there is no advice market and voting is entirely based on internal information, elections are determined by the median of funds that are self-informed. 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. Whether that leads to more or less informed representation depends on how the internally produced information is correlated with fund preferences. If, unlike the recommendations of advisory firms, internal information is uncorrelated with expressive voting benefits, then a market with no advisory firms may be less slanted toward the preferences of funds with high expressive benefits than one with advisory firms.

Empirically, a fund that purchases advice $\alpha$ will have a voting record that is perfectly aligned with the advisor firm's recommendation – sometimes called a "robo-voter" – if the fund has no internal information or if its preference $\theta$ is exactly the same as the policy $\alpha$. In fact, a substantial number of funds do appear to "robo-vote." Defining a robo-voter as a fund that followed its advisor's recommendation on 99 percent of its votes, Matsusaka and Shu (2024), using methods developed in Shu (2024), find that 35 percent of ISS customers and 33 percent of Glass Lewis customers were robo-voters in 2021. Even AQR, one of the world's largest hedge funds, appears to robo-vote. At the same time, many funds that purchase advice do not robo-vote. In the model, a fund casts some votes that differ from its advisor's recommendations if the fund has internal information ($\gamma > 0$) and its preference differs from the advising policy ($\theta \neq \alpha$). Our model predicts that the percentage of dissenting votes is increasing in the preference misalignment (difference between $\theta$ and $\alpha$) and on the amount of internal information $\gamma$.

### *C. 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 paying vote execution costs exceeds their private benefit from voting. To counteract this, government regulations force funds to vote: the Department of Labor ruled in 1988 that pensions had a fiduciary duty to vote under ERISA, and the SEC issued a no-action letter in 2003 stating that mutual funds had a fiduciary duty to vote. Moreover, if they voted based on advice from a proxy advisory firm, they enjoyed a presumption that their votes were free of conflicts of interest. These rulings led to a significant increase in voting by institutional investors and reliance on proxy advisory services.[17] In this section, we extend the model to consider the effect of a regulatory requirement to vote on shareholder democracy.

We now suppose that the cost of voting is strictly positive — each fund must pay a cost $FC > 0$ to maintain a voting system that allows it to cast all its votes. In this case, if funds are not required to vote, then funds that do not care enough about their voting records (funds with a sufficiently low $\lambda$) choose not to purchase proxy advice and abstain from voting. We continue to assume that some funds have internal information, captured by parameter $\gamma$.

If there is no regulation concerning voting, we show above that advisory firms offer policies that target funds with high $\tilde{\lambda} \equiv (1-\gamma)\lambda$, that is, funds that care a lot about their voting records (high $\lambda$) and have low information (low $\gamma$). Election outcomes are then determined by a combination of three forces: the internal information available to informed funds that vote, the slant of advice provided by advisory firms, and the number of funds that abstain from voting.

In this environment, a regulatory mandate to vote can improve informed representation,

[17]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.* SEC: *Proxy Voting by Investment Advisors, 68 Federal Register 6585, February 7, 2003.* Investment advisors, strictly speaking, are not required to vote their proxies and enforcement actions are rare, yet 90 percent of them choose to do so (SEC Staff Legal Bulletin No. 20, June 30, 2014; Broadridge + PwC, 2019). It is widely believed that funds can satisfy their fiduciary responsibilities by following the recommendations of a proxy advisor.

but can also harm it. To see how mandatory voting can be harmful, 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 $\lambda$ that do not purchase advice and abstain from voting. In this case, without regulation, the election outcomes are very close to the overall median preference and informed representation is high. Although relatively few funds vote, they have enough internal information, and their preferences are representative of the overall population.

If funds are required to vote, and their votes must be based on some information, the large group of funds with a low $\lambda$ that previously abstained now purchase advice from an advisory firm and vote according to this advice. Since this group is formed by funds with a low $\lambda$, the advisory firms do not change their policies much (recall that in order to maximize profit, their policies are targeted at funds with high $\lambda$). 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 overweights the preferences of funds with high $\lambda$, forcing funds to vote might decrease informed representation.

Intuitively, when voting is not mandatory and the number of funds that vote based on external advice is small, the election outcome is determined by a group of informed funds with a median preference similar to the population median preference. It is as if the economy delegates the election outcome to this representative group. When voting is mandatory and the number of funds using external advice is large, then it is as if the economy delegates the election outcome to the advisory firms, which may not be representative of the median voter.

## IX. Discussion

This paper develops a model to study the economics of the proxy advice market, with a special focus on understanding how the growing reliance on proxy advice affects the nature of outcomes in corporate elections. Our analysis attempts to bring to the fore two real-world factors that have received little attention in theoretical research so far. One critical assumption is that investors have heterogeneous preferences over the type of advice they wish to purchase: some care only about the financial returns generated by the issuer, while others care about nonfinancial aspects of a company's business, such as its environmental and human rights policies. The other important factor is the existence of competition: while proxy advice has been widely criticized, little attention has been paid to the fact that advice is produced in a competitive market and must survive a competitive test. We seek to understand the nature of advice that emerges from competition between profit-maximizing advisory firms.

Our central finding is that when investors have heterogeneous preferences, the profit-maximizing advice policies of advisory firms do not necessarily maximize the issuer's value or the utility of the funds that purchase advice. Instead, equilibrium proxy advice is slanted toward the preferences of funds with high expressive voting benefits, which we argue are likely to be those focused on "socially responsible" investment. Competition attenuates this bias, but except in the limit does not remove it. This bias in equilibrium advice and voting outcomes occurs even if SRI funds comprise only a small fraction of investors. In this sense, our analysis complements other research showing that while empowering shareholders seems essential for good governance, some governance processes can lead to unanticipated distortions in corporate decisions.

Our analysis suggests the need for more thinking about the appropriate normative criteria for evaluating corporate election outcomes. Value maximization is not an adequate criterion when investors receive nonpecuniary returns from corporate actions, such as those who receive a disutility if the company operates in a country that violates human rights. Hart and Zingales (2017) argue for the use of shareholder welfare (utility) maximization as a criterion in contexts like this. While conceptually appealing, this criterion is silent on the issue of heterogeneous preferences: when investors disagree, whose welfare should be maximized? It is natural to think of using voting to answer this question, but then we need a way to assess voting outcomes themselves. As a starting point for discussion, we develop the criterion of informed representation. Informed representation is maximized when election outcomes are those that would be favored by a majority of voters if all voters were fully informed. This is an analytically convenient concept since it boils down to comparing election outcomes to the ideal outcomes of a fully-informed median voter.

Nearly a century ago, Berle and Means (1932) called attention to the separation of ownership and control in modern corporations, which they attributed to the inability of dispersed shareholders to exert control over management. The situation has become even more complicated today, as the preponderance of stock in major corporations is now held by often-passive institutional investors with little interest in monitoring corporate policies, and limited capacity to assess management quality or decisions. To fill the gap, proxy advisory firms have emerged that specialize in monitoring companies and advising shareholders on how to vote. The hope is that these information intermediaries will allow dispersed shareholders to exercise effective control. One conclusion from our study is that even a competitive market for information intermediaries may not seamlessly fill the knowledge gap in the market. While competition among advice suppliers can lead to lower costs and informed voting, it can also lead to a lack of diversity in advice and slanted corporate elections.

## 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)

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

It is useful to define

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

and rewrite (A.1)

$$V_s'(\alpha) = (1-\alpha)K(\alpha). \quad \text{(A.3)}$$

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. We then have $V_s'(\alpha) > 0$.

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)

$$\begin{aligned} V_r'(\alpha) &= \int_{-\infty}^{\infty} \left[ -\frac{s}{(1-\alpha)^2} \frac{\alpha s}{(1-\alpha)} f\left(s, -\frac{\alpha s}{(1-\alpha)}\right) - \frac{s}{(1-\alpha)^2} \frac{\alpha s}{(1-\alpha)} f\left(s, -\frac{\alpha s}{(1-\alpha)}\right) \right] ds \\ &= -\alpha \left\{ \frac{2}{(1-\alpha)^3} \int_{-\infty}^{\infty} \left[ s^2 f\left(s, -\frac{\alpha s}{(1-\alpha)}\right)\right] ds. \right\} \end{aligned} \quad \text{(A.4)}$$

We next show that the term in brackets simplifies to $K(\alpha)$. Perform a change in variables $s = -(1-\alpha)r/\alpha$:

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

Rewrite (A.4)

$$V_r'(\alpha) = -\alpha K(\alpha), \quad \text{(A.5)}$$

which implies $V_r'(\alpha) < 0$. When $\alpha = 1$, the PA 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 $(\lambda, \theta)$ is willing to pay $U_{\lambda\theta}(\alpha)$ for advice $\alpha$.

Claims 1 and 2 imply that advice has a positive value, $U_{\lambda\theta}(\alpha) = \lambda \{\theta V_s(\alpha) + (1-\theta)V_r(\alpha)\} \geq 0$. The fund's outside option is zero, therefore the value of information is $U_{\lambda\theta}(\alpha)$.

**Claim 4)** $U_{\lambda\theta}'(\alpha) > 0$ if $\alpha < \theta$, and $U_{\lambda\theta}'(\alpha) < 0$ if $\alpha > \theta$.

Compute the marginal value of an advice and simplify by using (A.3) and (A.5):

$$\begin{aligned} U_{\lambda\theta}'(\alpha) &= \lambda [\theta V_s'(\alpha) + (1-\theta)V_r'(\alpha)] \\ &= \lambda [\theta(1-\alpha)K(\alpha) - (1-\theta)\alpha K(\alpha)] \\ &= \lambda [\theta - \alpha] K(\alpha). \end{aligned} \quad \text{(A.6)}$$

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 Lemma 2:*** 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 $(\alpha_1^*, \alpha_2^*)$ 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 $\alpha_2^*$, it must be the case that

$$\alpha_1^* \in \arg\max_{\alpha_1 \in [0, \alpha_2^*]} \mathcal{U}(\alpha_1, \alpha_2^*). \tag{A.7}$$

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 (so we can disregard it in the first-order condition) and $\tilde{\theta}$ is a function of $\alpha_1$, the partial derivative with respect to $\alpha_1$ has the same sign as

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

The term in brackets in the first line is zero because, by definition, a fund with preference $\tilde{\theta}$ is indifferent between advice policies $\alpha_1$ and $\alpha_2^*$ — 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 $\theta$ 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.3) and (A.5)

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

This derivative is similar to (A.6), and implies that the first-order condition must satisfy $\alpha_1^* = \hat{E}[\theta|\theta \leq \tilde{\theta}]$, which implies that (20) must hold.

Fixed $\alpha_1^*$, the same logic applied to $\alpha_2$ implies that (21) must hold. Finally, condition (22) defines the preference $\tilde{\theta}$ of the funds which are indifferent between the two policies, concluding the proof. □

***Proof of Proposition 2:*** **Part (i)** We start by writing the profit function of Firm 1, given policies $(\alpha_1, \alpha_2)$. Funds with preference $(\lambda, \theta)$ 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.[18] Therefore, the profit of Firm 1 is

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

Therefore, to maximize its own profit given $\alpha_2$, the best response of Firm 1 is to choose the policy $\alpha_1$ 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

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

[18] 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.

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

Consequently, if policies $(\alpha_1^*, \alpha_2^*)$ maximize $\mathcal{U}$, then inequalities (23) and (24) most hold; therefore, our results imply that $(\alpha_1^*, \alpha_2^*)$ form a competitive equilibrium. From Lemma 2, 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 (23) holds, which in turn also implies that (A.7) holds since $\alpha_1^* < \alpha_2^*$. In the proof of Lemma 2 we have shown that (A.7) implies $\alpha_1^* = \hat{E}[\theta|\theta \leq \tilde{\theta}]$. Similar logic applies to Firm 2: inequality (24) 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 (20), (21), and (22) 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.

**Part (iii)** Combine conditions (20), (21), and (22) into a single equation. First, using (18), rewrite condition (22) and then substitute $\alpha_1 = \hat{E}[\theta \leq \tilde{\theta}]$ and $\alpha_2 = \hat{E}[\theta \geq \tilde{\theta}]$:

$$\tilde{\theta} \left[ V_s(\hat{E}[\theta \leq \tilde{\theta}]) - V_s(\hat{E}[\theta \geq \tilde{\theta}]) \right] + (1 - \tilde{\theta}) \left[ V_r(\hat{E}[\theta \leq \tilde{\theta}]) - V_r(\hat{E}[\theta \geq \tilde{\theta}]) \right] = 0. \quad \text{(A.8)}$$

Therefore, defining 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 $\Gamma$ is continuous, it crosses zero at least once. If it only crosses zero once, then the equilibrium must be unique. From Lemma 2 we know that there is a pair of policies $\alpha_1^* < \alpha_2^*$ that maximize total expressive voting utility $\mathcal{U}$; from Proposition 2 we know that these policies must form a competitive equilibrium; therefore, if the competitive equilibrium is unique, then it must maximize $\mathcal{U}$, concluding the proof.

□

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

**Part (i)** Write the profit function of Firm $k$, given policies $(\alpha_k, \alpha_{-k})$. 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 $(\lambda, \theta)$ 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

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

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

Consequently, if a policy vector maximizes total expressive voting utility (and note that such vector exists), then it constitutes a competitive equilibrium, 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 and Lemma 2, therefore we omit the details.

**Part (iii)** The result follows for two reasons. First, if the median fund $\theta_m$ prefers to purchase the advice $\alpha_m^*$, 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 $\alpha_m^*$ 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 10). Therefore, a majority of funds will be following the same recommendation as $\alpha_m^*$. From $\alpha_m^*$ 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] Berle, Adolf and Gardiner Means, *The Modern Corporation and Private Property*, 1932.

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

[4] Brennan, Geoffrey and Loren Lomasky, *Democracy and Decision: The Pure Theory of Electoral Preference*, Cambridge, UK: Cambridge University Press, 1993.

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

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

[7] Downs, Anthony, *An Economic Theory of Democracy*, New York, NY: Harper & Row, 1957.

[8] Edelman, Paul H., Randall S. Thomas, and Robert B. Thompson, "Shareholder Voting in an Age of Intermediary Capitalism," *Southern California Law Review*, September 2014, Vol. 87, 1359-1434.

[9] Fama, Eugene F. and Merton H. Miller, *The Theory of Finance*, Hinsdale, IL: Dryden Press, 1972.

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

[11] Friedman, Milton, "The Social Responsibility of Business is to Increase its Profits", *The New York Times Magazine*, September 13, 1970.

[12] Gallagher, Daniel M., "Outsized Power and Influence: The Role of Proxy Advisors," *Washington Legal Foundation, Critical Legal Issues Working Paper Series* Number 1187, August 2014.

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

[14] Hart, Oliver and Luigi Zingales, "Companies Should Maximize Shareholder Welfare Not Market Value," *Journal of Law, Finance, and Accounting*, 2017, Vol. 2(2), 247-274.

[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] Larcker, David F., Allan L. McCall, and Brian Tayan, "And Then A Miracle Happens! How Do Proxy Advisory Firms Develop Their Voting Recommendations?," *Stanford Closer Look Series*, February 2013.

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

[20] Levit, Doron and Anton Tsoy, "A Theory of One-Size-Fits-All Recommendations," *American Economic Journal: Microeconomics*, forthcoming.

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

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

[23] Malenko, Andrey, Nadya Malenko, and Chester S. Spatt, "Creating Controversy in Proxy Voting Advice," University of Michigan and Carnegie Mellon University, *working paper*, 2021

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

[25] 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.

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

[27] 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.

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

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

[30] Shu, Chong, "The Proxy Advisory Industry: Influencing and Being Influenced," *Journal of Financial Economics*, 2024, 154, 103810.

[31] 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.

[32] U. S. House of Representatives, Climate Control: Exposing the Decarbonization Collusion in Environmental, Social, and Governance (ESG) Investing, June 11, 2024.

[33] 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.

[34] Zingales, Luigi, Jana Kasperkevic, and Asher Schechter, Milton Friedman: 50 Years Later, Chicago, IL: Stigler Center for the Study of the Economy and the State at the University of Chicago, 2020, e-book available at: https://promarket.org/wp-content/uploads/2020/11/Milton-Friedman-50-years-later-ebook.pdf.