Shareholder Preference Structure and Voting Rules

Ole-Andreas Elvik Næss*

April 29, 2026

Abstract

I study whether decentralized shareholder voting can aggregate heterogeneous investor preferences into stable collective choices. In ten survey experiments across investor populations, I elicit rankings over feasible outcomes, including settings constructed from actual proxy ballots. Preferences exhibit strong latent order: they are consistent with single-crossing, and Condorcet cycles are rare, even in richer choice sets and board elections. Comparing voting rules to a willingness-to-pay benchmark, plurality performs worst when broadly acceptable compromise options attract second-choice support, while ranking-based rules track the benchmark more closely. Thus, shareholder preferences are often aggregable, but the voting rule matters. A key implication is that shareholder disagreement over climate and social preferences is often sufficiently structured for voting institutions to recover coherent collective choices.

*University of Bergen and Norwegian School of Economics. I gratefully acknowledge research support from the Norwegian Finansmarkedsfondet. I have no financial relationships or other potential conflicts of interest to disclose. Preliminary and incomplete.

1 Introduction

Shareholder democracy is expanding, as large asset managers increasingly allow ultimate investors to influence how fund-held shares are voted. As of 2026, both Vanguard and BlackRock have extended voting choice to funds covering several trillion dollars in assets, marking a rapid shift toward more decentralized shareholder voice. At the same time, proxy advisors are moving away from standardized benchmark approaches toward more customized voting frameworks, most clearly illustrated by Glass Lewis' announcement that, beginning in 2027, it will offer multiple client-specific voting perspectives rather than a single house view (Glass Lewis, 2025). Together with the introduction of universal proxy cards, which turn previously binary board contests into settings where investors can choose among multiple candidates, these developments make it easier for owners to express non-binary preferences in corporate elections. At the same time, shareholders disagree not only about profit versus sustainability, but also across several non-financial dimensions (Barber et al., 2021; Bauer et al., 2021; Ericson, 2024; Hart et al., 2024; Riedl and Smeets, 2017). The central empirical question is therefore not whether shareholder preferences are multidimensional, but whether they nevertheless contain enough latent structure to be aggregated into stable collective choices.

This is not obvious. Classical social-choice theory emphasizes how difficult it can be to aggregate unrestricted preferences (Arrow, 1951). By contrast, when preferences satisfy structural restrictions such as single-crossing or single-peakedness, majority-based aggregation can become stable and informative (Black, 1948; Gans and Smart, 1996). Whether such structure is present in shareholder voting is an open empirical question. Answering it requires complete preference rankings over feasible alternatives, yet real-world voting data typically reveal at most a single cast vote. I therefore use survey experiments, which make it possible to elicit complete rankings over carefully defined shareholder outcomes while holding fixed the information available to respondents.

This paper asks two questions. First, do shareholder preferences over multidimensional corporate trade-offs nevertheless line up along a latent one-dimensional ordering of investors,

so that collective choices remain coherent? Second, when ultimate investors are allowed to express such preferences more directly, which voting rules best translate them into a collective vote? To answer these questions, I conduct ten survey experiments covering more than 13,000 respondent-level observations from both retail and institutional investors across different countries. I first construct a deterministic and auditable algorithm that maps actual U.S. proxy ballots into three feasible voting outcomes using fixed dictionaries and pre-specified selection rules. These outcomes correspond to support for a climate-oriented proposal, a social proposal, or the board's position of not supporting any of these proposals. I then elicit complete rankings over those outcomes from U.S equity investors (Experiment 1) and Norwegian sovereign-wealth-fund beneficiaries (Experiment 2). Experiments 3 to 5 test transportability by varying margins that should have overturned the result if it were merely an artifact of one particular vignette, such as the respondent population and the issue space. Together, across Experiments 1 to 5, respondents are asked to rank three outcomes—one climate-oriented, one profit-oriented, and one capturing a range of other non-climate objectives—across 27 distinct scenarios. These rankings form the basis for the analysis of preference structure.

The first set of findings shows that shareholder preferences are structured. I begin by constructing a test of single-crossing requirements. Intuitively, if investors at some point along the ideological spectrum (for example, in the middle) prefer a left alternative to a right alternative, then all investors further to the left must also prefer the left alternative. Symmetrically, if investors at some point prefer the right alternative, then all investors further to the right must do the same. In other words, preferences can shift once along the spectrum, but they cannot reverse. Across Experiments 1 to 5, the ideology-based single-crossing test yields zero violations across 522 testable implications. A global randomization test shows that observing zero violations under the null that ideology is unrelated to ranking behavior is unlikely ($p = 0.014$). Permutation-based inference leads to the same conclusion, and complementary monotone-trend diagnostics point in the same direction. In addition, tests for single-peakedness provide substantial support for structured preferences: for a given ordering of alternatives, most respondents do not rank the middle option below both extremes.

Informally, these results imply that shareholders can be organized along a latent left-to-right dimension, even when the underlying disputes concern profit, climate, and a wide range of other social objectives. In this sense, the median-voter logic re-emerges empirically despite the apparent multidimensionality of shareholder conflict.

The second, and more practical, question is how decentralized shareholder voice should be translated into a collective vote. To study this, I first use respondents' stated preference rankings to compare how different voting rules—plurality, Borda, runoff, and Condorcet-type rules—aggregate preferences. To evaluate these rules more formally, I then construct a cardinal benchmark based on respondents' willingness to pay to have the researcher cast a higher-ranked rather than a lower-ranked vote on their behalf in a future shareholder meeting.

I find that plurality performs poorly: it can fail to select the benchmark-maximizing outcome and can generate economically large welfare losses of up to roughly one-third. By contrast, Borda, Condorcet, and runoff perform substantially better and often select the same outcome. In several scenarios, the ranking patterns suggest that one of the three alternatives functions as a middle option. That middle option often maximizes the cardinal benchmark and wins pairwise majorities, yet plurality can still select an extreme because it ignores how supporters of each extreme rank the compromise relative to the opposite side. This mechanism is especially important in polarized environments, where many investors are split across opposing extremes but still view an intermediate option as preferable to the other side's extreme. In such settings, plurality can reward concentrated first-choice support at the poles even when a compromise ha broader overall support. The immediate policy implication is that institutions that decentralize shareholder voice should not rely mechanically on first-past-the-post aggregation when compromise options are feasible.

Experiment 6 extends the baseline design by showing that preferences remain structured and consistent when the set of possible outcomes expands from three to four. Experiment 7 further broadens the policy menu to six outcomes and shows that structured preferences continue to emerge, while plurality again performs poorly relative to ranking-based rules. This indicates that the baseline results are not driven by a restrictive choice set. Finally, in a universal-proxy board-election setting, investors frequently express mixed preferences over

board composition, and the majority-preferred outcome is a compromise slate. Yet candidate-by-candidate plurality-at-large can still translate these preferences into an all-activist slate. Taken together, these results show that the effectiveness of shareholder democracy depends not only on shareholder preferences themselves, but also on the ballot institutions that aggregate them.

The last part of the paper studies what happens when structure on preferences is absent. In explicitly cyclical environments, neither professionals (Experiment 8 surveys 16 asset managers who, on average, oversee more than $1 billion each) nor retail respondents (Experiment 9 and 10) converge on a single obvious aggregation rule: many abstain, while others split between middle mappings, profit maximization, and cardinal averaging.

Taken together, the paper identifies when shareholder preferences are structured enough for coherent aggregation and shows which voting rules perform best under that structure. Two implications follow. First, as pass-through voting, proxy customization, and universal-proxy elections give ultimate investors more direct voice, decentralized shareholder voting need not produce chaotic aggregation. Second, the choice of voting rule matters: institutions should use rules that can recover broadly supported compromise outcomes rather than relying mechanically on plurality.

Related Literature and Contribution The primary contribution of this paper is to the finance literature on shareholder democracy, delegated voting, and the representation of ultimate owners in corporate voting. As asset managers expand pass-through voting and related mechanisms, a growing literature studies when voting should be decentralized and how intermediaries represent end-investor preferences (Herrmann et al., 2024; Hu et al., 2024; Malenko and Malenko, 2023; Montagnes et al., 2024; Zingales et al., 2025; Zytnick, 2022). Related work analyzes the benefits and costs of shareholder voice through proposals, engagement, activism, and trading around votes (Gantchev and Giannetti, 2021; Gormley et al., 2024; Kakhbod et al., 2023; Levit et al., 2024; Li et al., 2022; McCahery et al., 2016). While the classic benchmark of shareholder-value maximization (Friedman, 1970; Jensen, 2001) effectively collapses shareholder interests into a common pecuniary objective, more

recent work allows firms to reflect shareholder welfare and non-pecuniary investor objectives (Gollier and Pouget, 2022; Hart and Zingales, 2017; Morgan and Tumlinson, 2019).[1] This paper contributes a prerequisite institutional-design question to that literature: once voting power is shifted from intermediaries to ultimate owners, are shareholder preferences sufficiently structured to permit stable aggregation, or will decentralized voice generate fragmented and unstable corporate choices?

A second contribution is to the literature on shareholder voting and heterogeneous investor preferences. A growing finance literature documents that investors differ in their preferences over climate, impact, and other moral or political dimensions in portfolio and governance choices (Bauer et al., 2021; Bolton et al., 2020; Bubb and Catan, 2022; Giglio et al., 2025; Hartzmark and Sussman, 2019; Heeb et al., 2023; Iliewa et al., 2024; Krueger et al., 2020; Riedl and Smeets, 2017; Starks, 2023). Recent experiments likewise show that investors are willing to pay for socially responsible assets and study how they jointly invest and vote on ESG issues (Brodback et al., 2025; Duffy et al., 2022). The contribution of this paper is not to add another data point showing that investors disagree, but to analyze the *structure* of that disagreement. The contribution is empirical: theory shows that stable aggregation is possible when preferences satisfy domain restrictions such as single-peakedness or single-crossing, but it is not known whether shareholder preferences in salient corporate settings actually satisfy those restrictions. Existing work on shareholder voting typically studies observed votes as equilibrium objects that reflect not only preferences, but also information acquisition, proxy advice, delegation, engagement, trading, and strategic behavior (He et al., 2023; Kakhbod et al., 2023; Levit et al., 2023, 2024; Malenko and Malenko, 2019; Maug and Rydqvist, 2009). Observed votes are not well suited to answering the question posed here, because the same underlying preference profile can generate different voting patterns depending on the institutional environment, the information available to shareholders, and whether they vote strategically or sincerely. This paper instead takes heterogeneous preferences over profits and non-pecuniary corporate objectives as the primitive object of interest and abstracts from information acquisition, strategic voting, and other equilibrium forces in order to study

[1] Recent theory also studies how shareholder voting over such objectives differs from political voting under one-person-one-vote rather than one-share-one-vote (Döttling et al., 2024).

whether those preferences satisfy domain restrictions that permit stable aggregation, and how much commonly used voting rules lose when they ignore the full ranking. By eliciting complete rankings directly, the paper isolates the preference component from confounding effects, abstracting from information aggregation, strategic voting, and intermediary frictions.

The paper further contributes to the design of voting rules in corporate governance and, more broadly, to social choice theory and the empirical comparison of aggregation mechanisms in economically consequential environments. Classic results show that unrestricted preference aggregation is difficult (Arrow, 1951), whereas structured domains such as single-peaked and single-crossing preferences generate sharper predictions (Black, 1948; Gans and Smart, 1996; Nehring and Puppe, 2007). Within corporate finance, Meirowitz et al. (2024) theoretically show that majority aggregation remains well behaved when shareholders trade off firm value against a single social objective but may break down once multiple social dimensions are in play. Maskin (2025) provides an axiomatic rationale for Borda-type aggregation when ranking positions contain information about intensity, and quadratic voting is another mechanism designed to incorporate preference intensity (Lalley and Weyl, 2018; Posner and Weyl, 2015). Related empirical work compares different voting rules in laboratory and survey settings (Bassi, 2015; Bouton et al., 2016; Van der Straeten et al., 2010; Forsythe et al., 1996; Granić, 2017; Igersheim et al., 2022). A distinctive contribution of this paper is the construction of a consequential cardinal benchmark. In the core modules, respondents state a willingness to pay to have the researcher cast a higher-ranked rather than a lower-ranked shareholder vote on their behalf, which yields a money-metric benchmark for comparing aggregation rules. This is unusual relative to most political and survey-based social-choice settings, where preference intensity is harder to discipline with consequential monetary tradeoffs over realized collective outcomes and is often inferred only indirectly. The shareholder setting therefore allows the paper to compare plurality, Borda, runoff, and Condorcet-type rules not only on Condorcet-efficiency or induced-preference grounds, but on how closely they track consequential willingness-to-pay tradeoffs. This comparison is especially relevant in the growing set of corporate settings in which shareholders face more than a binary yes/no choice, including voting-choice menus, bundled policy environments,

and universal-proxy board elections. It also relates to a nascent literature on the design of shareholder-democracy institutions, which increasingly confronts the question of how heterogeneous end-investor preferences should be translated into a collective vote once voting authority is decentralized (Herrmann et al., 2024; Meirowitz et al., 2024; Montagnes et al., 2024; Zingales et al., 2025).

Finally, the last part of the paper also speaks to the positive analysis of aggregation when preferences are not well structured. Ambuehl and Bernheim (2021) show that individuals often favor compromise-oriented aggregation rules and may infer intensity from ordinal rankings. Related experiments study how individuals choose voting institutions or majority thresholds behind the veil of ignorance and how those choices trade off efficiency, fairness, and minority protection (Engelmann et al., 2023; Weber, 2020). The paper brings these questions into a shareholder setting by studying how practitioners and lay respondents think unstructured rankings should be mapped into a corporate vote. In doing so, it complements the abstract social-choice literature with evidence from a domain in which the objects being aggregated are concrete corporate policies and board outcomes.

This paper is organized as follows. Section 2 develops a theoretical framework for preference aggregation in three-dimensional choice environments and derives the key empirical requirements. The empirical strategy is summarized in Figure 1. Section 3 presents the design of Experiments 1–5, and Section 4 reports their results. Section 5 extends the baseline setting to include a fourth alternative, while Section 6 presents the results from a three-module experiment that studies richer implementation questions. Section 7 presents Experiments 8–10, which examine aggregation in explicitly cyclical environments. Section 8 concludes.

2 Theoretical Framework

This section adapts social-choice theory to the shareholder-voting environments studied in the paper and turns its predictions into empirical diagnostics. It explains why aggregating rankings is difficult once there are at least three feasible outcomes, identifies preference

restrictions under which majority aggregation is well behaved, and clarifies what follows from theory versus what requires data. I focus on aggregating *preferences*, not *information*. Ballots may also reflect heterogeneous information and strategic voting, but the question here is how alternative rules aggregate ranked preferences.[2]

2.1 Preferences and the aggregation problem

A corporate decision environment has a finite set of feasible outcomes $X = \{x_1, \ldots, x_m\}$; the baseline analysis uses $X = \{A, B, C\}$. These labels carry no geometry unless one is imposed, so absent an explicit axis there is no distinguished "middle" alternative. There are n shareholders, indexed by $i \in \{1, \ldots, n\}$, each with a strict ranking $\succ_i$ over X; the profile is $\succ = (\succ_1, \ldots, \succ_n)$. Shareholder i has voting weight $s_i > 0$. A social choice rule is a map $f : \mathscr{R}(X)^n \to X$. A useful benchmark is pairwise majority: for $x, y \in X$, $x \succ_M y$ iff $\sum_i s_i \mathbf{1}\{x \succ_i y\} > \sum_i s_i \mathbf{1}\{y \succ_i x\}$. A *Condorcet winner* is an outcome that beats every other outcome pairwise.

2.2 Why aggregation is difficult

With $|X| \geq 3$, unrestricted rankings can generate cycles. In the Condorcet paradox, three equally weighted shareholders with $A \succ_1 B \succ_1 C$, $B \succ_2 C \succ_2 A$, and $C \succ_3 A \succ_3 B$ imply $A \succ_M B$, $B \succ_M C$, and $C \succ_M A$. Arrow's impossibility theorem formalizes the broader point: when $|X| \geq 3$, no social welfare function aggregates arbitrary rankings into a complete social ordering while satisfying standard desiderata such as Pareto efficiency, independence of irrelevant alternatives, and non-dictatorship (Arrow, 1951). The implication is not that aggregation is impossible, but that positive results require more assumptions, such as additional structure on preferences.

[2]For classic manipulation results, see Gibbard (1973); Satterthwaite (1975).

2.3 Domain restrictions that restore tractability

Two classic domain restrictions restore majority stability. Both replace potentially shifting majority coalitions with a median logic: once preferences are sufficiently structured, pairwise majority no longer cycles. Single-crossing orders *shareholders*; single-peakedness orders *alternatives*. In many corporate-voting settings, it is easier to justify a one-dimensional ordering of shareholders—for example, by political ideology or attitudes toward stakeholder governance—than a common one-dimensional axis for the outcomes, which are often multidimensional bundles. Because single-crossing requires only the former, whereas single-peakedness also requires a defensible geometry of outcomes, single-crossing is the primary domain restriction in the paper, while single-peakedness serves as a complementary benchmark when an outcome axis is substantively meaningful.

2.3.1 Single-crossing

Intuitively, single-crossing is a restriction on heterogeneity across shareholders: once shareholders are ordered from left to right, each pairwise contest can change hands at most once along that order. Fix a shareholder order $1 \prec_v 2 \prec_v \cdots \prec_v n$. Preferences are *single-crossing* with respect to $\prec_v$ if, for every pair $x, y \in X$, the set $\{i : x \succ_i y\}$ forms an interval under $\prec_v$; equivalently, as one moves along the order, the preference between x and y changes at most once.

2.3.2 Single-peakedness

Intuitively, single-peakedness is a restriction on the ordering of alternatives: all shareholders evaluate outcomes relative to a common axis, each has a most-preferred point on that axis, and alternatives become less attractive as one moves away from that peak on either side. Fix an axis $x_1 \prec x_2 \prec \cdots \prec x_m$ on X. A ranking $\succ_i$ is *single-peaked* with respect to $\prec$ if it has a peak x_i^* and preferences fall as one moves away from the peak on either side; for $m > 3$, this is referred to as the one-dimensional no-valleys condition.

2.3.3 Implications for majority aggregation

Under either single-crossing or single-peakedness, pairwise majority is transitive, so absent pairwise ties a Condorcet winner exists (Black, 1948; Gans and Smart, 1996). The same median logic underlies both results, but the relevant median differs: under single-crossing it is the median shareholder in the order $\prec_v$ (or weighted median shareholder with unequal weights), whereas under single-peakedness it is the median peak on the common axis.

The two restrictions are distinct. For example, with ordered types $L \prec_v M \prec_v R$ and rankings $L : A \succ B \succ C$, $M : A \succ C \succ B$, and $R : C \succ B \succ A$, the profile is single-crossing with respect to $L \prec_v M \prec_v R$ but not single-peaked on any common axis.

2.3.4 Empirical implementation in the three-outcome design

The baseline design has three alternatives, $X = \{A, B, C\}$, and a fixed partition of respondents into three groups ordered by ideology $L \prec_v M \prec_v R$, where L denotes left-wing, M is middle/center and R is right-wing.

Single-crossing diagnostic. Empirically, I test the grouped implication that pairwise support should move monotonically across the ordered groups. For an ordered pair $p_\ell \prec p_r$ (meaning ℓ is to the left of r), write $p_\ell \succ_g p_r$ if a strict majority in group $g \in \{L, M, R\}$ ranks p_ℓ above p_r.[3] Single-crossing requires six implications per pair:

$$p_\ell \succ_R p_r \Longrightarrow p_\ell \succ_M p_r, \quad p_\ell \succ_M p_r \Longrightarrow p_\ell \succ_L p_r, \quad p_\ell \succ_R p_r \Longrightarrow p_\ell \succ_L p_r,$$

$$p_r \succ_L p_\ell \Longrightarrow p_r \succ_M p_\ell, \quad p_r \succ_M p_\ell \Longrightarrow p_r \succ_R p_\ell, \quad p_r \succ_L p_\ell \Longrightarrow p_r \succ_R p_\ell.$$

Intuitively, once even the rightmost group still prefers the left alternative, the center and left groups must do so as well; symmetrically, once the leftmost group prefers the right alternative, so must the groups to its right. Each ordered pair therefore yields six implication checks, and with three pairwise comparisons each scenario yields $3 \times 6 = 18$ checks. When

[3] If a group is exactly split 50–50 on a pairwise comparison, I treat that comparison as non-binding for the corresponding implication test.

the ordering is only partially specified ex ante—so that the relative position of one pair remains unresolved, although its position relative to the third alternative is known—I evaluate both admissible completions rather than choose one ex post. In those scenarios, the extra admissible orientation adds six further implication checks, for a total of 24.

Implementation. A prediction is *binding* when its condition is met and *violated* when the implied outcome does not follow; if the condition is not met, the prediction is non-binding. Hence, for the first requirement above, it is binding if and only if $p_\ell \succ_R p_r$, and violated if $p_\ell \succ_R p_r$ holds while $p_\ell \succ_M p_r$ does not hold. The empirical section reports binding and violated checks and complements them with resampling-based inference.

Single-peakedness diagnostic. For three alternatives, single-peakedness admits a simple characterization once an axis is fixed. Let the axis be $x_L \prec x_M \prec x_R$. A ranking is single-peaked if and only if the middle alternative x_M is not ranked last; equivalently, the only forbidden rankings are those that place x_M last. Let $\hat{p}_j^{\text{last}}$ denote the sample share of respondents who rank alternative $j \in \{A,B,C\}$ last. Exact single-peakedness with respect to a fixed axis holds if and only if $\hat{p}_{x_M}^{\text{last}} = 0$. When the axis is not fixed *ex ante*, I define the best-fitting-axis statistic $\widehat{\text{MaxSP}} = 1 - \min_{j\in\{A,B,C\}} \hat{p}_j^{\text{last}}$, which provides a convenient diagnostic for assessing the extent to which observed rankings can be rationalized by single-peaked preferences.

2.3.5 Preference structure with more than three alternatives

With more than three alternatives, single-crossing remains a restriction on the ordering of shareholders rather than on the geometry of outcomes: for each pair of outcomes, support can switch at most once along the shareholder order. Standard single-peakedness also extends once a one-dimensional axis is specified. But when outcomes are multidimensional, one needs an explicit geometry, and the appropriate extension is generalized single-peakedness relative to a betweenness relation (Nehring and Puppe, 2007). In the four-outcome bundle extension studied later, I operationalize this idea by testing issue separability on the natural 2×2 bundle space.

2.4 Voting rules

I now turn to the voting rules used in practice. The analysis focuses on four commonly used ordinal procedures that will also be used in the empirical analysis. **Plurality** selects the alternative receiving the most first-place votes. **Borda** assigns scores $m-1, m-2, \ldots, 0$ to the alternatives ranked first through last on each ballot and selects the alternative with the highest total score. **Condorcet** selects an alternative that defeats every other alternative in pairwise majority comparisons, when such an alternative exists. **Instant-runoff voting (IRV)** iteratively eliminates the alternative with the fewest first-place votes and transfers those ballots to the next-ranked remaining alternative until one alternative obtains a majority.

Under the preference structures discussed above, majority aggregation is well behaved. If preferences are single-crossing or single-peaked, pairwise majority is transitive, so absent ties a Condorcet winner exists. Any voting rule satisfying the Condorcet criterion therefore selects that alternative. This is the familiar median-voter logic: once preferences are sufficiently ordered, majority rule identifies a stable middle outcome.

This conclusion is important, but limited. It shows that preference structure provides one way around Arrow's impossibility result by restricting the domain of preferences. It does not imply that Condorcet-consistent rules are necessarily superior from a welfare perspective, nor does it imply that non-Condorcet rules must perform poorly. The normative ranking of voting rules depends on which criteria one emphasizes.

Indeed, one need not accept Arrow's full list of criteria as the only basis for evaluating social choice rules. In particular, Maskin (2025) argues that Arrow's independence of irrelevant alternatives can be too restrictive when complete rankings contain information about relative strength of preference. With three alternatives A, B, C, IIA requires that the social ranking of A and B depend only on how voters rank A versus B. Thus, if a voter changes from $C \succ A \succ B$ to $A \succ C \succ B$, the social ranking of A and B must remain unchanged because the voter still ranks A above B, even though the change improves the overall position of A. Under a weaker independence requirement, Borda can use this additional ordinal information.

The theory therefore leaves open the central comparison in this paper. Preference structure

tells us when a Condorcet winner exists, but it does not tell us which ordinal rule is best in welfare terms. To evaluate the performance of plurality, Borda, IRV, and Condorcet against one another, a cardinal benchmark is needed.

2.5 A cardinal benchmark: willingness to pay

Ordinal rankings alone do not identify cardinal utilities or interpersonal intensity comparisons (Arrow, 1951). To evaluate the consequences of different ordinal voting rules, I therefore construct an empirical cardinal benchmark. In the incentivized environments studied here, I use willingness to pay (WTP) as a practical money-metric benchmark. WTP records how much respondents are willing to sacrifice for one outcome relative to another and therefore captures information about preference intensity that ordinal rankings omit.

At the same time, WTP is not identical to welfare. As emphasized in the literature, money-metric measures can diverge from welfare because they depend on income and liquidity and may also reflect reference points, entitlements, and other features of the choice environment (Bar-Gill, 2021; McFadden, 2017; Sunstein, 2007).

I therefore use WTP as an empirical benchmark rather than as a complete welfare concept. In that sense, the paper's question is separate from the existence of a Condorcet winner. Preference structure tells us when majority rule yields a coherent ordinal benchmark; the WTP analysis asks which implementable ordinal voting rule comes closest to a cardinal measure of preference intensity. This benchmark makes it possible to study both how often plurality, Borda, IRV, and Condorcet select different outcomes and how large the losses are when they select outcomes that do not maximize WTP.

From theory to design. The baseline experimental design uses three outcomes, $X = \{A, B, C\}$. Three is the minimal non-trivial domain for preference aggregation: with only two outcomes, pairwise majority is mechanically transitive, plurality coincides with majority rule, and the central rule comparisons in the paper largely collapse. A three-outcome design is therefore the smallest setting in which structured preferences such as single-crossing and

single-peakedness have empirical bite and in which plurality, Borda, IRV, and Condorcet can in principle diverge. The theory does not, however, determine the substantive content of A, B, and C. Instead, the alternatives should be anchored in a common institutional baseline reflecting realistic environments, so that respondents compare feasible votes within the same shareholder-voting context. Section 3 implements this logic by taking C to be a natural board-line baseline and A and B to be two substantively distinct one-item deviations from that baseline selected by a deterministic algorithm. Importantly, this construction does not impose an ex ante middle option or a particular geometry on the choice set. Later experiments relax this minimal benchmark by expanding the choice set, allowing richer geometries of outcomes, and testing whether the same patterns survive in broader environments.

3 Experiments 1–5: Baseline design and extensions

Experiments 1–5 all implement the baseline three-outcome task implied by the theory-to-design mapping above. The studies assess transportability by varying dimensions that should have overturned the pattern if it were driven by one particular framing: the source of the scenario (actual proxy ballots versus stylized but consequential or hypothetical settings), the shareholder population (U.S. direct owners, representative U.S. respondents, mutual-fund investors, and Norwegian beneficiaries), the object of choice (shareholder proposals versus board candidates), and the precise form of the baseline outcome C (the board line versus an explicit profit-first platform). Across experiments, respondents rank three feasible outcomes anchored in a common institutional baseline, with A and B representing substantively distinct deviations from that baseline. Experiment 1 provides the disciplined core test using algorithmically selected deviations from real proxy ballots. Experiment 2 keeps those ballot-based scenarios fixed while changing the country and ownership channel. Experiment 3 broadens the issue space and the surrounding environments. Experiments 4 and 5 shift from proposal voting to board elections, where non-binary aggregation is institutionally more natural, first in the United States and then in Norway. Figure 1 summarizes this progression.

I. Core three-outcome design and transportability (Experiments 1–5)

Purpose: establish the baseline ranking pattern and test transportability across the paper's main design margins. Respondents rank three feasible outcomes: *A* (social), *B* (climate), and *C* (profit/board).

Exp. 1: disciplined core test ($N = 1456$)
Population: U.S. direct shareholders.
Setting: algorithmically constructed scenarios from actual 2025 proxy ballots.
Role: disciplined baseline and ballot-based WTP benchmark.

Exp. 2: ownership-channel extension ($N = 3200$)
Population: Norwegian beneficiaries.
Setting: same ballot-based scenarios as Exp. 1, reframed for Norway's sovereign wealth fund.
Role: holds scenarios fixed while changing country and ownership channel.

Exp. 3: issue-space extension ($N = 2418$)
Population: representative U.S. respondents.
Setting: potentially consequential scenarios plus broader hypothetical settings.
Role: varies scenario source, issue space, and baseline framing.

Exp. 4: board-election module ($N = 1469$)
Population: U.S. mutual-fund investors.
Setting: hypothetical three-candidate board elections.
Role: changes the object of choice from proposals to candidates.

Exp. 5: cross-country board replication ($N = 1513$)
Population: Norwegian beneficiaries.
Setting: sovereign-wealth-fund version of the board-election module.
Role: cross-country replication of the board module.

II. Richer designs and implementation tests (Experiments 6–7)

Exp. 6: four-outcome bundle extension ($N = 980$)
Role: asks whether the Exp. 1 results survive when the choice set expands from three to four options by adding support for both proposals.

Exp. 7: richer implementation settings
Role: asks whether the core pattern survives larger choice sets (six-option policy menu) and richer implementation (contested board module).

III. Explicitly cyclical environments (Experiments 8–10)

Exp. 8: practitioner study ($N = 16$)
Role: studies how professional Scandinavian institutional investors and fund managers want to aggregate preferences when the ranking profile is explicitly cyclical rather than structured.

Exp. 9: Norway lay replication ($N = 500$)
Role: extends Exp. 8 to a Norwegian lay sample.

Exp. 10: U.S. lay replication ($N = 968$)
Role: extends Exp. 8 to a U.S. lay sample.

Figure 1: Overview of the empirical design across ten experiments.

3.1 Experiment 1: Algorithmic Scenarios from Real Shareholder Ballots

Experiment 1 studies whether shareholder voting preferences exhibit empirically useful structure across five different scenarios and, conditional on the observed preference profiles, which ordinal aggregation rules perform best, relative to a cardinal WTP benchmark.

This experiment took place in January 2026 on Prolific, a widely recognized platform for high-quality participant recruitment (Eyal et al., 2021). 1456 respondents were recruited among US respondents who had self reported being in the equity market, either through holding stocks or ETFs. The average self-reported investment is arond USD 250,000, as reported in Table A5. The pre-analysis plan can be found in the Appendix and AsPredicted #267097.

Data quality safeguard measures To safeguard data quality and minimize automated or AI-assisted participation, I implemented a multi-layer screening protocol prior to the main experimental tasks. First, respondents completed a Google reCAPTCHA verification, where I excluded submissions with suspicious reCAPTCHA v3 scores using a conservative cutoff of 0.7. Second, I administered an audio transcription check in which participants listened to a short recording and entered the spoken number, complementing recent evidence that simple media-based number-recall tasks can help screen out AI agents (Çelebi et al., 2025). Third, I included an instructed-response attention check to detect inattentive clicking. Finally, respondents completed a brief open-ended writing task where I screened out observations with implausibly rapid text entry, consistent with evidence that typing-speed and related interaction patterns in open-text responses are among the most informative indicators of automated responding (Çelebi et al., 2025). Respondents failing any screen were prevented from proceeding to the main study.

3.1.1 Scenario construction

A central design goal of this experiment is to expose respondents to realistic shareholder voting situations while minimizing researcher discretion in scenario design. I construct each scenario from an actual 2025 proxy ballot using a deterministic, pre-specified algorithm described below.

The algorithm starts from a pre-specified universe of the nine largest U.S. firms by market capitalization. For each shareholder proposal on a firm's 2025 ballot, I combine the title, resolved clause, and supporting statement, normalize the text, and classify the proposal using fixed climate and non-climate dictionaries. I then apply a fixed anti–ESG filter to exclude reverse-ESG proposals, such as proposals calling on firms to remove ESG or DEI policies. Among the remaining shareholder proposals opposed by the board, the algorithm selects the highest-support climate proposal and the highest-support non-climate proposal. These two proposals are then rewritten into short, plain-language vignettes.

Once the firm universe, dictionaries, and selection rules are fixed, the mapping from real ballots to survey scenarios is deterministic. The algorithm therefore substantially limits researcher discretion and ensures that the scenarios are tied to salient proposals that actually appeared on large-firm proxy ballots. Online Appendix B provides the full dictionaries, the exact selection rules, and an audit trail showing why firms are either retained or dropped.

In the implementation used in the paper, the algorithm is applied to the nine largest U.S. firms. Four firms are dropped by the algorithm because their 2025 ballots do not contain at least one eligible climate proposal and one eligible non-climate proposal. The remaining five firms are Alphabet, Amazon, Microsoft, Meta, and Berkshire Hathaway. In all five cases, management recommends voting AGAINST both selected proposals, so the two non-management options each represent a one-item departure from the board line. The resulting scenarios are:

- **Alphabet (Google):** a human-rights proposal related to AI-driven targeted advertising, and a climate-disclosure proposal.

- **Microsoft:** a human-rights proposal about data operations in human-rights hotspots, and a proposal on the climate and financial risks of providing AI tools for oil and gas development.
- **Amazon:** a worker-conditions proposal (an independent audit of warehouse working conditions), and a proposal on how data-center growth affects the firm's climate commitments.
- **Meta Platforms:** a proposal concerning hate targeting of marginalized communities, and a proposal on greenhouse-gas emissions reduction.
- **Berkshire Hathaway:** an AI-governance proposal (independent board oversight of AI), and a climate-financing disclosure proposal.

Ranking task and mapping to three strategies. Respondents first read a short introduction explaining that, at shareholder meetings, investors vote on proposals submitted by shareholders and that the board typically provides voting recommendations. They then see the five scenarios described above.

For each scenario, respondents are asked to rank three voting strategies by dragging and dropping them into their preferred order, from most preferred to least preferred. The three strategies are: (i) support the non-climate proposal and oppose the climate proposal (labeled *A* in the paper), (ii) support the climate proposal and oppose the non-climate proposal (labeled *B*), and (iii) follow management by opposing both proposals (labeled *C*). Respondents are explicitly informed that other voting patterns are possible—for example, supporting both proposals—but the experiment focuses on these three strategies because they cleanly compare a non-climate deviation from the board line, a climate deviation from the board line, and the board-line baseline itself. Experiment 6 extends this exact setup to a four-option ranking task by adding the possibility of supporting both proposals.

Incentivized willingness-to-pay elicitation. After completing the ranking tasks, each respondent is randomly assigned to *one* of the five firm scenarios. Respondents answer

incentivized questions designed to measure how much they value having a higher-ranked strategy implemented relative to a lower-ranked strategy. The elicitation uses a multiple price list (MPL) framed as a trade-off between (i) receiving no additional payment while the researchers cast a shareholder vote according to a respondent's higher-ranked strategy, or (ii) receiving a personal cash bonus while the researchers cast the vote according to a lower-ranked strategy. Respondents complete two MPL modules:

1. **First-vs.-third MPL.** Respondents make a series of binary choices between "no bonus and we vote for your highest-ranked strategy" versus "a bonus of $0,1,2,3,10,20 and we vote for your lowest-ranked strategy".

2. **Second-vs.-third MPL.** Respondents then make the analogous series of choices between "no bonus and we vote for your second-ranked strategy" versus "a bonus of $0,1,2,3,10,20 and we vote for your lowest-ranked strategy".

So for example, if a respondent says they prefer "no bonus and we vote for your highest-ranked strategy" over a bonus of 0,1, and 2 dollar, but they prefer a 3 dollar bonus over "no bonus and we vote for your highest-ranked strategy", I set the WTP to 2 dollar.

After each MPL, respondents are shown a short summary of their implied switching point (e.g., preferring the no-bonus option up to some bonus level and preferring the bonus thereafter) and are asked to confirm that the summary reflects their intended choices. If respondents indicate that the summary is not correct, they are given an opportunity to revise by completing the same MPL again.

For each of the six possible rankings, this approach yields WTP values for first- vs. third-choice and second- vs. third-choice comparisons. WTP for the first- vs. second-choice comparison is estimated by subtracting these values. For example, if a participant's preference ranking is $A \succ B \succ C$ and is willing to pay $5 to implement A over C and $3 to implement B over C, then their inferred WTP for A over B is $2.

Implementation and payoffs. To make choices consequential, the survey informs respondents that one participant and one of their MPL decisions will be randomly selected for

possible implementation. The selected participant's choice is implemented by (i) paying the cash bonus associated with the selected MPL row if the bonus option is chosen and (ii) purchasing one share of the assigned company and casting the shareholder vote at the next annual meeting according to the chosen strategy. If the relevant proposal(s) do not appear on the firm's next ballot, the vote cannot be implemented; however, any cash bonus from the selected MPL decision is still paid.

Demographics and order randomization. The survey collects two demographic variables: self-reported political ideology and the approximate total value of the respondent's stock and mutual-fund investments. Political ideology was measured using the following choices: *strongly right-wing*, *right-wing*, *center*, *left-wing*, *strongly left-wing*, and *prefer not to say*. To map responses into the three political categories used for testing single-crossing, *strongly right-wing* and *right-wing* are classified as R, *strongly left-wing* and *left-wing* as L, and all others as M. This classification was pre-registered, and this way of classifying ideology is also used in the other experiments of this paper.

To assess and mitigate potential priming effects, the timing of these demographic questions is randomized. For half of respondents the demographics are asked before the voting scenario, and for the other half they are asked after the scenario.[4]

3.2 Experiment 2

This experiment was conducted in March 2026 through Ipsos Norway ($N = 3{,}201$) and pre-registered as AsPredicted #276348. Respondents were drawn from a nationally representative panel using quotas on age, gender, and region (based on Statistics Norway).

Experiment 2 replicates the scenarios from Experiment 1 in a Norwegian setting, where respondents are asked how they would want Norges Bank Investment Management (NBIM), the manager of Norway's sovereign wealth fund, to vote. As of March 2026, the fund's market value is approximately \$2.1 trillion, corresponding to roughly \$380,000 per Norwegian

[4] I find no evidence of systematic order effects. Rankings and WTP distributions differ significantly in only one of twenty tests. Details are reported in Appendix Table A6.

resident. For context, the median Norwegian household's net worth—excluding assets in the sovereign wealth fund—is about $170,000 (SSB, 2024). Norwegian citizens do not hold direct ownership claims or voting rights in the fund, but are its ultimate economic beneficiaries through the state (NBIM, 2024a). A nationally representative sample therefore captures the preferences of these beneficiaries. For expositional convenience, we occasionally refer to them as *indirect shareholders*, reflecting their economic stake without implying individual equity claims.

3.2.1 Scenarios

After a short introduction explaining the role of NBIM in shareholder voting, the respondents were assigned to the same five scenarios as in Experiment 1, displayed in random order. The scenarios were framed in the context of NBIM, asking them to rank outcomes as how they would want NBIM to vote.

3.3 Experiment 3

Experiment 3 complements the real-ballot evidence in Experiments 1 and 2 in two ways. First, it moves to a demographically representative U.S. sample. Second, while the A-options in Experiments 1 and 2 are disciplined by actual shareholder ballots, this experiment broadens the content of A across a wider set of non-climate stakeholder concerns. The goal is to examine whether the same ranking structure survives in a broader population and when A varies much more. The main evidence comes from four potentially consequential Equinor scenarios (E1–E4), which remain close to the shareholder-voting environment and also permit a WTP elicitation. Four additional hypothetical scenarios (H1–H4) provide a broader test.

3.3.1 Sample

This experiment was fielded on Prolific in August and September 2024 and preregistered on AsPredicted (#188348; see Appendix F). I recruited a demographically representative U.S. sample on certain dimensions and obtained $N = 2{,}418$ valid responses. Each respondent

evaluated four of the eight scenarios, yielding roughly 1,200 observations per scenario.

3.3.2 Design

As in the earlier experiments, respondents ranked a non-climate option (*A*), a climate option (*B*), and a profit-oriented option (*C*). I focus below only on what is different compared to previous experiments.

Scenarios E1–E4. The central part of Experiment 3 is four potentially consequential Equinor scenarios. Respondents were told that the research team owned shares in Equinor and that one randomly selected participant could determine how those voting rights would be used if the corresponding issue came to a vote at Equinor's next shareholder meeting. These scenarios are not drawn from realized ballots, but are designed to reflect salient and plausible trade-offs facing Equinor at the time of the survey, based on contemporaneous public debates about the firm's climate strategy, energy security role, electrification policies, and offshore wind investments. Further details on the underlying debates are provided in Online Appendix C.1. Across all four scenarios, option *B* is the climate-oriented alternative and option *C* is the profit-oriented alternative. In each scenario, *A* captures a different non-climate shareholder concern. In E1, *A* supports developing new gas fields only to strengthen European energy security during the energy crisis, *B* prioritizes climate by stopping all new gas field development, and *C* supports developing all gas fields whenever profitable. In E2, *A* supports continuing oil exports to developing countries to reduce economic hardship, *B* prioritizes climate by reducing such exports, and *C* supports exporting oil whenever profitable. In E3, *A* opposes electrification of oil production due to concerns about higher household energy costs, *B* supports electrification to reduce emissions, and *C* supports electrification only when profitable. In E4, *A* opposes offshore wind development due to concerns about biodiversity and fishing communities, *B* supports offshore wind to mitigate climate change even if unprofitable, and *C* supports offshore wind only when profitable. E1 differs somewhat from the other three because its *A*-option can also be seen as an intermediate position, allowing gas development only in the crisis setting rather than supporting all profitable

expansion.

Scenarios H1–H4. The remaining four scenarios are fully hypothetical and mainly serve to widen the support of A beyond the Equinor context. They vary A across harmful labor practices (H1), water conservation in a water-scarce region (H2), Indigenous rights in a hydroelectric project (H3), and worker health and safety (H4). Full scenario wordings are reported in Online Appendix E.

3.3.3 WTP Elicitation for scenarios E1 to E4

As in Experiment 1, WTP was elicited using a multiple price list. After the ranking task, respondents completed a first-versus-third and second-versus-third elicitation for one of the Equinor scenarios they had previously ranked. The survey concluded with the same political-ideology and investment questions as in Experiment 1.

3.4 Experiment 4

Experiment 4 complements the earlier experiments by changing the object of choice. Instead of ranking shareholder proposals, respondents rank board candidates. This matters because much shareholder influence is exercised through board elections, which is also a setting where non-binarity of options arise naturally when there are more than two candidates. Furthermore, Experiment 4 deliberately extends the set of non-climate A-alternatives by focusing on six new issues that are ranked alongside climate and profit. This experiment was preregistered as AsPredicted (#208601; see Appendix F) and fielded on Prolific in January 2025 among $N = 1{,}469$ U.S. respondents who are self-reported mutual fund investors. Respondents were asked how their main fund should vote in six hypothetical board elections. In each scenario, they ranked three candidates: a profit-oriented candidate (C), a climate-oriented candidate (B), and a third candidate (A) who combined profit with another non-climate objective.

Scenarios B1–B6. The six *A*-candidate types were deliberately varied across distinct stakeholder concerns: reducing AI risks (B1), strengthening workers' union rights (B2), increasing women's representation in leadership (B3), expanding access to medicine (B4), strengthening data security and privacy (B5), and improving animal welfare (B6). Full scenario wordings are reported in Online Appendix E. Finally, respondents reported their political preferences.

3.5 Experiment 5

Experiment 5 mirrors the design of Experiment 4 in the same way that Experiment 2 builds on Experiment 1, replicating three scenarios in a Norwegian setting with indirect fund ownership. This experiment was conducted in February 2025 and surveyed a nationally representative sample of $N = 1,513$ Norwegian adults, with data collection managed by YouGov Norway. The experiment was preregistered (AsPredicted #210314), with details available in Online Appendix F.

Participants first reported their political ideology on a left–right scale. They then evaluated three scenarios (N1–N3), adapted from Scenarios B2, B3, and B6 in Experiment 4. The scenarios remained unchanged from Experiment 4, except that participants were asked how the Norwegian sovereign wealth fund should vote. Instructions are available in Online Appendix E.

To enhance the perceived impact of the survey, respondents were informed that findings would be presented to policymakers and disseminated through media outlets. This statement was accurate, as the results were subsequently disseminated in public-facing outlets, including an op-ed in a leading Norwegian business newspaper.[5] Due to the limited empirical evidence available to Norwegian policymakers on how ultimate economic beneficiaries prioritize other concerns relative to climate and profit, respondents may view their survey responses as influential, increasing their incentive to respond thoughtfully.[6]

[5]Reference omitted for anonymity and will be included in the final version.

[6]This approach aligns with prior research (Næss, 2025) on whether Norwegians support leveraging ownership power for green voting or divestment, a project that also involved broad dissemination of results, including to the fund's management. Moreover, the fund actively exercises its voting power across a wide range of

4 Experiments 1 to 5: Results

Since the structure of experiments 1 to 5 followed similar setup, I here present the results together. While all five contain ranking of different three outcome scenarios, the five scenarios in Experiment 1 as well as E1 to E4 in Experiment 3 additionally contain WTP analysis used to create cardinal benchmarks, which mean that these will get extra attention.

4.1 Descriptive Results

Table 1 reports the distribution of rankings across scenarios in Experiments 1 to 5. Taken together, the descriptive rankings show substantial heterogeneity across shareholder populations and institutional environments. Because the content of option A, the respondent pool, and the choice setting vary across experiments, the relative support for A, B, and C is not expected to be constant. Still, several regularities emerge. In the algorithmic ballot settings of Experiments 1 and 2, the non-climate social option A is usually the most popular alternative, except in the Berkshire scenario, where the climate option B receives more first-place support. Support for voting with the board is also markedly lower in Norway than in the United States. In Experiment 3, support shifts between the climate option and the non-climate option depending on the content of the third objective. In the board-election setting of Experiment 4, support is more concentrated on the profit-oriented candidate, consistent with a stronger emphasis on profit among respondents investing their own funds; a similar, though weaker, pattern is also present in Experiment 1 and is reinforced when responses are weighted by investment holdings (Online Appendix D.2). In the corresponding Norwegian setting of Experiment 5, preferences are less profit-centered. Overall, the descriptive results show that which alternative is most popular depends systematically on both the shareholder population and the institutional environment. This heterogeneity is substantively interesting in its own right and also motivates the central question of the paper, which is whether, despite these shifts in the most popular alternative, the observed rankings nevertheless exhibit a common

corporations, considering both climate and other issues, and publishes its voting record online (NBIM, 2024b), underscoring its strong emphasis on voting decisions.

underlying structure that permits coherent aggregation.

Table 1: Distribution of rankings across all scenarios (shares in %)

Panel A. Algorithm-based firm ballots: Experiments 1 and 2

	Experiment 1					Experiment 2				
Ranking	Alphabet	Microsoft	Amazon	Meta	Berkshire	Alphabet	Microsoft	Amazon	Meta	Berkshire
A≻ B ≻ C	41.5	34.7	46.0	32.6	20.8	61.8	62.3	60.5	54.4	34.7
A ≻ C ≻ B	12.6	12.3	19.5	9.1	11.5	8.0	8.3	8.3	7.4	6.8
B ≻ A ≻ C	17.4	22.3	15.7	26.2	33.3	17.3	17.6	21.4	25.6	35.8
B ≻ C ≻ A	2.7	2.7	1.1	5.8	7.6	1.8	1.4	2.0	2.7	8.5
C ≻ A ≻ B	20.3	21.8	14.9	15.7	18.7	6.9	6.8	4.8	5.6	7.2
C ≻ B ≻ A	5.5	6.2	2.8	10.6	8.0	4.2	3.5	3.0	4.2	7.0
N	1456	1456	1456	1456	1456	3200	3200	3199	3200	3200

Panel B. Experiments 3, 4, and 5

	Experiment 3								Experiment 4						Experiment 5		
Ranking	H1	H2	H3	H4	E1	E2	E3	E4	B1	B2	B3	B4	B5	B6	N1	N2	N3
A≻ B ≻ C	29.5	17.6	39.5	35.8	20.3	23.1	19.0	22.3	11.0	11.1	5.8	18.9	16.3	12.7	25.7	13.4	28.1
A ≻ C ≻ B	22.8	9.9	26.0	36.9	17.1	11.7	17.7	21.3	11.2	7.9	2.9	11.9	16.7	9.7	12.7	6.1	14.3
B ≻ A ≻ C	9.9	27.7	10.5	10.2	30.5	27.0	26.3	18.6	12.5	16.3	14.8	12.5	13.2	15.5	24.9	26.1	20.4
B ≻ C ≻ A	6.9	13.6	5.8	2.8	5.0	13.3	15.2	13.4	12.7	9.7	15.0	5.7	7.1	10.3	8.1	16.2	8.2
C ≻ A ≻ B	20.3	16.0	9.8	11.5	23.9	13.3	13.8	12.4	32.9	28.0	14.6	35.1	33.6	27.3	18.5	17.0	18.1
C ≻ B ≻ A	10.7	15.2	8.3	2.9	3.2	11.7	8.0	12.1	19.7	27.0	46.8	15.9	13.1	24.6	10.1	21.2	11.0
N	1226	1197	1148	1192	1235	1183	1207	1211	1469	1469	1469	1469	1469	1469	1348	1301	1324

Notes: Each cell reports the share of respondents with the given strict ranking in that scenario. Scenario sample sizes shown in the last row of each panel.

4.2 Preference Structure

4.2.1 Single-Crossing

I first examine whether preferences satisfy the single-crossing property with respect to respondents' self-reported political ideology.

Identifying the left–right ordering of alternatives. In Experiment 1, the ordering of alternatives is pre-specified as *B–A–C*, which implies $18 \times 5 = 90$ single-crossing constraints

across the five scenarios. The same ordering is then carried over to Experiment 2, which follows the pre-specified design of Experiment 1. This ordering is also used as the baseline in the remaining experiments. However, alternative orderings are in some cases also defensible. To ensure that the results are not driven by the imposed ordering, I conduct a stricter test that evaluates single-crossing under all plausible rankings.

Online Appendix Section D.1 describes five complementary methods for classifying the ideological ordering of alternatives. Most methods identify *B* as the most left-leaning alternative in the majority of scenarios. However, in scenarios H1, H3, H4, B4, N1, and N3, *A* can also plausibly be classified as the most left-leaning option. Accordingly, I test single-crossing under both orderings, *A–B–C* and *B–A–C*, in these cases. This introduces six additional requirements per scenario, increasing the total number of constraints by 18 in Experiment 3, 6 in Experiment 4, and 12 in Experiment 5.

Requirement checks. Table 2 summarizes the results. Strikingly, none of the constraints is violated in any experiment. Of the 522 tested requirements, 174 are binding, and all of these are satisfied.

Table 2: Single-crossing constraints by experiment and assumed order

Scenario	# Constraints	Violated	Binding	Not binding
Exp1				
Exp1 (BAC)	90	0	29	61
Exp2				
Exp2 (BAC)	90	0	43	47
Exp3				
Exp3 (ABC)	18	0	9	9
Exp3 (BAC)	144	0	37	107
Exp4				
Exp4 (ABC)	6	0	3	3
Exp4 (BAC)	108	0	30	78
Exp5				
Exp5 (ABC)	12	0	6	6
Exp5 (BAC)	54	0	17	37
Total	522	0	174	348

A global randomization test. Table 2, however, does not indicate how often we would expect violations of the single-crossing requirements under a null in which political ideology does not structure shareholder preferences. For example, if all respondents shared the same ranking, there would be zero single-crossing violations even though ideology and ranking behavior were unrelated. To evaluate the results more formally, I therefore construct a global test that speaks directly to the requirement framework.

First, I define a test statistic V equal to the total number of *violated binding* single-crossing implications, summed over all relevant pairs and scenarios. In the observed data, $V_{\text{obs}} = 0$ in each experiment, and hence $V_{\text{obs}}^{\text{All}} = 0$ when summing across all scenarios.

To assess how unusual $V_{\text{obs}}^{\text{All}} = 0$ is under a null in which ideology does not structure preferences, I implement a permutation test. Within each experiment, I randomly reassign ideology labels across respondents while keeping each respondent's full ranking profile fixed. I then recompute V for each permuted dataset and repeat this procedure 5,000 times.

Under this permutation null, the median draw produces 10 violations, and $V^{\text{All}} = 0$ occurs in only 69 out of 5,000 permutations, corresponding to a randomization p-value of $(69+1)/(5000+1) \approx 0.014$. Thus, obtaining zero requirement violations simultaneously across all five experiments is unlikely under a null in which political ideology is unrelated to ranking behavior.

Result 1. *Preferences satisfy the single-crossing condition: across the five experiments, none of the 522 tested single-crossing requirements is violated. Permutation analysis shows that observing zero violations is unlikely ($p = 0.014$) if ideology does not structure preferences.*

Formal trend tests. As an additional check on preference structure, I complement the single-crossing requirement checks with two further formal tests that examine whether pairwise support tends to move monotonically with political ideology. For any ordered pair (x, y) with x to the left of y on the assumed axis, I test whether the probability of ranking x above y displays a monotone ideological gradient. For each scenario and relevant pairwise comparison, I therefore use (i) a one-parameter logit trend regression, reported as odds ratios, and (ii) a nonparametric test for ordered trends across ideology groups applied to the

corresponding binary pairwise indicator.

The full set of estimates is reported in Tables A7 and A8 in the Online Appendix. Overall, the results strongly support monotone pairwise gradients in the predicted direction. The logit estimates generally imply the expected ideological slope, given the coding of each pairwise comparison, and the nonparametric trend-test p-values are typically small, with only a small number of comparisons showing weak or near-zero gradients.

Table 3: Permutation test for zero violations of binding single-crossing requirements

Dataset	# Scenarios	V_{obs}	B_{obs}	$\Pr_{\text{perm}}(V \leq V_{\text{obs}})$	median(V)	V_{95}
Exp. 1	5	0	29	0.469	2	4
Exp. 2	5	0	43	0.720	0	4
Exp. 3	8	0	46	0.114	4	10
Exp. 4	6	0	33	0.351	2	6
Exp. 5	3	0	23	0.974	0	0
All experiments (sum)	27	0	174	0.014	10	16

Notes. V counts the number of violated *binding* single-crossing implications, aggregated across all relevant ordered pairs and scenarios; B counts the number of binding implications. Permutation draws randomly reassign ideology labels across respondents within each experiment, holding each respondent's full set of rankings fixed. Reported permutation p-values use the standard randomization correction $(\#\{V_{\text{perm}} \leq V_{\text{obs}}\}+1)/(R+1)$ with $R = 5000$ draws. For V is summed across experiments within each draw.

4.2.2 Single-Peakedness

With three alternatives, single-peakedness along a one-dimensional ordering implies a simple restriction: once a middle alternative is fixed ex ante, preferences are single-peaked if and only if no respondent ranks that middle alternative last.

This implication is already visible in Table 1. In Experiments 1 and 2, where the axis is pre-specified and A is the middle alternative, few respondents rank A last, consistent with single-peaked preferences.

Empirical implementation. Following the diagnostic in Section 2, I interpret the reported SP statistic as a measure of *approximate* single-peakedness in noisy data. For Experiments 1

Table 4: Single-peakedness results

Scenario	N	SP (%)	Boot. 95% CI	Middle agreement (%)	OOS median (%)	MC p-value
Experiment 1						
Alphabet	1456	91.8	[90.3, 93.1]	100.0	91.8	< 0.00005
Microsoft	1456	91.1	[89.7, 92.5]	100.0	91.1	< 0.00005
Amazon	1456	96.1	[95.1, 97.0]	100.0	96.1	< 0.00005
Meta	1456	83.6	[81.6, 85.4]	100.0	83.6	< 0.00005
Berkshire Hathaway	1456	84.3	[82.4, 86.2]	100.0	84.3	< 0.00005
Experiment 2						
Alphabet	3200	94.0	[93.2, 94.8]	100.0	94.0	< 0.00005
Microsoft	3200	95.1	[94.4, 95.8]	100.0	95.1	< 0.00005
Amazon	3199	95.0	[94.2, 95.7]	100.0	95.0	< 0.00005
Meta	3200	93.0	[92.1, 93.9]	100.0	93.0	< 0.00005
Berkshire Hathaway	3200	84.4	[83.2, 85.7]	100.0	84.4	< 0.00005
Experiment 3						
H1	1226	82.5	[80.2, 84.7]	100.0	82.5	< 0.00005
H2	1197	74.1	[72.0, 76.7]	84.0	74.1	< 0.00005
H3	1148	85.9	[83.7, 87.8]	100.0	85.9	< 0.00005
H4	1192	94.4	[93.0, 95.6]	100.0	94.4	< 0.00005
E1	1235	91.8	[90.2, 93.3]	100.0	91.8	< 0.00005
E2	1183	75.1	[74.2, 77.9]	5.5	74.4	< 0.00005
E3	1207	76.8	[74.6, 79.1]	100.0	76.8	< 0.00005
E4	1211	74.6	[72.1, 76.9]	99.5	74.6	< 0.00005
Experiment 4						
B1	1469	76.5	[74.3, 78.7]	100.0	76.5	< 0.00005
B2	1469	72.6	[70.4, 74.9]	100.0	72.6	< 0.00005
B3	1469	82.4	[80.7, 84.4]	93.0	82.4	< 0.00005
B4	1469	78.4	[76.3, 80.5]	100.0	78.4	< 0.00005
B5	1469	79.8	[77.8, 81.8]	100.0	79.8	< 0.00005
B6	1469	71.8	[69.5, 74.0]	99.5	71.8	0.00005
Experiment 5						
N1	1348	81.8	[79.7, 83.9]	100.0	81.8	< 0.00005
N2	1301	76.9	[74.6, 79.0]	100.0	76.9	< 0.00005
N3	1324	80.8	[78.7, 82.8]	100.0	80.8	< 0.00005

Notes: For Experiments 1 and 2, the middle alternative is fixed ex ante at A, so SP is $100 \times (1 - p_A)$, where p_A is the sample share placing A last. For Experiments 3 to 5, the middle alternative is selected endogenously to maximize fit, so SP is $100 \times (1 - \min\{p_A, p_B, p_C\})$. Bootstrap confidence intervals are percentile intervals based on 2,000 resamples. Middle agreement reports the percentage of random splits in which the selected middle alternatives overlap across the two half-samples. In fixed-axis rows (Experiments 1 and 2), middle agreement is mechanically 100. OOS median reports the median across 200 random splits of the cross-fitted holdout fit: the middle alternative is chosen in one half-sample and evaluated on the other half, averaging both directions. In fixed-axis rows (Experiments 1 and 2), OOS median equals the split-sample fit. MC p-values are based on 20,000 simulations under the null of uniformly random rankings, using the same axis-selection rule as in the corresponding row.

and 2, where A is pre-specified as the middle alternative, the statistic is simply one minus the share of respondents who rank A last. For Experiments 3 to 5, where no ex ante axis is imposed, I instead use the best-fitting-axis statistic $\widehat{\text{MaxSP}} = 1 - \min\{\hat{p}_A^{last}, \hat{p}_B^{last}, \hat{p}_C^{last}\}$, which chooses as the middle alternative the option that is least often ranked last.

Because this best-fitting procedure mechanically improves fit, inference in Experiments 3 to 5 must account for the same selection rule. I therefore use Monte Carlo simulations under a sharp null of random rankings, uniformly distributed across the six strict orderings. For each scenario and sample size, I simulate many datasets under this null and recompute the same statistic using the same axis-selection rule as in the observed data. The resulting p-value therefore asks how often a degree of fit at least as large as the observed one would arise by chance under unstructured rankings.[7]

Table 4 shows strong support for approximate single-peakedness in Experiments 1 and 2 under the pre-specified A-middle axis. The SP statistic ranges from 83.6 to 96.1 percent in Experiment 1 and from 84.4 to 95.1 percent in Experiment 2. In Experiments 3 to 5, where the middle alternative is selected to maximize fit, approximate single-peakedness ranges from 71.8 to 94.4 percent across the 17 scenarios. The Monte Carlo p-values are very small in all scenarios, indicating that the observed degree of fit is far greater than would be expected under random rankings.

Result 2. *Across scenarios, preferences exhibit strong approximate single-peakedness in the disciplined fixed-axis settings of Experiments 1 and 2, and moderate-to-strong approximate single-peakedness in the more flexible settings of Experiments 3 to 5. In all cases, the observed levels are far too large to be explained by random rankings.*

[7]Table 4 also reports bootstrap confidence intervals and simple split-sample stability checks. The bootstrap intervals quantify sampling uncertainty in the estimated fit. The split-sample diagnostics ask whether the same middle alternative tends to be selected across random halves of the sample and whether the selected axis continues to fit well out of sample, thereby providing a check against overfitting.

Table 5: WTP-based welfare analysis for WTP-elicited scenarios in Experiment 1(Berkshire) and Experiment 3 (E1–E3).

	Berkshire	E1	E2	E3
Panel A: WTP-based welfare loss				
Plurality	29.5 [12.2, 48.5]	19.6 [0.0, 72.9]	8.5 [0.0, 27.7]	5.4 [0.0, 24.7]
Borda	14.4 [0.0, 42.3]	0.0 [0.0, 0.0]	7.5 [0.0, 25.6]	2.3 [0.0, 16.6]
Condorcet	20.6 [0.0, 44.7]	0.0 [0.0, 0.0]	6.2 [0.0, 25.2]	2.8 [0.0, 18.6]
IRV (Runoff)	22.4 [0.0, 46.4]	0.0 [0.0, 0.0]	6.2 [0.0, 25.2]	2.8 [0.0, 18.6]
Panel B: Rule selects a non-WTP-maximizing outcome				
Plurality	99.7	32.3	61.9	51.6
Borda	51.8	0.0	58.0	30.3
Condorcet	70.6	0.0	49.6	33.6
IRV (Runoff)	74.1	0.0	49.6	33.6
Panel C: Rule selects alternative B				
Plurality	100.0	21.8	98.6	97.1
Borda	3.0	0.0	91.5	19.3
Condorcet	20.8	0.0	90.2	35.2
IRV (Runoff)	20.9	0.0	90.2	35.2

Notes: Panels A and B use the matched WTP sample. Panel A reports mean normalized WTP-based welfare loss across bootstrap replications in which the rule yields a unique winner; brackets show 95% percentile intervals from the same nonmissing bootstrap draws. Thus, no-unique-winner replications are omitted from Panel A. Panel B reports the percentage of bootstrap replications in which the rule selects a unique winner that does not maximize WTP in that replication. Bootstrap replications with no unique winner are not counted as wrong in Panel B. Panel C uses the full sample and reports the percentage of bootstrap replications in which each rule selects alternative B, the pole option in these scenarios. Bootstrap resampling is at the respondent level, conducted within WTP scenario for Panels A and B and within experiment for Panel C.

4.3 Aggregation Rules

I next turn to the paper's second question: how different voting rules translate the observed rankings into collective outcomes. Appendix Table A1 reports the point-estimate winner under plurality, Borda, Condorcet, and instant-runoff voting (IRV) in every scenario. In most scenarios, the four ordinal rules select the same winner.

Point estimates alone, however, do not reveal how robust a given outcome is. Two rules may both select alternative A in the full sample, yet differ sharply in how sensitive that result is to small perturbations in the observed rankings. If a small change in the sample frequently overturns the winner under one rule but not under another, the first rule rests on more fragile support. The analysis below therefore focuses not only on point estimates, but also on respondent-level bootstrap winner frequencies, which show how often each rule selects each outcome under repeated resampling of the observed electorate.

4.3.1 Bootstrap design and inference

To evaluate the stability of the voting-rule results, I use a respondent-level clustered bootstrap. In each replication, I resample respondents with replacement and reconstruct the implied ranking distribution for each scenario. I then reapply each voting rule—plurality, Borda, Condorcet, and IRV—to the resampled electorate and record the resulting winner. This procedure serves two purposes. First, it provides a natural measure of outcome stability: for each scenario and rule, the bootstrap frequencies show how often a given alternative is selected across nearby realizations of the sample. Second, it is well suited to winner-based objects, which are inherently non-smooth: small changes in ranking shares can discretely change the winning alternative, making standard linear approximations unreliable. Formally, let $b = 1, \ldots, B$ index bootstrap replications, with $B = 5{,}000$. In each replication, I resample respondents at the individual level and recompute: (i) the ranking shares in each scenario, (ii) the winner under each voting rule, and (iii) whether pairwise majority yields a Condorcet winner or a Condorcet cycle. Throughout, a rule is treated as selecting an outcome only when it has a unique winner. In Experiment 1 and in the consequential scenarios E1–E4 of

Experiment 3, I additionally recompute respondents' willingness-to-pay (WTP) aggregates in each bootstrap replication. This allows me to compare not only which alternative each rule selects, but also how closely each rule tracks the cardinal benchmark introduced in Section 4.4. The welfare bootstrap is implemented on the matched WTP sample within each WTP-elicited scenario, so that the ordinal winner and the cardinal benchmark are computed from the same resampled respondent pool.

4.3.2 Outcome Frequencies Across All Five Experiments

Appendix Table A2 reports bootstrap outcome frequencies for all 27 scenarios across Experiments 1–5. A central pattern is that, in most scenarios, the selected outcome is stable and invariant across voting rules. When one alternative is sufficiently well supported, plurality, Borda, Condorcet, and IRV typically yield nearly identical outcomes. The analysis therefore focuses on the subset of scenarios in which support is more evenly divided and the choice of voting rule becomes consequential.

4.3.3 Where Rule Choice Matters

Table 5, Panel C, highlights the scenarios in which the voting rules meaningfully disagree among those with WTP elicitation, which are the Berkshire scenario in Experiment 1 and scenarios E1–E3 in Experiment 3. The panel reports the bootstrap frequency with which each rule selects alternative B, which is the option most likely to differ across rules in these cases.

The Berkshire scenario in Experiment 1 provides a clear illustration. Plurality selects B in 100% of bootstrap replications, whereas Borda selects B in only 3.0% of replications, and Condorcet and IRV do so in 20.8% and 20.9%, respectively. The ranking distribution in Table 1 explains this divergence. Alternative B has the largest first-place support and therefore wins under plurality, but A has substantially broader support once lower-ranked positions are taken into account. In the full sample, 18.7% of respondents have ranking $C \succ A \succ B$, whereas only 8.0% have $C \succ B \succ A$. Likewise, 33.3% have ranking $B \succ A \succ C$, whereas only 7.6% have $B \succ C \succ A$. Thus, both respondents who most prefer C and respondents who most

prefer B tend to place A ahead of the opposite extreme. Plurality ignores this information and therefore favors the alternative with the strongest first-place support even when another alternative is more broadly acceptable across the electorate.

The same mechanism appears in scenario E1, which was designed to include a designated middle alternative, A. Here first-place support for A and B is relatively close: 37.4% of respondents rank A first and 35.5% rank B first. Even so, plurality selects B in 21.8% of bootstrap replications. Once the full ranking is taken into account, however, support for A is clearly broader. In the full sample, 23.9% of respondents have ranking $C \succ A \succ B$, whereas only 3.2% have $C \succ B \succ A$. Likewise, 30.5% have ranking $B \succ A \succ C$, whereas only 5.0% have $B \succ C \succ A$. Accordingly, Borda, Condorcet, and IRV all select A, and in the bootstrap they do so in 100% of replications.

Scenario E3 shows the same logic in a somewhat noisier form. Here B receives the most first-place support and plurality therefore selects it in 97.1% of bootstrap replications. By contrast, Borda selects B in only 19.3% of replications, while Condorcet and IRV do so in 35.2%. Again, the reason is that A attracts broader second-place support than plurality can use. In the full sample, 13.8% of respondents have ranking $C \succ A \succ B$, compared with 8.0% who have $C \succ B \succ A$. Similarly, 26.3% have ranking $B \succ A \succ C$, compared with 15.2% who have $B \succ C \succ A$. Thus, although A trails B in first-place votes, it is more broadly supported as a compromise outcome.[8]

The broader lesson is that plurality systematically disadvantages compromise candidates because it uses only first-place support and ignores the rest of the ranking. This matters especially in environments where preferences are polarized. When voters on opposing sides tend to rank an intermediate option second, plurality can select an extreme outcome even though the compromise commands broader overall support in the full ranking profile.

[8]Scenario E2 is somewhat different. There, all four ordinal rules mostly select B. However, the same mechanism is still present in weaker form: Borda, Condorcet, and IRV give A a non-trivial probability of winning, whereas plurality almost never does.

4.3.4 Condorcet Cycles

A potential concern with Condorcet-based aggregation is the possibility of cycles. In the data, however, this concern is empirically negligible. There are no Condorcet cycles in the point estimates in any of the 27 scenarios, and Appendix Table A2 shows that the bootstrap frequency of strict Condorcet cycles is essentially zero throughout. In other words, respondent-cluster bootstrap replications almost always yield a well-defined Condorcet winner.

Result 3. *Strict Condorcet cycles are absent, or empirically negligible, across scenarios.*

4.4 Welfare Effects of Voting Rules

The previous subsection examined when ordinal voting rules agree with one another, but did not assess their relative performance. This section addresses a different question: in the subset of scenarios where cardinal information is available, which rule selects the WTP-maximizing outcome? The welfare analysis therefore focuses on the scenarios with WTP elicitation in which meaningful differences arise, namely the Berkshire scenario from Experiment 1 and scenarios E1–E3 in Experiment 3.

Welfare benchmark and loss measure. For each respondent, the WTP data provide intensity information about their ranking. I map these data into respondent-level utilities by normalizing the third-ranked alternative to zero, assigning the top-ranked alternative WTP_f (first vs. third), and assigning the second-ranked alternative WTP_s (second vs. third). Averaging these respondent-level utilities across the matched WTP sample yields a valuation for each alternative. Let μ_j denote the resulting valuation of alternative j, and let $\widetilde{\text{WTP}}_j = \mu_j - \min_{k \in X} \mu_k$ denote the normalized aggregate WTP of alternative j, so that the lowest-valued alternative in the choice set has normalized value zero. Let $\widetilde{\text{WTP}}^* = \max_{k \in X} \widetilde{\text{WTP}}_k$ denote the highest normalized aggregate WTP in the sample. For a voting rule that selects alternative j, I define normalized welfare loss as $\text{WL} = 100 \times \frac{\widetilde{\text{WTP}}^* - \widetilde{\text{WTP}}_j}{\widetilde{\text{WTP}}^*}$. Under this normalization, a welfare loss of 0 indicates that the rule selects a WTP-maximizing alternative,

while a welfare loss of 100 indicates that it selects the lowest-WTP alternative in the choice set. Welfare loss is computed only for bootstrap replications in which the rule yields a unique winner.

Bootstrap inference for welfare. Because the identity of the WTP-optimal alternative and the voting-rule winner can change discontinuously with small perturbations in the sample, I again use the respondent-cluster bootstrap. In each of 5,000 replications, I resample respondents, recompute aggregate WTP by alternative, determine the WTP-optimal benchmark i^*, apply each voting rule, and compute the implied welfare loss. The benchmark i^* is recomputed in every replication rather than fixed from the original sample. I then report mean welfare loss across replications together with 95% bootstrap percentile intervals.

4.4.1 Main Welfare Results

Table 5 focuses on the four scenarios in which non-trivial differences across voting rules arise. Panels A and B are based on the matched WTP sample, while Panel C reports the corresponding ordinal outcome frequencies in the full sample to illustrate the underlying ranking patterns.

Plurality is consistently the worst-performing rule in these scenarios. In the Berkshire scenario, plurality yields a mean welfare loss of 29.5% of WTP-optimal welfare and selects a non-WTP-maximizing outcome in 99.7% of bootstrap replications. By contrast, Borda reduces the mean loss to 14.4%, while Condorcet and IRV reduce it to 20.6% and 22.4%, respectively. In scenario E1, plurality generates a mean welfare loss of 19.6% and selects a non-WTP-maximizing outcome in 32.3% of replications, whereas Borda, Condorcet, and IRV all coincide with the WTP benchmark and yield zero welfare loss in every replication. In E3, plurality again performs worst, with a mean welfare loss of 5.4%, compared with 2.3% for Borda and 2.8% for both Condorcet and IRV.

The earlier analysis showed that plurality differs from the other rules because it relies solely on first-place support and ignores the remainder of the ranking. The welfare results

demonstrate that this difference is economically meaningful: in the scenarios where aggregation is non-trivial, plurality systematically selects outcomes that are further from the WTP benchmark than those selected by ranking-based rules.

Scenario E2 provides a complementary insight. In that case, all four ordinal rules typically select the same alternative, B, when only ordinal information is considered. However, the WTP analysis reveals that A yields higher aggregate welfare. Thus, even when ordinal rules agree, they need not recover the cardinally optimal outcome.

Result 4. *In the WTP-elicited scenarios, non-trivial welfare differences arise in four out of nine cases. In these scenarios, plurality systematically yields the largest expected welfare loss. Ranking-based rules—Borda, Condorcet, and IRV—consistently perform better, although no ordinal rule fully matches the WTP benchmark in all cases.*

Which voting rule performs best? The evidence so far establishes that plurality is the weakest rule in the non-trivial cases, but the relative performance of Borda, Condorcet, and IRV remains to be systematically compared. Across these three ranking-based rules, no single rule dominates in every scenario. In E1, all three coincide and recover the WTP benchmark in every bootstrap replication. In Berkshire and E3, however, Borda performs best: Appendix Table A3 shows that, conditional on producing different welfare outcomes, Borda almost always outperforms Condorcet and IRV in Berkshire and does so in a clear majority of replications in E3. By contrast, in E2, Condorcet and IRV perform best and each outperforms Borda when they differ. More generally, when Condorcet and IRV select different outcomes, Condorcet weakly outperforms IRV in the pairwise welfare-dominance comparisons.

5 Experiment 6: Adding the joint-support bundle

5.1 Design

Experiment 6 revisits the five algorithmically constructed firm scenarios from Experiment 1, but adds the previously omitted outcome D: supporting both shareholder proposals. Respondents therefore rank the full four-outcome bundle set. The sample consists of 980 U.S. respondents who report holding stocks or mutual funds, matching the investor-screened population in Experiment 1. The goal is to test whether the preference structure documented in the three-outcome design survives once the full 2×2 bundle set is available. Unlike Experiment 1, this follow-up does not include WTP elicitation, so the analysis here focuses on robustness of the preference-structure results rather than welfare comparisons.

I retain the ideology-based shareholder order from the main analysis. For the single-crossing diagnostic, I extend the baseline substantive ordering to $D \prec B \prec A \prec C$; this ordering is used only to orient pairwise comparisons. With four alternatives there are $\binom{4}{2} = 6$ unordered pairs, and applying the same six implication checks per pair as in the main three-group diagnostic yields $6 \times 6 = 36$ single-crossing requirements per firm scenario, or a total of 180 requirements. For generalized single-peakedness, I instead use the natural square geometry of the bundle space, where $C = (0,0)$, $A = (1,0)$, $B = (0,1)$, and $D = (1,1)$, with coordinates indicating support for the social and climate proposal. On this space, Nehring–Puppe consistency is equivalent to issue separability: $A \succ C \iff D \succ B$ and $B \succ C \iff D \succ A$. With four outcomes there are $4! = 24$ possible strict rankings, but only eight satisfy both issue-separability conditions. I report the shares of respondent-scenario rankings that satisfy both conditions, exactly one, or neither.

5.2 Results

Single-crossing. The single-crossing pattern survives the expansion of the choice set. Across the five firm scenarios, 64 of the 180 implication checks are binding and none is violated.

Nehring–Puppe consistency. Table 6 reports consistency with the square-geometry benchmark. The share of NP-consistent rankings ranges from 60.6% to 70.2% across the five scenarios, with a pooled share of 64.5%. Most departures are limited rather than wholesale: 23.0% of pooled rankings violate exactly one separability condition, whereas 12.5% violate both. Thus, adding the omitted joint-support bundle introduces some departures from strict separability, but a clear majority of rankings remains consistent with the natural two-dimensional bundle structure. Overall, Experiment 6 shows that the main pattern in Experiment 1 is not an artifact of excluding $D = (1,1)$. Once the full four-outcome bundle space is admitted, preferences remain highly structured: single-crossing continues to hold, and most rankings remain consistent with the natural two-dimensional bundle structure.

Table 6: Nehring–Puppe consistency in the four-outcome design

Scenario	NP consistent	Violates one NP condition	Violates two NP conditions
Alphabet	70.2%	19.4%	10.4%
Microsoft	67.1%	20.7%	12.2%
Amazon	62.8%	24.2%	13.0%
Meta	61.9%	24.1%	14.0%
Berkshire	60.6%	26.6%	12.8%
Pooled	64.5%	23.0%	12.5%

Notes: Each observation is a respondent–scenario ranking over the four bundle outcomes $C = (0,0)$, $A = (1,0)$, $B = (0,1)$, and $D = (1,1)$. A ranking is classified as *NP-consistent* if both issue-separability conditions implied by the square geometry hold, as violating one NP condition if exactly one fails, and as violating two NP conditions if both fail. Shares are reported as percentages of respondent–scenario observations.

6 Experiment 7

Experiment 7 was conducted in January 2026 through Prolific, and is preregistered as AsPredicted 269073. This experiment contained three preregistered modules. This section

focuses on the two modules that directly extend the paper's central aggregation question to richer choice sets and board-election settings. A third module studies a complementary implementation question—how pass-through voting through delegated voting profiles affects realized proposal-level votes relative to direct voting under participation costs—and is reported in Online Appendix C.

6.1 Part 1: Extending Choice Sets

Experiments 1 to 5 study voting-rule performance in settings with three feasible outcomes. This experiment extends the choice set from three to six outcomes by allowing respondents not only to prioritize a single objective, but also to select policies that combine two objectives. I focus on a scenario inspired by E1 from Experiment 3, in which respondents evaluate how an energy firm should trade off profit (C), a climate objective (B), and a social/geopolitical objective (A). In this version, respondents rank six policy options: the three core objectives and three additional options that jointly prioritize two of the three dimensions (e.g., $B+C$). They are asked to provide a complete ranking over all six alternatives. This richer choice environment allows me to examine whether the preference structure documented in the main analysis persists when the set of feasible outcomes is expanded. After completing the ranking task, respondents also provide cardinal evaluations of each of the six policy options on a 0–10 scale. Unlike in the main welfare analysis, these evaluations are not tied to a willingness-to-pay benchmark and therefore do not admit a direct monetary interpretation. Nevertheless, they provide a useful proxy for preference intensity and allow for a welfare comparison across voting rules using the same logic as in Experiments 1–5.

Policy ordering and single-crossing. To study whether preferences remain structured in this richer environment, I again impose a left-to-right policy ordering and test for single-crossing with respect to political ideology. As elsewhere in the paper, alternatives involving profit (C) are naturally located furthest to the right, while the climate option B is placed to the left. The combination options are also naturally interpreted as intermediate compromises rather than additive bundles, since the instructions make clear that placing greater emphasis

on one goal typically requires placing less emphasis on the others. Under this interpretation, there are two plausible one-dimensional orderings of the six policies, differing only in the position of $B+C$ relative to A. To be cautious, I therefore test for single-crossing under both orderings.[9] As in the main analysis, respondents are grouped into three ordered ideology categories: Left, Middle, and Right. Single-crossing is then tested using pairwise-majority comparisons, as in Experiments 1 to 5. With six outcomes, there are $\binom{6}{2} = 15$ ordered pairs for any given policy ordering, and each pair yields six implication requirements, for a total of 90 requirements. Since the two plausible orderings differ only in the relative position of A and Profit+B, they share 84 requirements and differ on 12, for a total of 96 distinct requirements when both orderings are considered jointly.

6.1.1 Results

Single-crossing In the data, 34 of the single-crossing requirements are binding and 62 are non-binding, and there are zero violations. Thus, even in this richer six-outcome setting, the majority preference profile remains fully consistent with single-crossing under both plausible ideology orderings.

Voting rules and welfare. I next compare voting-rule outcomes to a cardinal benchmark, following the same logic as in the main analysis. Let μ_j denote the mean score on the cardinal question for option j, and let $\mu^\star = \max_j \mu_j$ denote the welfare benchmark. Welfare loss is then defined, as before, as $100 \times (\mu^\star - \mu_j)/\mu^\star$.

In the data, the option $B+C$ maximizes mean scores, with $\mu^\star = 7.585$, and is also the unique Condorcet winner. Borda and IRV select this option as well. By contrast, plurality selects $A+B$, which has a substantially lower mean score.

To assess robustness, I implement a respondent-level clustered bootstrap with 5,000 replications, as in the analysis underlying Experiments 1 to 5. Plurality selects the welfare-optimal option in only 26% of draws. By contrast, Borda, Condorcet, and IRV almost always select the welfare-optimal option. As a result, the bootstrap mean welfare loss is 5.9% under

[9] The two orderings are $B \prec A{+}B \prec A \prec B{+}C \prec A{+}C \prec C$ and $B \prec A{+}B \prec B{+}C \prec A \prec A{+}C \prec C$.

plurality and essentially zero under the other voting rules. The intuition for for why plurality fails is the same as previously, but in some sense extending the set of outcomes makes it easier for a more extreme alternative to win, as a lower share of support is necessary to win the plurality vote in this case.

Taken together, these results show that the main findings of the paper extend to richer choice environments. Even when the number of policy options is doubled, preferences remain highly structured in the sense of single-crossing, and voting rules that use more ordinal information continue to select outcomes that are close to the welfare benchmark, while plurality performs substantially worse.

Table 7: Voting-rule outcomes and welfare in the six-option module

Rule	Full-sample winner	$P(\text{hit})$ (%)	Mean WL (%)	WL p2.5–p97.5 (%)
Plurality	A+B	26.06	5.932	[0.00, 9.41]
Borda	B+C	100.00	0.00	[0.00, 0.00]
Condorcet	B+C	100.00	0.00	[0.00, 0.00]
IRV	B+C	99.66	0.007	[0.00, 0.00]

Notes: Full-sample winner reports the point-estimate winner in the full sample. $P(\text{hit})$ is the share of respondent-cluster bootstrap replications in which the rule selects a welfare-optimal option. Mean welfare loss (WL) and percentile intervals are computed from the bootstrap distribution of welfare losses across replications in which the rule yields a unique winner. Welfare is benchmarked using the cardinal 0–10 evaluation question.

6.2 Part 2: Board module

The second part of the survey studies board elections. The introduction of the universal proxy card allows investors to express preferences in settings that would otherwise often be perceived as binary. The purpose of this module is to examine how this institutional change transforms preferences over board control from a seemingly binary choice into a richer outcome space, and then to study whether plurality-at-large, the rule typically used in such elections, is able to aggregate these richer preferences in a meaningful way.

I consider a simplified version of the Exxon versus Engine No. 1 proxy contest in 2021, with four open board seats and eight candidates, of which four are company nominees and four activist nominees. The activist nominees are shown using the actual names of the Engine No. 1 nominees, together with a short description of the perspective each would bring to the board. The company nominees, by contrast, are labeled Company Nominee A to D. Respondents are asked to select up to four candidates, and told that the four candidates receiving the most votes are elected under a standard plurality-at-large rule.

After the candidate-selection task, respondents are asked to rank three salient outcome bundles: a *Company* bundle with zero activist directors, a *Compromise* bundle with two activist directors, and an *Activist* bundle with four activist directors.[10]

Results The candidate-selection task shows that preferences over board composition are not well described as a binary company-versus-activist split. Among respondents casting a valid ballot, only 31.8% choose an extreme outcome—that is, either zero or four activist nominees—whereas 68.2% choose mixed boards with one to three activist directors. The modal ballot corresponds to exactly two activist directors, chosen by 28.9% of respondents. Despite this, plurality-at-large elects the all-activist slate, because the four activist nominees receive most votes. In other words, candidate-level aggregation produces an extreme board even though most respondents do not cast an extreme ballot.

The subsequent bundle-ranking task makes this disconnect even clearer. The *Compromise* bundle is the modal first choice, receiving 49.8% of first-place rankings, compared with 20.8% for the *Company* bundle and 29.5% for the *Activist* bundle. Preferences are also highly structured. Fully 94.4% of respondents are single-peaked on the natural Company $\prec$ Compromise $\prec$ Activist axis, and the ideology-based single-crossing diagnostic yields seven binding implications and zero violations. The *Compromise* bundle also defeats both alternatives in pairwise majority voting, and all standard aggregation rules applied to the bundle rankings select the *Compromise* bundle.

[10]The compromise bundle is not arbitrary, as Glass Lewis recommended support for these two activist nominees rather than full board turnover (Reuters, 2021).

Taken together, the results show that once investors can express preferences over board composition, those preferences are neither binary nor unstructured. Instead, they are centered on an intermediate outcome: the compromise bundle is the most commonly preferred allocation of board seats. The aggregation problem arises at the candidate level. Under plurality-at-large, votes are cast for individual nominees rather than for bundles, and support for compromise outcomes is therefore dispersed across multiple candidates. As a result, a majority that prefers partial activist representation can still produce an extreme outcome when votes are aggregated nominee by nominee. The key implication is that when the object of choice is a bundle, such as board composition, candidate-level plurality need not recover the majority-preferred outcome.

7 Experiment 8-10

A key question arising from the previous experiments is how investors and practitioners think preferences should be aggregated when the preference profile is not well structured. Classical results in social-choice theory (Arrow, 1951) demonstrate that no voting rule can perfectly translate preferences into a single outcome. However, in real-world settings, institutional investors voting on behalf of clients must make decisions even when preferences lack a structure that facilitates straightforward aggregation. To explore how aggregation rules are applied in such cases, Experiments 8-10 examine how both practitioners and laypeople navigate scenarios involving potential preference cycles.

7.1 Experiment 8: Design

Experiment 8 was conducted in August 2024 with $N = 16$ Scandinavian institutional investors and mutual fund managers, managing an average of more than $1 billion in assets.[11] Because the survey was anonymous, I cannot verify for each respondent whether they personally

[11] As data collection relied on direct personal contact with potential respondents and the sample size was unknown, preregistration was not conducted. Respondents had previously answered a survey on green shareholder voting as part of a separate research project.

execute proxy votes or instead work in organizations where voting is handled by a dedicated stewardship team. The questionnaire was administered via Qualtrics, presenting respondents with a modified version of Scenario E1 from Experiment 3. They were asked to choose among three alternatives: (i) C: exporting coal as long as it remains profitable, (ii) B: ceasing all exports, and (iii) A: exporting coal only when it supports geopolitical stability. The labels A, B, and C differ from those in the actual experiment (see Online Appendix E). However, the notation from previous sections is retained for consistency and clarity.

Respondents were then informed that three clients (Anne, Kari and Knut), each holding an equal stake in the fund, had conflicting preference rankings; Anne had ranking $C \succ A \succ B$, Kari had ranking $A \succ B \succ C$, while Knut had ranking $B \succ C \succ A$. By construction, these rankings do not produce a Condorcet winner. Notably, alternative A represents a "middle" position in line with E1, implying that Knut's ranking is *not* single-peaked.

First, participants were asked: *"How would you vote on behalf of your three clients in this situation?"*. They were given four response options: vote for A, vote for B, vote for C, or abstain from voting/unsure.

Next, they were asked: *"Which of the following strategies would you use to decide how to vote on behalf of the owners in such cases?"*. They were given the following alternatives:

- *"If it is difficult to determine what the majority truly prefers, it feels right to vote for the option that provides the highest profitability."*
- *"If it is difficult to determine what the majority truly prefers, I would abstain from voting."*
- *"If it is difficult to determine what the majority truly prefers, I would choose the option I personally prefer."*
- *"If it is difficult to determine what the majority truly prefers, I would try to assess the strength of the owners' preferences and choose the option with the highest average."*

The survey concluded with an open-ended question: *"Assume you voted for either A, B, or C, and were later asked to explain your decision. How would you justify your choice?"*

7.2 Experiment 9: Design

Experiment 9 extends the practitioner study in Experiment 8 to a Norwegian lay sample. The experiment was conducted in March and April 2026, using Norstat Norway, and was pre-registered as AsPredicted #276586. The sample consists of 500 respondents recruited to be representative of the Norwegian adult population. Respondents received the same scenario and the same two questions described in the preceding section for Experiment 8.

7.3 Experiment 10: Design

Experiment 10 was conducted in January 2025 using Prolific, with a sample of $N = 968$ respondents, broadly representative of the U.S. population in terms of political affiliation, gender, and age. Data collection was preregistered (AsPredicted #206840), with details available in Online Appendix F. Participants were presented with the same scenario as in Experiment 8. Since they were not institutional fund managers, participants were asked to *imagine* making voting decisions on behalf of a mutual fund. Additionally, the names of the hypothetical clients were changed to Anne, Peter, and Mary to better resonate with a U.S. audience.

7.4 Results

Experiment 8 In Experiment 8, 44% of fund managers indicated uncertainty or preferred to *abstain* rather than select one of the three proposed alternatives. In contrast, 31% supported the "middle" option (A), while 13% chose the profit-oriented option (C) and another 13% selected the climate option (B). When asked which strategy they would use to aggregate client preferences in a cyclical setting, 38% stated they would assess the intensity of each client's preferences and select the option with the highest overall 'average'. Meanwhile, 25% would abstain, another 25% would choose the profit-maximizing outcome, and 13% would vote for their own most-preferred alternative. Open-ended responses aligned with these patterns: 38% cited efforts to achieve consensus or an 'average-based' decision, 13% described adopting

a more paternalistic approach (e.g., applying their own judgment), and 25% emphasized profit maximization in the presence of cyclical preferences. These open-ended justifications align closely with the selected aggregation principles: respondents choosing the profitability rule invoke return maximization, those selecting the average-based rule refer to consensus or average reasoning, and those choosing the own-preference rule invoke personal judgment.

Experiment 9 In the Norwegian lay sample in Experiment 9, 32% of respondents chose the middle alternative, 28% selected the profit-oriented alternative, 12% chose the no-export alternative, and 28% were unsure or would abstain. Thus, relative to the practitioner sample in Experiment 8, Norwegian lay respondents were less likely to abstain and somewhat more likely to select one of the substantive alternatives. When asked about the decision principle they would use in such cases, 42% stated that profitability should prevail when it is difficult to determine what the majority truly prefers, 35% said they would assess the strength of investors' preferences and choose the alternative with the highest average, 22% would choose their own most-preferred alternative, and 9% would abstain.

Experiment 10 In the U.S. sample in Experiment 10, 38% of participants favored the profit-maximizing option (C), 35% chose the middle alternative (A), 9% selected the no-export policy (B), and 17% opted to abstain. Respondents' concrete voting choices are strongly related to the decision principles they endorse. Those who state that profitability should prevail are far more likely to vote for the profit-maximizing option, those who say they would abstain are far more likely to abstain, and those who favor average-based aggregation are much more likely to choose the middle option.

Political orientation was strongly correlated with voting behavior. Among self-identified Democrats, 26% selected profit (C), 39% selected the middle option (A) , whereas 17% chose no export (B). Among Republicans, 55% favored profit (C), 27% opted for the middle policy (A), and only 3% picked the climate alternative (B). While partisan differences in preferences are expected, this finding highlights how even the choice of *aggregation rule* may be shaped by political leanings. When asked about their preferred strategy for resolving

cyclical preferences, 50% of Republicans agreed that *profit should prevail* in the absence of a clear majority, compared to 27% of Democrats. Conversely, 32% of Democrats indicated they would assess the strength of each stakeholder's preferences, compared to 20% of Republicans. Similarly, 32% of Democrats—versus 19% of Republicans—reported they would select their own most-preferred outcome. Collectively, these findings suggest that individuals' political beliefs influence how they conceptualize the aggregation of shareholders' cyclical preferences.

8 Conclusion

This paper shows that shareholder preferences are often much more structured than the multidimensional nature of corporate conflict might suggest. That finding is important because such structure is a prerequisite for shareholder democracy to aggregate preferences coherently once voting power shifts from intermediaries to ultimate investors. In that sense, the results offer a cautiously optimistic message: even when investors disagree over profit, climate, and other non-pecuniary objectives, decentralized shareholder voting need not lead to chaotic or unstable outcomes.

At the same time, the paper shows that coherent aggregation does not follow mechanically from giving investors more voice. Institutional design matters. Once shareholders can express more than a binary yes-or-no choice, the voting rule becomes consequential. In the paper's empirical settings, plurality is the rule most likely to miss broadly supported compromise outcomes and to produce more extreme winners, because it ignores information contained in the rest of the ranking. By contrast, ranking-based rules such as Borda, Condorcet, and runoff perform substantially better and come much closer to the cardinal benchmark.

The broader implication is that the future of shareholder democracy will depend not only on whether ultimate investors are given a greater voice, but also on how that voice is delegated, expressed, and aggregated. As institutional changes increasingly allow investors to express non-binary preferences—most clearly in contested director elections under universal proxy and in other multi-option shareholder votes—the design of aggregation rules will play

an increasingly important role in determining whether shareholder voting recovers broadly supported outcomes or instead mechanically rewards vote-splitting and extremes. While the paper shows that coherent aggregation of shareholder preferences is possible in the settings studied here, an important direction for future research is to examine how these insights generalize to environments where information aggregation, strategic voting, and different types of investors play a more central role.

References

Alphabet Inc., "Definitive Proxy Statement (Schedule 14A) – 2025 Annual Meeting of Stockholders," U.S. Securities and Exchange Commission (SEC) EDGAR 2025. Accession No. 0001308179-25-000511. Accessed 2025-12-18.

_ , "Form 8-K (Item 5.07) – Voting Results of the 2025 Annual Meeting," U.S. Securities and Exchange Commission (SEC) EDGAR 2025. Accession No. 0001193125-25-139948. Accessed 2025-12-18.

Amazon.com, Inc., "Definitive Proxy Statement (Schedule 14A) – 2025 Annual Meeting of Shareholders," U.S. Securities and Exchange Commission (SEC) EDGAR 2025. Accession No. 0001104659-25-033442. Accessed 2025-12-18.

_ , "Form 8-K (Item 5.07) – Voting Results of the 2025 Annual Meeting," U.S. Securities and Exchange Commission (SEC) EDGAR 2025. Accession No. 0001104659-25-052014. Accessed 2025-12-18.

Ambuehl, Sandro and B Douglas Bernheim, "Interpreting the will of the people: a positive analysis of ordinal preference aggregation," Technical Report, National Bureau of Economic Research 2021.

Apple Inc., "Definitive Proxy Statement (Schedule 14A)," 2024.

_ , "Form 8-K: Current Report (Item 5.07 Submission of Matters to a Vote of Security Holders; 2025 Annual Meeting of Shareholders)," U.S. Securities and Exchange Commission (SEC), EDGAR February 2025. Filed February 25, 2025. SEC Accession No. 0001140361-25-005876.

Arrow, Kenneth J, *Social choice and individual values*, Vol. 12, Yale university press, 1951.

Bar-Gill, Oren, "Willingness to Pay: A Welfarist Reassessment," *Yale J. on Reg.*, 2021, *38*, 503.

Barber, Brad M, Adair Morse, and Ayako Yasuda, "Impact investing," *Journal of Financial Economics*, 2021, *139* (1), 162–185.

Bassi, Anna, "Voting Systems and Strategic Manipulation: An Experimental Study," *Journal of Theoretical Politics*, 2015, *27* (1), 58–85.

Bauer, Rob, Tobias Ruof, and Paul Smeets, "Get real! Individuals prefer more sustainable investments," *The Review of Financial Studies*, 2021, *34* (8), 3976–4043.

Berkshire Hathaway Inc., "Definitive Proxy Statement (Schedule 14A) – 2025 Annual Meeting of Shareholders," U.S. Securities and Exchange Commission (SEC) EDGAR 2025. Accession No. 0001193125-25-054877. Accessed 2025-12-18.

_ , "Form 8-K (Item 5.07) – Voting Results of the 2025 Annual Meeting," U.S. Securities and Exchange Commission (SEC) EDGAR 2025. Accession No. 0001193125-25-113900. Accessed 2025-12-18.

Black, Duncan, "On the rationale of group decision-making," *Journal of political economy*, 1948, *56* (1), 23–34.

Bolton, Patrick, Tao Li, Enrichetta Ravina, and Howard Rosenthal, "Investor ideology," *Journal of Financial Economics*, 2020, *137* (2), 320–352.

Bouton, Laurent, Micael Castanheira, and Aniol Llorente-Saguer, "Divided Majority and Information Aggregation: Theory and Experiment," *Journal of Public Economics*, 2016, *134*, 114–128.

Broadcom Inc., "Form 8-K: Current Report (Item 5.07 Submission of Matters to a Vote of Security Holders; 2025 Annual Meeting of Stockholders)," U.S. Securities and Exchange Commission (SEC), EDGAR April 2025. Filed April 22, 2025. SEC Accession No. 0001730168-25-000037.

Brodback, Daniel, Nadja Guenster, Sébastien Pouget, and Ruichen Wang, "Investor valuation for socially responsible assets: A willingness to pay experiment," *Management Science*, 2025.

Bubb, Ryan and Emiliano M Catan, "The party structure of mutual funds," *The Review of Financial Studies*, 2022, *35* (6), 2839–2878.

Dagens Næringsliv, "Fisk trues av elektromagnetisk stråling fra havvindanlegg," 2022. Accessed: 2024-11-05.

der Straeten, Karine Van, Jean-François Laslier, Nicolas Sauger, and André Blais, "Strategic, sincere, and heuristic voting under four election rules: an experimental study," *Social Choice and Welfare*, 2010, *35* (3), 435–472.

Döttling, Robin J., Doron Y. Levit, Nadya Malenko, and Magdalena A. Rola-Janicka, "Voting on Public Goods: Citizens vs Shareholders," NBER Working Paper 32605, National Bureau of Economic Research 2024. Revised March 2025.

Duffy, John, Daniel Friedman, Jean Paul Rabanal, and Olga Rud, "Trade, Voting, and ESG policies: Theory and Evidence," *Available at SSRN*, 2022.

Engelmann, Dirk, Hans Peter Grüner, Timo Hoffmann, and Alex Possajenikov, "Minority Protection in Voting Mechanisms–Experimental Evidence," Technical Report, Discussion Paper 2023.

Equinor, "New measures to maintain high gas exports to Europe," 2022. Accessed: 2024-11-05.

_ , "Havvind i Norge," 2024. Accessed: 2024-11-05.

Ericson, Keith Marzilli, "What Do Shareholders Want? Consumer Welfare and the Objective of the Firm," Technical Report, National Bureau of Economic Research 2024.

Europower, "Elektrifisering vil øke strømregningen med 1 milliard i året i Midt-Norge," 2023. Accessed: 2024-11-05.

Exxon Mobil Corporation, "Definitive Proxy Statement (Schedule 14A)," 2024.

Eyal, Peer, Rothschild David, Gordon Andrew, Evernden Zak, and Damer Ekaterina, "Data quality of platforms and panels for online behavioral research," *Behavior research methods*, 2021, pp. 1–20.

Forsythe, Robert, Thomas A. Rietz, Roger B. Myerson, and Robert J. Weber, "An Experimental Study of Voting Rules and Polls in Three-Candidate Elections," *International Journal of Game Theory*, 1996, *25* (3), 355–383.

Friedman, Milton, "The social responsibility of business is to increase its profits," in "Corporate ethics and corporate governance," Springer, 1970, pp. 173–178.

Gans, Joshua S and Michael Smart, "Majority voting with single-crossing preferences," *Journal of public Economics*, 1996, *59* (2), 219–237.

Gantchev, Nickolay and Mariassunta Giannetti, "The Costs and Benefits of Shareholder Democracy: Gadflies and Low-Cost Activism," *The Review of Financial Studies*, 2021, *34* (12), 5629–5675.

Gibbard, Allan, "Manipulation of voting schemes: a general result," *Econometrica: journal of the Econometric Society*, 1973, pp. 587–601.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, Zhenhao Tan, Stephen Utkus, and Xiao Xu, "Four facts about ESG beliefs and investor portfolios," *Journal of Financial Economics*, 2025, *164*, 103984.

Glass Lewis, "Glass Lewis Leads Change in Proxy Voting Practices," `https://www.glasslewis.com/news-release/glass-lewis-leads-change-in-proxy-voting-practices` October 2025. Announces that beginning in 2027 clients will be able to access multiple voting perspectives. Accessed 2026-03-29.

Gollier, Christian and Sébastien Pouget, "Investment strategies and corporate behaviour with socially responsible investors: A theory of active ownership," *Economica*, 2022, *89* (356), 997–1023.

Gormley, Todd A., Manish Jha, and Meng Wang, "The Politicization of Social Responsibility," NBER Working Paper 32869, National Bureau of Economic Research 2024. Revised January 2026.

Granić, Đura-Georg, "The Problem of the Divided Majority: Preference Aggregation under Uncertainty," *Journal of Economic Behavior & Organization*, 2017, *133*, 21–38.

Hart, Oliver and Luigi Zingales, "Companies should maximize shareholder welfare not market value," *ECGI-Finance Working Paper*, 2017, (521).

_ , David Thesmar, and Luigi Zingales, "Private sanctions," *Economic Policy*, 2024, *39* (117), 203–268.

Hartzmark, Samuel M and Abigail B Sussman, "Do investors value sustainability? A natural experiment examining ranking and fund flows," *The Journal of Finance*, 2019, *74* (6), 2789–2837.

He, Yazhou Ellen, Bige Kahraman, and Michelle Lowry, "ES Risks and Shareholder Voice," *The Review of Financial Studies*, 2023, *36* (12), 4824–4863.

Heeb, Florian, Julian F Kölbel, Falko Paetzold, and Stefan Zeisberger, "Do investors care about impact?," *The Review of Financial Studies*, 2023, *36* (5), 1737–1787.

Herrmann, Nathan D, John M McInnis, Brian Monsen, and Laura T Starks, "Decentralizing Proxy Voting Power," *Available at SSRN 5007107*, 2024.

Hu, Edwin, Nadya Malenko, and Jonathon Zytnick, "Custom proxy voting advice," Technical Report, National Bureau of Economic Research 2024.

Igersheim, Herrade, François Durand, Aaron Hamlin, and Jean-François Laslier, "Comparing Voting Methods: 2016 US Presidential Election," *European Journal of Political Economy*, 2022, *71*, 102057.

Iliewa, Zwetelina, Elisabeth Kempf, and Oliver G Spalt, "Corporate Actions as Moral Issues," *Available at SSRN*, 2024.

Jensen, Michael C, "Value maximization, stakeholder theory, and the corporate objective function," *Business ethics quarterly*, 2001, pp. 235–256.

Kakhbod, Ali, Uliana Loginova, Andrey Malenko, and Nadya Malenko, "Advising the management: A theory of shareholder engagement," *The Review of Financial Studies*, 2023, *36* (4), 1319–1363.

Krueger, Philipp, Zacharias Sautner, and Laura T Starks, "The importance of climate risks for institutional investors," *The Review of Financial Studies*, 2020, *33* (3), 1067–1111.

Lalley, Steven P. and E. Glen Weyl, "Quadratic Voting: How Mechanism Design Can Radicalize Democracy," *AEA Papers and Proceedings*, 2018, *108*, 33–37.

Levit, Doron, Nadya Malenko, and Ernst Maug, "The voting premium," Technical Report, National Bureau of Economic Research 2023.

_ , _ , and _ , "Trading and shareholder democracy," *The Journal of Finance*, 2024, *79* (1), 257–304.

Li, Sophia Zhengzi, Ernst Maug, and Miriam Schwartz-Ziv, "When shareholders disagree: Trading after shareholder meetings," *The Review of Financial Studies*, 2022, *35* (4), 1813–1867.

Malenko, Andrey and Nadya Malenko, "Proxy advisory firms: The economics of selling information to voters," *The Journal of Finance*, 2019, *74* (5), 2441–2490.

_ and _ , "Voting choice," Technical Report, National Bureau of Economic Research 2023.

Maskin, Eric, "Borda's rule and Arrow's independence condition," *Journal of Political Economy*, 2025, *133* (2), 385–420.

Maug, Ernst and Kristian Rydqvist, "Do shareholders vote strategically? Voting behavior, proposal screening, and majority rules," *Review of Finance*, 2009, *13* (1), 47–79.

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* (6), 2905–2932.

McFadden, Daniel, "The Foundations of Welfare Economics and Cost-Benefit Analysis," NBER Working Paper 23535, National Bureau of Economic Research 2017.

Meirowitz, Adam, Shaoting Pi, and Matthew C Ringgenberg, "Voting for socially responsible corporate policies," *Available at SSRN 4057792*, 2024.

Meta Platforms, Inc., "Definitive Proxy Statement (Schedule 14A) – 2025 Annual Meeting of Stockholders," U.S. Securities and Exchange Commission (SEC) EDGAR 2025. Accession No. 0001326801-25-000040. Accessed 2025-12-18.

_ , "Form 8-K (Item 5.07) – Voting Results of the 2025 Annual Meeting," U.S. Securities and Exchange Commission (SEC) EDGAR 2025. Accession No. 0001326801-25-000090. Accessed 2025-12-18.

Microsoft Corporation, "Definitive Proxy Statement (Schedule 14A) – 2025 Annual Shareholders Meeting," U.S. Securities and Exchange Commission (SEC) EDGAR 2025. Accession No. 0001193125-25-245150. Accessed 2025-12-18.

_ , "Form 8-K (Item 5.07) – Voting Results of the 2025 Annual Shareholders Meeting," U.S. Securities and Exchange Commission (SEC) EDGAR 2025. Accession No. 0001193125-25-311196. Accessed 2025-12-18.

Montagnes, B Pablo, Zachary Peskowitz, and Suhas A Sridharan, "How Well Do Voting Choice Policies Represent Public and Investor Preferences?," 2024.

Morgan, John and Justin Tumlinson, "Corporate provision of public goods," *Management Science*, 2019, *65* (10), 4489–4504.

NBIM, "About the Fund," 2024.

_ , "Our voting records," 2024.

Nehring, Klaus and Clemens Puppe, "The structure of strategy-proof social choice—Part I: General characterization and possibility results on median spaces," *Journal of Economic Theory*, 2007, *135* (1), 269–305.

Norges vassdrags- og energidirektorat, "Kunnskapsgrunnlag om virkninger av havvind: Fiskeri," 2023. Accessed: 2024-11-05.

NRK, "Olje for de fattiges skyld," 2013. Accessed: 2024-11-05.

_ , "Rapport: Elektrifisering av sokkelen har massiv klimaeffekt," 2023. Accessed: 2024-11-05.

NVIDIA Corporation, "Form 8-K: Current Report (Item 5.07 Submission of Matters to a Vote of Security Holders; 2025 Annual Meeting of Stockholders)," U.S. Securities and Exchange Commission (SEC), EDGAR July 2025. Filed July 1, 2025. SEC Accession No. 0001045810-25-000179.

Næss, Ole-Andreas Elvik, "Owner Support for Green Shareholder Voting and Divestment," 2025. Working Paper.

Posner, Eric A and E Glen Weyl, "Voting squared: Quadratic voting in democratic politics," *Vand. L. Rev.*, 2015, *68*, 441.

Reuters, "Advisory firm Glass Lewis backs two dissident nominees in Exxon battle," May 2021. Accessed 2026-03-29.

Riedl, Arno and Paul Smeets, "Why do investors hold socially responsible mutual funds?," *The Journal of Finance*, 2017, *72* (6), 2505–2550.

Satterthwaite, Mark Allen, "Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions," *Journal of economic theory*, 1975, *10* (2), 187–217.

SSB, "Hva er vanlig formue?," *Statistics Norway*, 2024. Accessed: 2024-05-31.

Starks, Laura T, "Presidential Address: Sustainable Finance and ESG Issues—Value versus Values," *The Journal of Finance*, 2023.

Statistics Norway, "Norsk olje- og gassproduksjon," 2013. Accessed: 2023-11-05.

Sunstein, Cass R., "Willingness to Pay vs. Welfare," *Harvard Law & Policy Review*, 2007, *1*, 303–330.

Tesla, Inc., "Definitive Proxy Statement (Schedule 14A)," 2024.

_ , "Form 8-K: Current Report (Item 5.07 Submission of Matters to a Vote of Security Holders; 2025 Annual Meeting of Shareholders)," U.S. Securities and Exchange Commission (SEC), EDGAR November 2025. Filed November 7, 2025. SEC Accession No. 0001104659-25-108507.

Weber, Matthias, "Choosing the rules: Preferences over voting systems for assemblies of representatives," *Journal of Economic Behavior & Organization*, 2020, *174*, 420–434.

Zingales, Luigi, Oliver Hart, and Helene E Landemore, "How To Implement Shareholder Democracy," *George J. Stigler Center for the Study of the Economy & the State Working Paper*, 2025, (350).

Zytnick, Jonathon, "Do Mutual Funds Represent Individual Investors?," *NYU Law and Economics Research Paper*, 2022, (21-04).

Çelebi, Can, Christine Exley, Sören Harrs, Hannu Kivimaki, Marta Serra-Garcia, and Jeffrey Yusof, "Mission Possible: The Collection of High-Quality Data," Technical Report, Working Paper December 2025. Working paper.

Summary of the Online Appendix

Section A presents additional tables, while Section B describes the algorithm for generating scenarios in Experiment 1 and 2. Section C provides supplementary results, while Section D provides robustness tests. Section E outlines the experimental instructions, and Section F provides the pre-analysis plans.

A Additional tables

Table A1: Point estimates of winners by voting rule

Experiment	Scenario	Plurality	Borda	Condorcet	IRV	WTP-opt
Exp 1 (all)	Alphabet	A	A	A	A	
Exp 1 (all)	Microsoft	A	A	A	A	
Exp 1 (all)	Amazon	A	A	A	A	
Exp 1 (all)	Meta	A	A	A	A	
Exp 1 (all)	Berkshire	B	A	A	A	
Exp 1 (wtp)	Alphabet	A	A	A	A	A
Exp 1 (wtp)	Microsoft	A	A	A	A	A
Exp 1 (wtp)	Amazon	A	A	A	A	A
Exp 1 (wtp)	Meta	A	A	A	A	A
Exp 1 (wtp)	Berkshire	B	B	B	B	A
Exp 2 (all)	Alphabet	A	A	A	A	
Exp 2 (all)	Microsoft	A	A	A	A	
Exp 2 (all)	Amazon	A	A	A	A	
Exp 2 (all)	Meta	A	A	A	A	
Exp 2 (all)	Berkshire	B	B	B	B	
Exp 3 (all)	H1	A	A	A	A	
Exp 3 (all)	H2	B	B	B	B	
Exp 3 (all)	H3	A	A	A	A	
Exp 3 (all)	H4	A	A	A	A	
Exp 3 (all)	E1	A	A	A	A	
Exp 3 (all)	E2	B	B	B	B	
Exp 3 (all)	E3	B	A	A	A	
Exp 3 (all)	E4	A	A	A	A	
Exp 3 (wtp)	E1	A	A	A	A	A
Exp 3 (wtp)	E2	B	B	B	B	A
Exp 3 (wtp)	E3	B	A	B	B	A
Exp 3 (wtp)	E4	A	A	A	A	A
Exp 4 (all)	B1	C	C	C	C	
Exp 4 (all)	B2	C	C	C	C	
Exp 4 (all)	B3	C	C	C	C	
Exp 4 (all)	B4	C	C	C	C	
Exp 4 (all)	B5	C	C	C	C	
Exp 4 (all)	B6	C	C	C	C	
Exp 5 (all)	N1	A	A	A	A	
Exp 5 (all)	N2	B	B	B	B	
Exp 5 (all)	N3	A	A	A	A	

Note: This table reports the winning alternative in each scenario under four voting rules, together with the willingness-to-pay maximizing alternative (WTP-opt). In Experiments 1 and 3, WTP is elicited only for a subset of respondents; accordingly, results are reported both for the full sample (*all*) and for the subset assigned to WTP elicitation in the given scenario (*wtp*).

Table A2: Outcome frequencies across bootstrap replications by voting rule

	Plurality			Borda			Condorcet				IRV (Runoff)		
Scenario	A	B	C	A	B	C	A	B	C	Cycle	A	B	C
Eksperiment 1													
Alphabet	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	100.0	0.0	0.0
Microsoft	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	99.6	0.0	0.0
Amazon	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	100.0	0.0	0.0
Meta	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	100.0	0.0	0.0
Berkshire	0.0	100.0	0.0	96.7	3.0	0.0	77.7	20.8	0.0	0.0	77.6	20.9	0.0
Eksperiment 2													
Alphabet	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	100.0	0.0	0.0
Microsoft	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	100.0	0.0	0.0
Amazon	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	100.0	0.0	0.0
Meta	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	100.0	0.0	0.0
Berkshire	4.2	95.5	0.0	1.9	98.0	0.0	6.4	93.4	0.0	0.0	6.4	93.4	0.0
Eksperiment 3													
E1	77.2	21.8	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	100.0	0.0	0.0
E2	1.2	98.6	0.0	8.2	91.5	0.0	9.3	90.2	0.0	0.0	9.3	90.2	0.0
E3	2.7	97.1	0.0	80.3	19.3	0.0	63.6	35.2	0.0	0.0	63.6	35.2	0.0
E4	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	100.0	0.0	0.0
H1	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	100.0	0.0	0.0
H2	0.0	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	99.6	0.0
H3	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	100.0	0.0	0.0
H4	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	100.0	0.0	0.0
Eksperiment 4													
B1	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	100.0
B2	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	100.0
B3	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	100.0
B4	0.0	0.0	100.0	0.2	0.0	99.8	0.0	0.0	100.0	0.0	0.0	0.0	100.0
B5	0.0	0.0	100.0	11.0	0.0	88.5	0.2	0.0	99.8	0.0	0.2	0.0	99.8
B6	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	100.0
Eksperiment 5													
N1	99.0	0.9	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	99.9	0.0	0.0
N2	0.0	94.1	5.5	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	100.0	0.0
N3	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	0.0	98.8	0.0	0.0

Note: This table reports the distribution of winners under plurality, Borda, Condorcet, and IRV (runoff) for each scenario across Experiments 1 to 5, based on 5,000 bootstrap replications. A rule is recorded as selecting a winner only when it yields a unique winner; replications with ties are omitted. "N" indicates that no Condorcet winner exists due to cyclical majority preferences.

Table A3: Pairwise welfare dominance across voting rules.

	Plurality	Borda	Condorcet	IRV
Panel 1: E1				
Plurality	–	0.00	0.00	0.00
Borda	1.00	–	.	1.00
Condorcet	1.00	.	–	1.00
IRV	1.00	0.00	0.00	–
Panel 2: E2				
Plurality	–	0.31	0.14	0.14
Borda	0.69	–	0.24	0.24
Condorcet	0.86	0.76	–	.
IRV	0.86	0.76	.	–
Panel 3: E3				
Plurality	–	0.28	0.24	0.24
Borda	0.72	–	0.73	0.73
Condorcet	0.76	0.27	–	1.00
IRV	0.76	0.27	0.00	–
Panel 4: Berkshire				
Plurality	–	0.00	0.00	0.00
Borda	1.00	–	0.99	1.00
Condorcet	1.00	0.01	–	1.00
IRV	1.00	0.00	0.00	–

Note: Off-diagonal entries report the share of bootstrap replications in which the row rule yields lower welfare loss than the column rule, conditional on the two rules producing different welfare outcomes. A value above 0.50 indicates that the row rule tends to outperform the column rule when they differ. Diagonal entries are omitted.

Table A4: Ranking shares and willingness-to-pay by scenario (WTP subsample)

Panel A: Experiment 3 (Scenarios E1–E4)

	E1			E2			E3			E4		
Ranking	Share	WTP_f	WTP_s	Share	WTP_f	WTP_s	Share	WTP_f	WTP_s	Share	WTP_f	WTP_s
ABC	18.80	4.87	4.10	25.79	5.38	4.81	19.27	4.31	3.28	24.22	5.00	4.19
ACB	18.48	4.36	4.09	11.74	4.97	3.70	16.94	3.32	3.18	20.93	5.17	3.91
BAC	31.44	4.94	4.42	25.79	4.61	3.82	27.24	5.18	4.43	16.96	5.10	4.15
BCA	4.27	3.67	2.37	13.39	2.69	2.60	15.12	4.62	3.10	12.80	4.72	4.23
CAB	23.85	3.64	3.02	12.23	3.41	2.57	13.29	2.86	2.55	12.80	2.91	2.42
CBA	3.16	4.15	3.25	11.07	2.36	2.13	8.14	2.67	2.18	12.28	2.58	1.92

Panel B: Experiment 1

	Alphabet			Microsoft			Amazon			Meta			Berkshire		
Ranking	Share	WTP_f	WTP_s	Share	WTP_f	WTP_s	Share	WTP_f	WTP_s	Share	WTP_f	WTP_s	Share	WTP_f	WTP_s
ABC	43.22	5.14	5.31	36.36	7.25	7.01	46.39	6.45	5.97	32.69	6.89	6.35	19.61	7.93	6.03
ACB	12.45	3.50	4.41	12.46	4.08	3.43	17.87	3.37	2.52	10.68	4.67	5.94	9.97	2.94	2.32
BAC	14.65	6.70	5.85	21.89	5.05	4.92	14.43	4.52	3.50	26.86	6.54	5.31	36.33	5.96	6.37
BCA	1.47	8.75	7.75	2.36	2.00	4.57	2.06	2.83	3.67	6.80	1.43	2.19	7.07	3.68	2.27
CAB	23.44	5.05	4.11	19.87	3.15	2.90	17.18	4.50	3.82	14.56	4.27	3.49	18.33	4.19	3.14
CBA	4.76	0.38	0.38	7.07	5.19	3.57	2.06	7.00	6.83	8.41	4.69	5.08	8.68	3.93	2.22

Notes: Shares are percentages within the WTP subsample. WTP_f and WTP_s denote mean willingness-to-pay for moving from the third-ranked alternative to the first- and second-ranked alternatives, respectively. In the firm experiment, WTP is elicited for one firm per respondent.

Table A5: Distribution of self-reported investment values across experiments.

Investment Value	Exp. 1 (%)	Exp. 3 (%)	Exp. 6 (%)	Exp. 7 (%)
\$0	5.54	25.64	5.51	4.46
\$0 – \$10,000	27.21	30.56	31.43	29.10
\$10,000 – \$100,000	29.37	19.44	29.39	29.05
\$100,000 – \$1,000,000	27.08	15.30	25.20	26.66
More than \$1,000,000	8.51	4.59	5.82	7.00
Prefer not to say	2.30	4.47	2.65	3.73
Average (USD)	257,422	134,344	218,485	260,833

Note: Average investment values are estimated using category midpoints (\$0, \$5,000, \$55,000, \$550,000) based on survey responses. The top category ("More than \$1,000,000") is assigned a value of \$1,000,000. "Prefer not to say" responses are excluded from the average calculations. Experiment 3 is conducted among representative Americans, while the others are conducted among self-reported investors. Entries are percentages.

Table A6: Tests of Question Order Effects by Scenario

Scenario	Exp. 1 rankings	WTPfirst	WTPsecond	Exp. 6 rankings
Alphabet	0.442	0.536	0.235	0.551
Amazon	0.551	0.877	0.493	0.380
Berkshire	0.327	0.831	0.297	0.425
Meta	0.462	0.003	0.099	0.413
Microsoft	0.728	0.118	0.557	0.841

Notes: Each entry reports the p-value from a Pearson chi-squared test of whether the distribution of responses differs by randomized question order (demographics before vs. after the scenario evaluations). In Experiment 1, ranking tests compare the distribution over the six possible rankings of alternatives A, B, and C, and WTP tests compare the distribution of responses in the multiple price list used to elicit willingness to pay to obtain the first-ranked (WTPfirst) or second-ranked (WTPsecond) alternative rather than the third-ranked alternative. Each respondent provided WTP responses for one randomly selected scenario only. In Experiment 6, ranking tests compare the distribution over the 24 possible rankings of alternatives A, B, C, and D.

Table A7: Logit and nonparametric trend tests (Experiments 1–2)

Panel	Scenario	Pair	Ordering	OR	logit p-value	nptrend p-value	Left %	Center %	Right %	N
Experiment 1	Alphabet	A over C		0.448	0.0000	0.00000	84.5	70.9	52.2	1456
		B over C		0.331	0.0000	0.00000	81.7	59.7	32.9	1456
		B over A	B-A-C	0.685	0.0000	0.00000	31.7	24.1	17.9	1456
	Microsoft	A over C		0.423	0.0000	0.00000	83.8	68.6	48.0	1456
		B over C		0.355	0.0000	0.00000	79.1	57.4	32.4	1456
		B over A	B-A-C	0.676	0.0000	0.00000	38.2	29.5	22.1	1456
	Amazon	A over C		0.551	0.0000	0.00000	88.6	81.0	70.2	1456
		B over C		0.370	0.0000	0.00000	80.9	61.1	36.8	1456
		B over A	B-A-C	0.672	0.0000	0.00000	24.8	18.2	13.0	1456
	Meta	A over C		0.381	0.0000	0.00000	84.3	67.2	43.8	1456
		B over C		0.347	0.0000	0.00000	83.4	63.5	37.6	1456
		B over A	B-A-C	0.905	0.1159	0.11507	44.7	42.3	39.9	1456
	Berkshire Hathaway	A over C		0.564	0.0000	0.00000	76.3	64.5	50.6	1456
		B over C		0.349	0.0000	0.00000	81.0	59.9	34.3	1456
		B over A	B-A-C	0.504	0.0000	0.00000	63.1	46.2	30.2	1456
Experiment 2	Alphabet	A over C		0.527	0.0000	0.00000	93.3	88.0	79.5	3200
		B over C		0.443	0.0000	0.00000	91.5	82.6	67.8	3200
		B over A	B-A-C	0.876	0.0068	0.00682	25.6	23.2	20.9	3200
	Microsoft	A over C		0.514	0.0000	0.00000	94.2	89.2	81.0	3200
		B over C		0.449	0.0000	0.00000	91.6	83.0	68.6	3200
		B over A	B-A-C	0.820	0.0001	0.00007	25.8	22.2	19.0	3200
	Amazon	A over C		0.607	0.0000	0.00000	94.1	90.7	85.6	3199
		B over C		0.422	0.0000	0.00000	93.5	85.8	71.9	3199
		B over A	B-A-C	0.795	0.0000	0.00000	30.7	26.0	21.9	3199
	Meta	A over C		0.515	0.0000	0.00000	93.7	88.4	79.7	3200
		B over C		0.427	0.0000	0.00000	92.8	84.6	70.1	3200
		B over A	B-A-C	0.804	0.0000	0.00000	37.2	32.3	27.7	3200
	Berkshire Hathaway	A over C		0.698	0.0000	0.00000	83.1	77.5	70.6	3200
		B over C		0.458	0.0000	0.00000	90.0	80.6	65.5	3200
		B over A	B-A-C	0.658	0.0000	0.00000	61.2	51.0	40.6	3200

Notes: Each row reports one pairwise comparison within a scenario. “Pair” indicates the binary outcome, coded as 1 when the first option is ranked above the second; for ambiguous cases, both admissible A–B orderings are shown in the “Ordering” column. “OR” and “logit p-value” come from a logit trend model in which political ideology is coded 0, 1, and 2 for Left, Center, and Right. “nptrend p-value” reports Stata’s nonparametric test for an ordered trend across the same three groups. “Left %”, “Center %”, and “Right %” are the shares in each ideological group for whom the first option in the pair is ranked above the second. “N” is the number of non-missing observations used in the test.

Table A8: Logit and nonparametric trend tests (Experiments 3–5)

Panel	Scenario	Pair	Ordering	OR	logit p-value	nptrend p-value	Left %	Center %	Right %	N
Experiment 3	H1	A over C		0.599	0.0000	0.00000	71.0	59.5	46.8	1226
		B over C		0.432	0.0000	0.00000	61.0	40.3	22.6	1226
		B over A	B-A-C	0.839	0.0326	0.03268	30.0	26.4	23.2	1226
		A over B	A-B-C	1.192	0.0326	0.03268	70.0	73.6	76.8	1226
	H2	A over C		0.564	0.0000	0.00000	65.3	51.5	37.5	1197
		B over C		0.485	0.0000	0.00000	71.4	54.7	37.0	1197
		B over A	B-A-C	0.743	0.0001	0.00006	61.8	54.6	47.2	1197
	H3	A over C		0.465	0.0000	0.00000	86.5	74.9	58.1	1148
		B over C		0.526	0.0000	0.00000	68.0	52.8	37.0	1148
		B over A	B-A-C	1.369	0.0002	0.00017	20.1	25.6	32.1	1148
		A over B	A-B-C	0.731	0.0002	0.00017	79.9	74.4	67.9	1148
	H4	A over C		0.652	0.0000	0.00000	87.5	82.0	74.8	1192
		B over C		0.481	0.0000	0.00000	62.3	44.3	27.6	1192
		B over A	B-A-C	0.845	0.0983	0.09648	17.5	15.2	13.1	1192
		A over B	A-B-C	1.184	0.0983	0.09648	82.5	84.8	86.9	1192
	E1	A over C		0.387	0.0000	0.00000	82.9	65.3	42.2	1235
		B over C		0.369	0.0000	0.00000	73.7	50.8	27.6	1235
		B over A	B-A-C	0.483	0.0000	0.00000	51.5	33.9	19.9	1235
	E2	A over C		0.550	0.0000	0.00000	72.0	58.6	43.8	1183
		B over C		0.344	0.0000	0.00000	80.3	58.4	32.6	1183
		B over A	B-A-C	0.615	0.0000	0.00000	60.7	48.7	36.9	1183
	E3	A over C		0.970	0.6828	0.68475	63.5	62.8	62.1	1207
		B over C		0.456	0.0000	0.00000	74.3	56.8	37.5	1207
		B over A	B-A-C	0.459	0.0000	0.00000	63.7	44.7	27.1	1207
	E4	A over C		0.817	0.0062	0.00608	65.7	61.0	56.1	1211
		B over C		0.376	0.0000	0.00000	71.3	48.3	26.0	1211
		B over A	B-A-C	0.548	0.0000	0.00000	54.8	39.9	26.7	1211
Experiment 4	B1	A over C		0.700	0.0000	0.00000	41.7	33.3	25.9	1469
		B over C		0.448	0.0000	0.00000	51.5	32.2	17.6	1469
		B over A	B-A-C	0.528	0.0000	0.00000	58.1	42.3	27.9	1469
	B2	A over C		0.549	0.0000	0.00000	46.8	32.6	21.0	1469
		B over C		0.441	0.0000	0.00000	52.8	33.0	17.9	1469
		B over A	B-A-C	0.746	0.0000	0.00000	59.2	52.0	44.7	1469
	B3	A over C		0.497	0.0000	0.00000	33.9	20.3	11.3	1469
		B over C		0.447	0.0000	0.00000	50.9	31.6	17.1	1469
		B over A	B-A-C	0.969	0.6758	0.67243	77.1	76.6	76.0	1469
	B4	A over C		0.578	0.0000	0.00000	54.7	41.1	28.7	1469
		B over C		0.448	0.0000	0.00000	52.5	33.1	18.2	1469
		B over A	B-A-C	0.690	0.0000	0.00000	41.2	32.6	25.0	1469
		A over B	A-B-C	1.449	0.0000	0.00000	58.8	67.4	75.0	1469
	B5	A over C		0.690	0.0000	0.00000	54.1	44.9	36.0	1469
		B over C		0.428	0.0000	0.00000	52.8	32.4	17.0	1469
		B over A	B-A-C	0.459	0.0000	0.00000	47.8	29.6	16.2	1469
	B6	A over C		0.582	0.0000	0.00000	48.7	35.6	24.3	1469
		B over C		0.439	0.0000	0.00000	54.5	34.5	18.8	1469
		B over A	B-A-C	0.695	0.0000	0.00000	58.2	49.2	40.2	1469
Experiment 5	N1	A over C		0.413	0.0000	0.00000	82.9	66.7	45.2	1337
		B over C		0.412	0.0000	0.00000	79.6	61.7	39.9	1337
		B over A	B-A-C	0.867	0.0390	0.03792	46.9	43.4	39.9	1337
		A over B	A-B-C	1.153	0.0390	0.03792	53.1	56.6	60.1	1337
	N2	A over C		0.491	0.0000	0.00000	64.2	46.9	30.2	1292
		B over C		0.429	0.0000	0.00000	76.5	58.3	37.4	1292
		B over A	B-A-C	0.724	0.0000	0.00001	71.4	64.4	56.7	1292
	N3	A over C		0.589	0.0000	0.00000	75.3	64.2	51.4	1314
		B over C		0.422	0.0000	0.00000	77.5	59.3	38.1	1314
		B over A	B-A-C	0.733	0.0000	0.00001	47.6	39.9	32.8	1314
		A over B	A-B-C	1.365	0.0000	0.00001	52.4	60.1	67.2	1314

Notes: This table is constructed in the same way as Table A7, but reports results for Experiments 3–5.

B Algorithm for Constructing Real-World Scenarios

This appendix describes a fully deterministic and auditable procedure that, for each firm $i \in U$, selects one non-climate shareholder proposal A_i, one climate shareholder proposal B_i, and three voting strategies (A, B, C) used to write the experimental vignettes. It is on purpose written to be very detailed (and hence also long) to ensure all necessary details are included to allow for replication.

B.1 Inputs, Parameters, and Output

Firm universe. Let U denote the set of the K largest U.S. firms by market capitalization as of a pre-specified date. The algorithm is applied to each firm $i \in U$.

Input data. Let i index firms and p proposals on firm i's 2025 ballot. For each (i, p) we observe: (i) proposer type (management vs. shareholder); (ii) proposal text (title, resolved clause, supporting statement); (iii) ballot agenda number; (iv) board recommendation $r_{ip} \in \{\text{FOR}, \text{AGAINST}, \text{NO_REC}\}$ (any other values are treated as missing unless mapped deterministically); and (v) vote outcome (for/against/abstain counts).

Support. Define support as

$$s_{ip} = \frac{\text{votes_for}_{ip}}{\text{votes_for}_{ip} + \text{votes_against}_{ip}}.$$

Abstentions are excluded from the denominator. If the denominator is zero or missing, set $s_{ip} = \text{NA}$ and exclude that proposal from candidate-set construction and selection.

Let τ denote the minimum shareholder support required for a proposal to be eligible for final selection. In the implementation, $\tau = 0$, so all proposals opposed by the board are eligible on the support dimension; support is therefore used only to rank eligible proposals, not to screen them out.

Board-line definition (explicit recommendations only). Define $r^*_{ip} \in \{\text{FOR}, \text{AGAINST}\}$ using only explicit board recommendations.[12]

$$r^*_{ip} = \begin{cases} \text{FOR}, & r_{ip} = \text{FOR}, \\ \text{AGAINST}, & r_{ip} = \text{AGAINST}, \\ \text{NA}, & r_{ip} = \text{NO_REC or missing}. \end{cases}$$

Proposals with $r^*_{ip} = \text{NA}$ are skipped (excluded) from selection in Step 5 and from the ballot item set used to define strategies in Step 6. This ensures that the one-item deviations in A and B are always deviations from an *explicit* board position, rather than an imputed position under `NO_REC`.

Output. For each retained firm i, the algorithm outputs a selected non-climate proposal A_i, a selected climate proposal B_i, and the corresponding ballot-level voting strategies (A, B, C).

B.2 Step 1: Normalization and Matching

Overview. All text matching is performed on a standardized text string. The same normalization is applied to both proposal text and dictionary terms before any match is checked.

Standardized proposal text. Define a deterministic function $\textsc{Normalize}(\cdot)$ that: (i) lowercases text; (ii) replaces every non-alphanumeric character, including hyphens, with a space; (iii) collapses consecutive whitespace to a single space; and (iv) trims leading and trailing spaces. For each proposal (i, p), define the missing-safe raw text

$$\texttt{raw_text}_{ip} := \texttt{title}_{ip} \parallel \texttt{resolved}_{ip} \parallel \texttt{supporting}_{ip},$$

[12] The algorithm treats the board's recommendation as a baseline voting line rather than as an identifying assumption about profit maximization. Strategies A and B are defined as one-item deviations from this baseline in a non-climate and a climate direction, respectively. Any interpretation of the board line as more profit-oriented belongs to the substantive discussion in the main text, not to the algorithmic construction itself.

where any missing field is deterministically recoded to the empty string before concatenation and $\|$ denotes concatenation with single-space separators. Then define

$$\texttt{full_text}_{ip} := \textsc{Normalize}(\texttt{raw_text}_{ip}),$$

and let $\texttt{tokens}_{ip}$ be the whitespace split of $\texttt{full_text}_{ip}$.

Matching convention. Each dictionary term is normalized once using the same $\textsc{Normalize}(\cdot)$ function and is then matched against $\texttt{full_text}_{ip}$. Thus `net-zero` and `net zero` are treated identically after normalization. A single-word term matches if it appears as an exact token in $\texttt{tokens}_{ip}$. A multi-word term matches if it appears as an exact substring in $\texttt{full_text}_{ip}$.

B.3 Step 2: Domain Classification

Overview. Each proposal is classified as `CLIMATE`, `SOCIAL`, or `OTHER` using fixed dictionaries applied to the full proposal text from Step 1, i.e. the title, resolved clause, and supporting statement. Each domain dictionary contains *strong* and *weak* terms. A strong term is sufficient on its own. A weak term is never sufficient on its own; absent a strong hit, at least two distinct matched terms from that domain are required. Climate takes precedence over social. Thus a proposal that matches both `climate` and `ai` is classified as `CLIMATE`, not `SOCIAL`. In the later steps, the relevant "non-climate" proposals are exactly those classified as `SOCIAL`; proposals classified as `OTHER` are never eligible.

Formal rule. Assign $D_{ip} \in \{\text{CLIMATE}, \text{SOCIAL}, \text{OTHER}\}$ as follows.

Climate rule. Let $\mathscr{K}^{\text{climate,strong}}$ and $\mathscr{K}^{\text{climate,weak}}$ be fixed, and let $\text{hits}_{ip}^{\text{clim}}$ denote the set of distinct matched terms in their union. Set $D_{ip} = \text{CLIMATE}$ iff either: (i) at least one strong climate term is matched, or (ii) $|\text{hits}_{ip}^{\text{clim}}| \geq 2$. Hence weak climate terms never classify a proposal on their own.

Social rule. If the climate rule is not satisfied, let $\mathscr{K}^{\text{social,strong}}$ and $\mathscr{K}^{\text{social,weak}}$ be fixed, and let $\text{hits}_{ip}^{\text{soc}}$ denote the set of distinct matched terms in their union. Set $D_{ip} = \text{SOCIAL}$ iff either: (i) at least one strong social term is matched, or (ii) $|\text{hits}_{ip}^{\text{soc}}| \geq 2$. Otherwise set $D_{ip} = \text{OTHER}$.

Full dictionaries (shown in normalized form). All dictionary terms below are shown after applying $\text{NORMALIZE}(\cdot)$, so punctuation and hyphen variants are handled automatically at the matching stage.

Climate dictionaries. **Strong climate terms** $\mathscr{K}^{\text{climate,strong}}$: `climate`, `global warming`, `greenhouse gas`, `ghg`, `emission`, `emissions`, `carbon`, `co2`, `methane`, `net zero`, `decarbonization`, `paris agreement`, `scope 1`, `scope 2`, `scope 3`, `science based target`.

Weak climate terms $\mathscr{K}^{\text{climate,weak}}$: `paris aligned`, `energy transition`, `transition risk`, `fossil fuel`, `oil sands`, `tar sands`, `coal`, `renewable energy`, `clean energy`, `solar`, `wind`.

Social dictionaries (used only if the climate rule fails). **Weak social triggers** $\mathscr{K}^{\text{social,weak}}$ (only help classify when at least one additional social term is also matched): `union`, `unions`, `equity`, `inclusion`, `ai`, `pac`, `super pac`.

For readability, the strong social dictionary $\mathscr{K}^{\text{social,strong}}$ is listed in eight thematic groups below; the groups themselves play no algorithmic role, and the algorithm uses only their union.

Workers & workplace ($\mathscr{K}^{\text{workers}}$): `worker`, `workers`, `employee`, `employees`, `collective bargaining`, `freedom of association`, `living wage`, `minimum wage`, `fair wage`, `overtime pay`, `working conditions`, `workplace safety`, `warehouse safety`, `gig worker`, `gig workers`, `contract worker`, `temporary worker`.

Human rights & supply chain ($\mathscr{K}^{\text{human_rights}}$): `human rights`, `labor rights`, `labour rights`, `forced labor`, `forced labour`, `child labor`, `child labour`, `modern slavery`, `slavery`, `human trafficking`, `supply chain due diligence`, `responsible sourcing`, `conflict minerals`, `indigenous rights`, `land rights`, `free prior and informed consent`, `community consent`.

Diversity, equity & inclusion ($\mathscr{K}^{\text{diversity}}$): `diversity`, `dei`, `diversity equity and inclusion`, `gender pay gap`, `racial equity`, `racial justice`, `civil rights`, `equal opportunity`, `non discrimination`,

marginalized communities, board diversity, lgbt, lgbtq, sexual orientation, gender identity, minority representation.

Digital rights, privacy & AI ($\mathscr{K}^{\text{digital_ai}}$): privacy, data protection, data security, cybersecurity, cyber security, artificial intelligence, machine learning, algorithmic bias, ai ethics, responsible ai, facial recognition, surveillance, content moderation, hate speech, misinformation.

Public health, product and environmental harm (non-climate) ($\mathscr{K}^{\text{health_safety}}$): public health, health impact, health impacts, health risk, health risks, product safety, consumer safety, food safety, opioid, opioids, tobacco, smoking, alcohol harm, pesticide, pesticides, chemical safety, toxic, toxic exposure, biodiversity, plastic waste, ocean plastic.

Animal welfare ($\mathscr{K}^{\text{animal_welfare}}$): animal welfare, animal testing, animal experimentation, vivisection, factory farming, intensive farming, gestation crate, battery cage, cage free, broiler chicken, humane slaughter, fur farming.

Political spending & lobbying ($\mathscr{K}^{\text{political_spending}}$): political contributions, political spending, political expenditures, corporate political activity, campaign finance, lobbying, lobbying expenditures, trade association, chamber of commerce, think tank, political action committee.

Other social ($\mathscr{K}^{\text{other_social}}$): community impact, community relations, community benefits, local community, local communities, environmental justice, social justice, human capital management, customer privacy, data breach, affordable housing, access to medicine, access to healthcare, food security.

B.4 Step 3: Anti–ESG Filter and Domain Eligibility

Overview. The anti–ESG step is applied only after a proposal has been classified as CLIMATE or SOCIAL. It is a separate text-based exclusion rule, not an additional dictionary used to define those domains. It removes proposals framed as opposition to ESG-type policies. For example, end dei is flagged directly by the anti-action dictionary, whereas a proposal containing both woke and esg is flagged by the framing-plus-target rule.

Formal rule. For proposals with $D_{ip} \in \{\text{CLIMATE}, \text{SOCIAL}\}$, define $\text{anti}_{ip} = 1$ iff either: (i) any term in the fixed anti-action dictionary matches $\texttt{full_text}_{ip}$ under the matching

rule in Step 1; or (ii) at least one framing term and at least one target term both match $\texttt{full_text}_{ip}$ under the matching rule in Step 1. Otherwise set $\text{anti}_{ip} = 0$.

Domain eligibility. Define

$$E_{ip} = 1 \quad \Longleftrightarrow \quad D_{ip} \in \{\text{CLIMATE}, \text{SOCIAL}\} \text{ and } \text{anti}_{ip} = 0.$$

Otherwise set $E_{ip} = 0$. Thus proposals classified as `OTHER` are automatically ineligible at this stage. Final candidate sets in Step 4 additionally require that the proposal is shareholder-sponsored, has non-missing support, and has an explicit board recommendation.

Anti–ESG dictionaries (shown in normalized form). All terms below are shown after applying $\text{NORMALIZE}(\cdot)$.

Targets $\mathscr{K}^{\text{anti,target}}$: `esg`, `environmental social governance`, `dei`, `diversity equity and inclusion`, `net zero`, `sustainability`, `sustainable`.

Explicit anti-actions $\mathscr{K}^{\text{anti,action}}$: `anti esg`, `against esg`, `oppose esg`, `end esg`, `stop esg`, `eliminate esg`, `abolish esg`, `ban esg`, `prohibit esg`, `discontinue esg`, `remove esg`, `phase out esg`, `reverse esg`, `rollback esg`, `end esg program`, `end esg programs`, `reverse esg policies`, `prohibit the use of esg factors`, `ban the use of esg factors`, `against the use of esg factors`, `anti dei`, `against dei`, `oppose dei`, `end dei`, `stop dei`, `eliminate dei`, `abolish dei`, `ban dei`, `prohibit dei`, `discontinue dei`, `remove dei`, `phase out dei`, `reverse dei`, `rollback dei`, `oppose diversity equity and inclusion`, `reverse diversity policies`, `reverse racial equity`, `end net zero`, `stop net zero`, `ban net zero`, `prohibit net zero`, `reverse net zero`.

Anti–ESG framing terms $\mathscr{K}^{\text{anti,frame}}$: `woke`, `wokeism`, `politicized`, `ideological`, `corporate activism`, `political agenda`, `stakeholder capitalism`, `critical race theory`, `crt`, `reverse discrimination`, `viewpoint discrimination`, `political discrimination`, `risks of esg`, `costs of esg`, `risks of dei`, `costs of dei`.

B.5 Step 4: Firm-Level Candidate Sets and Ordering

For each firm i, define the climate and non-climate candidate sets

$$\mathscr{P}_i^{\text{climate}} = \left\{ p : \begin{array}{l} p \text{ is a shareholder proposal}, \\ D_{ip} = \text{CLIMATE}, \\ E_{ip} = 1, \\ s_{ip} \neq \text{NA}, \\ r^*_{ip} \neq \text{NA} \end{array} \right\},$$

and

$$\mathscr{P}_i^{\text{social}} = \left\{ p : \begin{array}{l} p \text{ is a shareholder proposal}, \\ D_{ip} = \text{SOCIAL}, \\ E_{ip} = 1, \\ s_{ip} \neq \text{NA}, \\ r^*_{ip} \neq \text{NA} \end{array} \right\}.$$

Order each set by decreasing support s_{ip}, tie-breaking by increasing ballot agenda number. If either candidate set is empty, drop the firm and record the drop reason.

B.6 Step 5: Select One Climate Proposal and One Non-Climate Proposal

For each firm i, let B_i be the first element of the ordered set $\mathscr{P}_i^{\text{climate}}$ such that $r^*_{ip} = \text{AGAINST}$ and $s_{ip} \geq \tau$. Analogously, let A_i be the first element of the ordered set $\mathscr{P}_i^{\text{social}}$ such that $r^*_{ip} = \text{AGAINST}$ and $s_{ip} \geq \tau$.

If either A_i or B_i does not exist, drop the firm and record the drop reason. Let

$$\mathscr{S} = \{ i \in U : A_i \text{ and } B_i \text{ both exist} \}$$

denote the set of retained firms.

B.7 Step 6: Define Strategies A, B, and C

For each retained firm $i \in \mathscr{S}$, let $\mathscr{P}_i$ denote the set of all ballot items with r^*_{ip} defined. Define the board-line strategy C by

$$\text{vote}^C_{ip} = r^*_{ip} \qquad \forall p \in \mathscr{P}_i.$$

Define strategy B by

$$\text{vote}^B_{ip} = \begin{cases} \text{FOR}, & p = B_i, \\ r^*_{ip}, & \text{otherwise}, \end{cases}$$

and strategy A by

$$\text{vote}^A_{ip} = \begin{cases} \text{FOR}, & p = A_i, \\ r^*_{ip}, & \text{otherwise}. \end{cases}$$

By construction, $r^*_{iA_i} = r^*_{iB_i} = \text{AGAINST}$, so A and B each implement exactly one salient one-item deviation from the board line.

B.8 Audit Log

The replication package records, for each firm i: (i) proposal ID, agenda number, and title; (ii) r_{ip} and r^*_{ip}; (iii) s_{ip}; (iv) matched dictionary terms for domain classification and for the anti–ESG filter; (v) the ordered climate and non-climate candidate lists; and (vi) either the selected pair (A_i, B_i) or the precise drop reason.

B.9 Step-by-Step Algorithm (Pseudocode)

Algorithm 1: Deterministic scenario construction

Input: Ballot-level dataset for year 2025 with fields above; fixed dictionaries

Output: Scenario firms and (A_i, B_i) with strategies (A, B, C)

Normalize all dictionary terms using NORMALIZE$(\cdot)$

foreach *proposal* (i, p) **do**

Construct $\texttt{full_text}_{ip}$ and $\texttt{tokens}_{ip}$;

Compute domain $D_{ip} \in \{\text{CLIMATE}, \text{SOCIAL}, \text{OTHER}\}$ (Step 2);

Compute anti flag $\text{anti}_{ip} \in \{0, 1\}$ (Step 3);

Set eligibility $E_{ip} = 1$ if $D_{ip} \in \{\text{CLIMATE}, \text{SOCIAL}\}$ and $\text{anti}_{ip} = 0$;

foreach *firm* i **do**

Build candidate lists:

$$\mathscr{P}_i^{\text{climate}} = \{p : \text{shareholder proposal}, D_{ip} = \text{CLIMATE}, E_{ip} = 1, s_{ip} \neq NA, r_{ip}^* \neq NA\}$$

$$\mathscr{P}_i^{\text{social}} = \{p : \text{shareholder proposal}, D_{ip} = \text{SOCIAL}, E_{ip} = 1, s_{ip} \neq NA, r_{ip}^* \neq NA\}$$

If either set is empty: drop firm and record reason;

Sort each list by decreasing s_{ip}, tie-break by increasing agenda number;

Select B_i as the first climate proposal with $r_{ip}^* = \text{AGAINST}$ and $s_{ip} \geq \tau$;

Select A_i as the first social proposal with $r_{ip}^* = \text{AGAINST}$ and $s_{ip} \geq \tau$;

If either selection fails: drop firm and record reason;

Else keep firm in set $\mathscr{S}$ and record (A_i, B_i);

Define strategies (A, B, C) on each ballot in $\mathscr{S}$ (Step 6);

B.10 Implementing the Algorithm

I set $\tau = 0$ and $K = 9$. Thus, U consists of the nine largest U.S. firms by market capitalization as of January 2026. Each of these firms had a market capitalization above approximately $1 trillion on that date. The resulting firm universe is: NVIDIA, Apple, Alphabet, Microsoft, Amazon, Meta, Tesla, Broadcom, and Berkshire Hathaway.

B.11 Dropped Firms

The following four firms are dropped by the algorithm in Step 4 because at least one required candidate set is empty. In Tesla, Apple, and NVIDIA, the climate candidate set is empty; in Broadcom, both the climate and non-climate candidate sets are empty because the ballot contains no shareholder proposals.

- **Tesla:** Tesla's 2025 annual-meeting ballot includes shareholder proposals on topics such as integrating *sustainability metrics* into executive compensation and requesting a *child labor audit*, plus governance/bylaw items, but no climate-themed shareholder proposal; hence $\mathscr{P}^{\text{climate}}_{\text{Tesla}} = \varnothing$ and the firm is dropped in Step 4 (Tesla, Inc., 2025).[13]
- **Broadcom:** Broadcom's 2025 annual-meeting ballot includes only management matters (director elections, auditor ratification, and say-on-pay) and therefore has no shareholder proposals; consequently $\mathscr{P}^{\text{social}}_{\text{Broadcom}} = \mathscr{P}^{\text{climate}}_{\text{Broadcom}} = \varnothing$ and the firm is dropped in Step 4 (Broadcom Inc., 2025).
- **Apple:** Apple's 2025 ballot includes shareholder proposals on ethical AI data acquisition, CSAM-identifying software and user privacy, DEI, and charitable giving; none are climate proposals under the algorithm's climate dictionary, so $\mathscr{P}^{\text{climate}}_{\text{Apple}} = \varnothing$ and the firm is dropped in Step 4 (Apple Inc., 2025).
- **NVIDIA:** NVIDIA's 2025 ballot includes shareholder proposals on special-meeting thresholds, a director resignation policy, and workforce-data reporting, but no climate proposal; thus $\mathscr{P}^{\text{climate}}_{\text{NVIDIA}} = \varnothing$ and the firm is dropped in Step 4 (NVIDIA Corporation, 2025).

B.12 Retained Firms and Selected Proposal Pairs

I successfully implemented Algorithm 1 on the 2025 annual-meeting ballots of five firms: Alphabet, Amazon, Microsoft, Meta, and Berkshire Hathaway. For each firm, I collected two

[13] The term `sustainability` was intentionally excluded from the climate dictionary because it can refer to non-climate sustainability concepts as well.

SEC filings: (i) the definitive proxy statement (Schedule 14A), which provides each ballot item's proposal text and the board's voting recommendation, and (ii) the post-meeting Form 8-K (Item 5.07), which reports vote tabulations for each matter submitted to shareholders.[14]

Following Algorithm 1, I concatenated each proposal's title, resolved clause, and supporting statement from the proxy statement, normalized the resulting text, and applied the fixed climate/social dictionaries and the anti–ESG filter. For each firm i, this produced two ordered candidate sets of shareholder proposals, $\mathscr{P}_i^{\text{climate}}$ and $\mathscr{P}_i^{\text{social}}$. Within each set, proposals were ranked by s_{ip}, with ties broken by ballot agenda number. I then selected B_i as the highest-ranked climate shareholder proposal with $r_{ip}^* = \text{AGAINST}$ and $s_{ip} \geq \tau$, and analogously selected A_i from the social candidate set. I then defined strategy C to follow the board recommendation on all ballot items with an explicit board recommendation, strategy A to differ from C only by voting `FOR` on A_i, and strategy B to differ from C only by voting `FOR` on B_i. Finally, for survey administration, I rewrote each selected proposal into a short vignette that preserves the operative request in the resolved clause while removing non-essential legal or technical wording.

Selected proposal pairs. Applying the procedure yielded the following firm-specific scenario pairs (A_i, B_i) (proposal descriptions and board recommendations in the proxy statements; vote outcomes in the Forms 8-K):

- **Alphabet (GOOGL):** A_i = Proposal 12 (human-rights impact assessment of AI-driven targeted advertising policies); B_i = Proposal 7 (enhanced disclosure on Alphabet's climate goals) (Alphabet Inc., 2025a,b).
- **Amazon.com (AMZN):** A_i = Item 10 (independent audit/report on warehouse working conditions); B_i = Item 7 (report on how data-center growth affects Amazon's ability to meet climate commitments) (Amazon.com, Inc., 2025a,b).
- **Microsoft (MSFT):** A_i = Proposal 8 (report on data operations in human-rights hotspots); B_i = Proposal 10 (report on climate and financial risks from providing

[14] For each firm, the cited proxy statement provides the proposal description (title, resolved clause, and supporting statement) and the board recommendation; the cited Form 8-K provides the voting results.

AI and machine-learning tools for oil and gas development/production) (Microsoft Corporation, 2025a,b).

- **Meta Platforms (META):** A_i = Proposal 8 (report on hate targeting marginalized communities); B_i = Proposal 12 (report on GHG emissions reduction actions) (Meta Platforms, Inc., 2025a,b).
- **Berkshire Hathaway (BRK.A/BRK.B):** A_i = Proposal 8 (independent board committee for oversight of AI); B_i = Proposal 7 (report on Berkshire's clean-energy supply financing ratio) (Berkshire Hathaway Inc., 2025a,b)

The content of these scenarios as described to survey respondents is displayed on Appendix E.

C Supplementing results

C.1 Detailed Analysis of Scenarios

C.1.1 Real-World Basis of Scenarios E1–E4

All four scenarios mirror actual trade-offs faced by Equinor, whose majority shareholder is the Norwegian state (holding 67% of its shares).[15] In E1, Norwegian climate activists support redirecting Equinor's investments toward renewable energy, whereas others contend that developing new oil and gas fields strengthens geopolitical stability by reducing Europe's reliance on Russian gas (Equinor, 2022). A similar debate arises in E2: while some dispute the claim that developing countries benefit from continued Norwegian oil exports (Statistics Norway, 2013), it remains central to arguments put forth by influential politicians across the political spectrum (NRK, 2013). Scenario E3 reflects a key policy debate in Norway: proponents of electrification emphasize its climate benefits, whereas critics argue that it raises household energy costs (Europower, 2023; NRK, 2023). In E4, Equinor is actively pursuing

[15] Although Equinor and the Norwegian Sovereign Wealth Fund are both owned by the Norwegian state, they operate under distinct governance structures. Equinor is a publicly traded company with government majority ownership.

offshore wind projects (Equinor, 2024), though critics express concerns about their potential effects on fishing communities and biodiversity (Dagens Næringsliv, 2022; Norges vassdrags- og energidirektorat, 2023).

C.2 Experiment 7, Part 3: The effect of Voting Choice

This third part of the experiment (which was chronologically first in Experiment 7) studies how investor preferences translate into *realized votes* in pass-through voting systems where participation is costly and investors may instead delegate their votes through a *Voting Choice* program. The motivation is the rapid expansion of pass-through voting among large asset managers. To study how such programs affect realized voting outcomes, I construct an experimental module based on Vanguard's Voting Choice program.

Under Voting Choice, investors select a voting profile that determines how their shares are voted across shareholder proposals. I focus on the four most commonly used profiles: the Vanguard default policy, a board/management-aligned policy, the Glass Lewis ESG profile, and an Egan-Jones profile with a stronger focus on financial performance.[16]

The module begins by eliciting respondents' preferences over these four profiles. Respondents rank the profiles according to how they would want their shares to be voted. They are then randomly assigned to one of the following treatments.

In the **costless item-by-item voting** treatment, respondents vote directly on three shareholder proposals, one by one. This provides a benchmark in which participation is effectively full and preferences can be expressed separately on each proposal. In the **Voting Choice** treatment, respondents are asked whether they would use Voting Choice. Those who opt in have their votes determined by the profile they ranked first in the initial ranking task, while those who opt out receive the institutional default vote. Finally, in the **costly voting** treatment, respondents are asked whether they are willing to wait 40 seconds without compensation before proceeding to vote on the same three proposals. Those who choose to wait then vote

[16] A fifth profile exists in practice but has very limited adoption, so it is omitted here in order to keep the environment tractable.

on each proposal individually, while those who decline receive the institutional default vote.[17] This treatment aims to capture elements from current environment in which voting requires effort and non-participation leads to the default vote.

The three shareholder proposals correspond to real ballot items from recent U.S. annual meetings (ExxonMobil 2024, Apple 2024, and Tesla 2024). These ballots items are chosen because the recommendations of the board, Glass Lewis, Vanguard, and Egan-Jones are publicly observable, meaning I can recover how each Voting Choice profile would vote on each proposal. Furthermore, they span compensation (Tesla), governance and shareholder rights (ExxonMobil), and AI transparency (Apple), while generating meaningful disagreement across profiles (Apple Inc., 2024; Exxon Mobil Corporation, 2024; Tesla, Inc., 2024). This makes it possible to compare realized item-by-item voting with the counterfactual outcome implied by profile-based delegation.

Profile choice and take-up. Preferences over Voting Choice profiles are heterogeneous. 29.1% rank Glass Lewis (ESG) first, 19.6% rank the Vanguard-advised default profile first, 6.9% rank the company/board-aligned profile first, and 44.3% rank the Egan-Jones (wealth-focused) profile first.

Take-up of Voting Choice is substantial and varies systematically with profile preferences. Overall, 67.8% of respondents in the Voting Choice treatment state that they would actively select a voting profile. Among those who rank the ESG profile first, 86.1% say that they would opt in. In contrast, only 44.3% of respondents who rank the Vanguard-advised default profile first would actively select a profile. Take-up is 67.7% among respondents preferring the company/board-aligned profile and 66.2% among those preferring the Egan-Jones profile. These patterns unsurprisingly suggest that profile-based delegation is most attractive to respondents whose preferred profile differs from the default. All else equal, this would suggest that introducing Voting choice through a turnout effect. However, introducing Voting Choice also has other effects as described below.

[17]There is also a fourth treatment, similar to the costly-voting treatment but without the 40-second delay. For simplicity, I do not analyze that treatment here.

Decomposition of ESG-aligned voting outcomes. The main outcome in this part of the experiment is the expected ESG-aligned vote share under each regime. Let m^A, m^B, and m^C denote ESG-aligned vote shares in the Voting Choice treatment, the costly-voting treatment, and the costless item-by-item treatment, respectively. A useful way to understand the total effect of introducing Voting Choice relative to the status quo is to separate it into $m^A - m^B = (m^C - m^B) + (m^A - m^C)$. The first term, $m^C - m^B$, can in turn be decomposed into a turnout component and a selection component. Let π^B denote turnout in the costly-voting treatment, and let μ^B denote ESG-aligned support among those respondents in treatment B who actually cast a vote. Then $m^B = \pi^B \mu^B$, since respondents who do not continue to vote effectively leave the outcome at the default. Substituting this back in yields

$$m^A - m^B = \underbrace{(1-\pi^B)\mu^B}_{\text{turnout/default-displacement}} + \underbrace{(m^C - \mu^B)}_{\text{Selection}} + \underbrace{(m^A - m^C)}_{\text{Voting Choice wedge}} .$$

This decomposition clarifies what changes when Voting Choice is introduced. One margin operates through participation: lowering the effective cost of voting moves investors away from the default and toward active choice. A second margin operates through selection: those who vote under costly direct voting are not necessarily representative of the full sample. A third margin operates through delegation itself: selecting a profile aggregates preferences differently than costless item-by-item voting.

Table A9 reports both the total effect, $m^A - m^B$, and its three components. The net effect of introducing Voting Choice is modest. However, the decomposition shows that this small net effect masks substantial offsetting forces: lower voting costs increase ESG-aligned voting through turnout, while selection into participation and delegation through policy profiles push outcomes in the opposite direction.

Table A9: Decomposition of the effect of Voting Choice

	Outcome (ESG aligned)			Decomposition		
	Voting Choice	Costly	Δ	Turnout	Selection	VC wedge
Exxon	25.1	23.2	1.9	43.9	-6.1	-35.9
Apple	25.1	30.5	-5.4	57.7	-4.5	-58.6
Tesla	25.1	21.9	3.2	41.4	7.2	-45.4
Average	25.1	25.2	-0.1	47.6	-1.2	-46.6

D Robustness

D.1 Identifying left-wing alternative

This section describes five methods (labeled M1 to M5) for selecting the most left-wing alternative in Experiment 3 to 5 (since Experiment 1 and 2 pre-registered the ranking B-A-C, this is not included here). The specification M1 compares only outcomes A and B by counting the number of left-wing supporters, whereas the alternative methods (M2 to M5) evaluate all three outcomes (A, B, and C) as follows:

- **M1 (Most L supporters, A vs. B only)** Compares only outcomes A and B by counting left-wing supporters and selects the outcome with the higher count.
- **M2 (Largest $(L-R)$)**

  For each outcome, computes the difference between the number of left-wing and right-wing supporters, i.e., $L-R$, and chooses the outcome with the largest difference.
- **M3 (Highest L/R ratio)**

  Calculates the ratio of left-wing to right-wing supporters (L/R) for each outcome. The outcome with the highest ratio is then selected.
- **M4 (Largest fraction $L/(L+M+R)$)**

Computes the fraction of left-wing supporters among all supporters (left, middle, and right) for each outcome and picks the outcome where this fraction is highest.

- **M5 (Largest $(L+M-R)$)**

  Calculates a net score for each outcome as the sum of left and middle supporters minus right supporters, i.e., $(L+M-R)$, and selects the outcome with the highest net score.

Table A10 shows the results.

Table A10: Left-wing alternative across experiments

(a) Experiment 3

	M1	M2	M3	M4	M5
H1	A	A	B	B	A
H2	B	B	B	B	B
H3	A	A	A	A	A
H4	A	A	B	B	A
E1	B	B	B	B	B
E2	B	B	B	B	B
E3	B	B	B	B	B
E4	B	B	B	B	B

(b) Experiment 4

	M1	M2	M3	M4	M5
B1	B	B	B	B	B
B2	B	B	B	B	B
B3	B	B	B	B	B
B4	A	B	B	B	A
B5	B	B	B	B	B
B6	B	B	B	B	B

(c) Experiment 5

	M1	M2	M3	M4	M5
N1	B	B	B	B	A
N2	B	B	B	B	B
N3	A	B	B	B	B

This table shows how five different methods (M1 to M5) identify the most left-wing alternative. Panel (a) reports scenarios H1–H4 and E1–E4 (Experiment 3), Panel (b) reports scenarios B1–B6 (Experiment 4), and Panel (c) reports scenarios N1–N3 (Experiment 5).

D.2 Value-weighted analysis of Experiment 1 and 3

Ownership weighting. I assess the robustness of the main results to ownership weighting. Respondents are weighted by their reported total investments in stocks and mutual funds, so that individuals with larger financial stakes receive greater influence in the analysis. To avoid giving excessive weight to extreme observations or misreporting in the upper tail, I use a compressed tiered weighting scheme rather than weights proportional to reported wealth. Specifically, respondents reporting zero investments, or selecting "Prefer not to say," are assigned weight 0; respondents reporting investments between \$0 and \$10,000 are assigned weight 1; those reporting \$10,000–\$100,000 are assigned weight 2; those reporting \$100,000–\$1,000,000 are assigned weight 3; and those reporting more than \$1,000,000 are

assigned weight 4. This approach increases the influence of respondents with larger holdings while limiting the sensitivity of the results to outliers. Weighted results are reported for Experiments 1 and 3, where the relevant investment-value question is available.

Table A11: Weighted distribution of rankings across scenarios in Experiments 1 and 3 (shares in %)

Panel A. Experiment 1 (weighted)

Ranking	Alphabet	Microsoft	Amazon	Meta	Berkshire
A≻ B ≻ C	36.3	31.9	42.6	30.9	20.1
A ≻ C ≻ B	9.8	10.4	18.2	7.6	11.8
B ≻ A ≻ C	18.3	21.0	16.0	25.1	30.1
B ≻ C ≻ A	3.7	2.9	0.8	5.4	6.7
C ≻ A ≻ B	23.0	23.2	17.7	17.0	21.2
C ≻ B ≻ A	8.9	10.7	4.7	14.0	10.2
Weighted N	1424.0	1424.0	1424.0	1424.0	1424.0

Panel B. Experiment 3 (weighted)

Ranking	H1	H2	H3	H4	E1	E2	E3	E4
A≻ B ≻ C	20.5	16.3	28.7	27.0	18.0	16.6	11.6	16.2
A ≻ C ≻ B	23.2	13.3	28.4	41.6	19.9	9.9	16.6	20.5
B ≻ A ≻ C	7.1	18.2	11.9	9.2	22.2	25.5	21.7	14.1
B ≻ C ≻ A	7.8	8.8	4.8	4.3	3.7	15.0	18.4	13.5
C ≻ A ≻ B	28.0	23.9	16.6	14.9	34.9	16.0	21.6	18.0
C ≻ B ≻ A	13.4	19.5	9.5	3.1	1.2	17.0	10.1	17.6
Weighted N	1180.0	1074.1	1152.5	1130.0	1199.9	1110.1	1123.8	1186.2

Notes: Each cell reports the ownership-weighted share with the given strict ranking in that scenario. Only Experiments 1 and 3 are included. The last row reports the total weighted mass in each scenario.

Ownership-weighted results. Tables A11–A12 show that the main qualitative findings are robust to ownership weighting. Weighting respondents by reported investment holdings does not alter the overall structure of preferences: mass remains concentrated on a limited set of

Table A12: Single-crossing constraints by experiment and assumed order: weighted results

Scenario	# Constraints	Violated	Binding	Not binding
Exp1				
Exp1 (BAC)	90	0	25	65
Exp3				
Exp3 (ABC)	12	0	6	6
Exp3 (BAC)	144	0	38	106
Total	246	0	69	177

Table A13: Weighted single-peakedness results

Scenario	Weighted N	Max SP (%)	Boot. 95% CI	OOS axis median (%)	MC p-value
Experiment 1					
Alphabet	2939.0	90.4	[88.5, 92.3]	90.4	< 0.00005
Microsoft	2939.0	89.4	[87.5, 91.2]	89.4	< 0.00005
Amazon	2939.0	95.4	[94.1, 96.7]	95.4	< 0.00005
Meta	2939.0	82.0	[79.7, 84.3]	82.0	< 0.00005
Berkshire Hathaway	2939.0	84.1	[81.9, 86.2]	84.1	< 0.00005
Experiment 3					
H1	1622.0	79.9	[76.7, 83.1]	79.9	< 0.00005
H2	1560.0	70.2	[68.6, 73.8]	68.8	0.08025
H3	1571.0	84.6	[81.6, 87.3]	84.6	< 0.00005
H4	1611.0	93.4	[91.4, 95.2]	93.4	< 0.00005
E1	1678.0	93.3	[91.4, 94.9]	93.3	< 0.00005
E2	1555.0	73.6	[71.0, 76.9]	73.6	0.00035
E3	1597.0	73.5	[70.0, 76.9]	73.5	0.00025
E4	1636.0	70.9	[68.0, 74.4]	70.9	0.02655

Notes: Results are ownership-weighted using the weight variable in the cleaned data. Weighted N reports the total weight mass, not the raw number of respondents. Max SP = $100 \times (1 - \min\{p_A, p_B, p_C\})$, where p_A, p_B, and p_C are the weighted shares that place A, B, and C last, respectively. Bootstrap confidence intervals are percentile intervals based on 2,000 resamples. OOS axis median reports the median across 200 random splits of the cross-fitted holdout fit: the middle alternative is chosen in one half-sample and evaluated on the other half, averaging both directions. When the training half-sample yields multiple tied middle alternatives, the holdout fit is averaged across those tied middles. MC p-values are based on 20,000 simulations under the null of uniformly random rankings, respecting the observed integer weight distribution.

rankings, the single-peakedness results remain strong, and the single-crossing restrictions continue to hold in all binding cases.

Relative to the baseline results, ownership weighting shifts some support toward the C option, especially in Experiment 3 and, to a lesser extent, in several Experiment 1 scenarios. In Experiment 1, the weighted ranking distributions in Table A11 continue to show that A-first rankings are most common for Alphabet, Microsoft, and Amazon, while $B \succ A \succ C$ remains relatively stronger for Meta and Berkshire Hathaway. At the same time, first-place support for C rises modestly in most cases. The same pattern is clear in Experiment 3.

The weighted single-crossing results in Table A12 remain fully consistent with the baseline analysis. There are no violated binding constraints in either experiment under any ordering considered.

Finally, the weighted single-peakedness results in Table A13 are very similar to the baseline findings. In Experiment 1, the maximum single-peaked fit ranges from 82.0% to 95.4%, with all Monte Carlo p-values below 0.00005. In Experiment 3, the fit ranges from 70.2% to 93.4%.

E Experimental instructions

E.0.1 Experiment 1

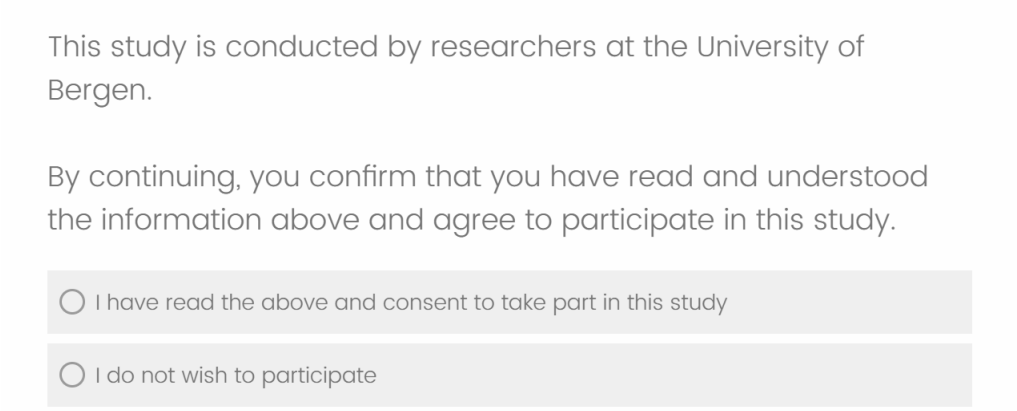


Figure A1: Introduction

Note: The same procedure was used in Experiments 9 and 10.

E.0.2 Experiment 1

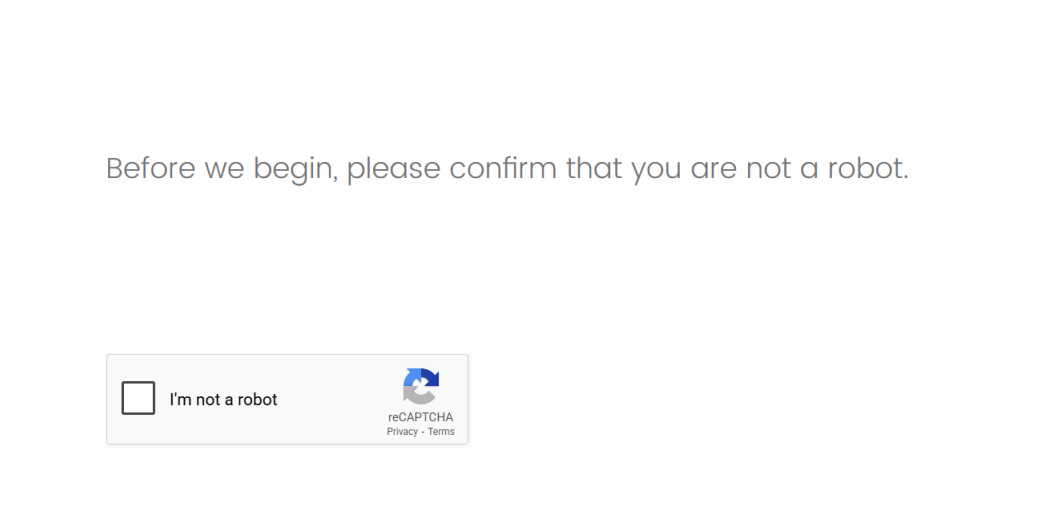


Figure A2: Bot check.

Note: The same procedure was used in Experiments 9 and 10.

E.0.3 Experiment 1

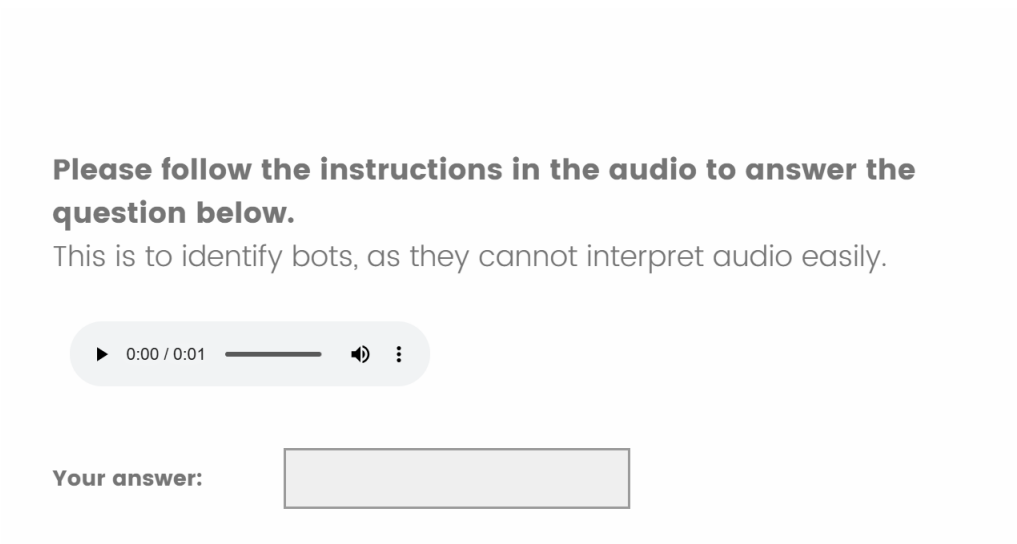


Figure A3: Bot check (audio check).

Note: The same procedure was used in Experiments 9 and 10.

E.0.4 Experiment 1



Figure A4: Bot check (typing speed and pattern).

Note: The same procedure was used in Experiments 9 and 10.

E.0.5 Experiment 1

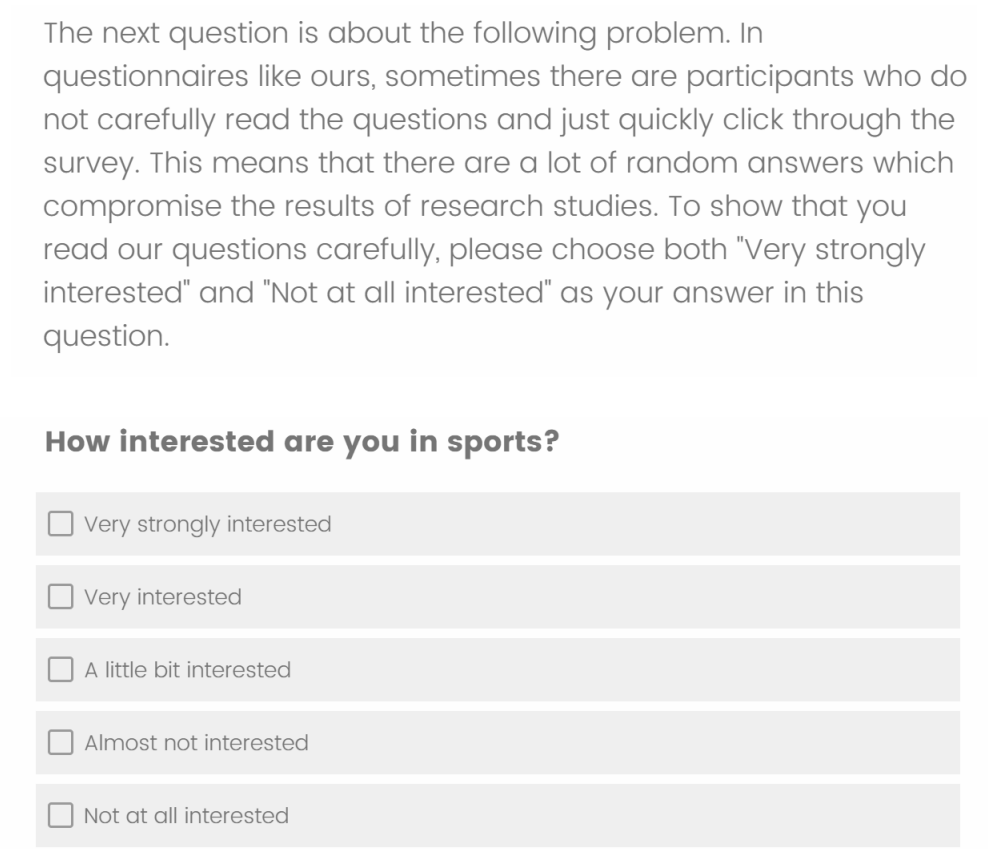


Figure A5: Attention check.

Note: The same procedure was used in Experiments 9 and 10.

E.0.6 Experiment 1 (demographic questions)



Figure A6: Political ideology.

Note: Half respondents were shown this question at the end of the survey. The same procedure was used in Experiments 9 and 10.

E.0.7 Experiment 1 (demographic questions)



Figure A7: Value of investments

Note: Half respondents were shown this question at the end of the survey. The same procedure was used in Experiments 9 and 10.

E.1 Experiment 1

In the following questions, you will be asked about shareholder voting at five major U.S. companies. At shareholder meetings, investors can vote on proposals submitted by other shareholders. The company's management, including the board of directors, typically recommends how shareholders should vote on each proposal. For each company, we will describe two shareholder proposals and ask you to rank **three voting strategies.**

Important: These are not the only possible voting strategies. Other combinations may exist. For the purposes of this study, however, we ask you to rank only the three strategies shown.

Figure A8: Introduction to scenarios

Imagine you own shares in Alphabet (Google). At the upcoming shareholder meeting, investors will vote on how the company should respond to two shareholder proposals. The management (Board) recommends voting AGAINST both proposals below:

• **Human rights**: Conduct a human rights impact assessment of Alphabet's AI-driven targeted advertising policies.
• **Climate**: Provide enhanced disclosure on Alphabet's climate goals.

Rank the strategies according to the voting choice you would most want taken as an owner of the company, by dragging them into your preferred order.

Prioritize human rights (Vote FOR human rights, AGAINST climate)

Prioritize climate (Vote FOR climate, AGAINST human rights)

Support management (Vote AGAINST both proposals)

Figure A9: Alphabet scenario

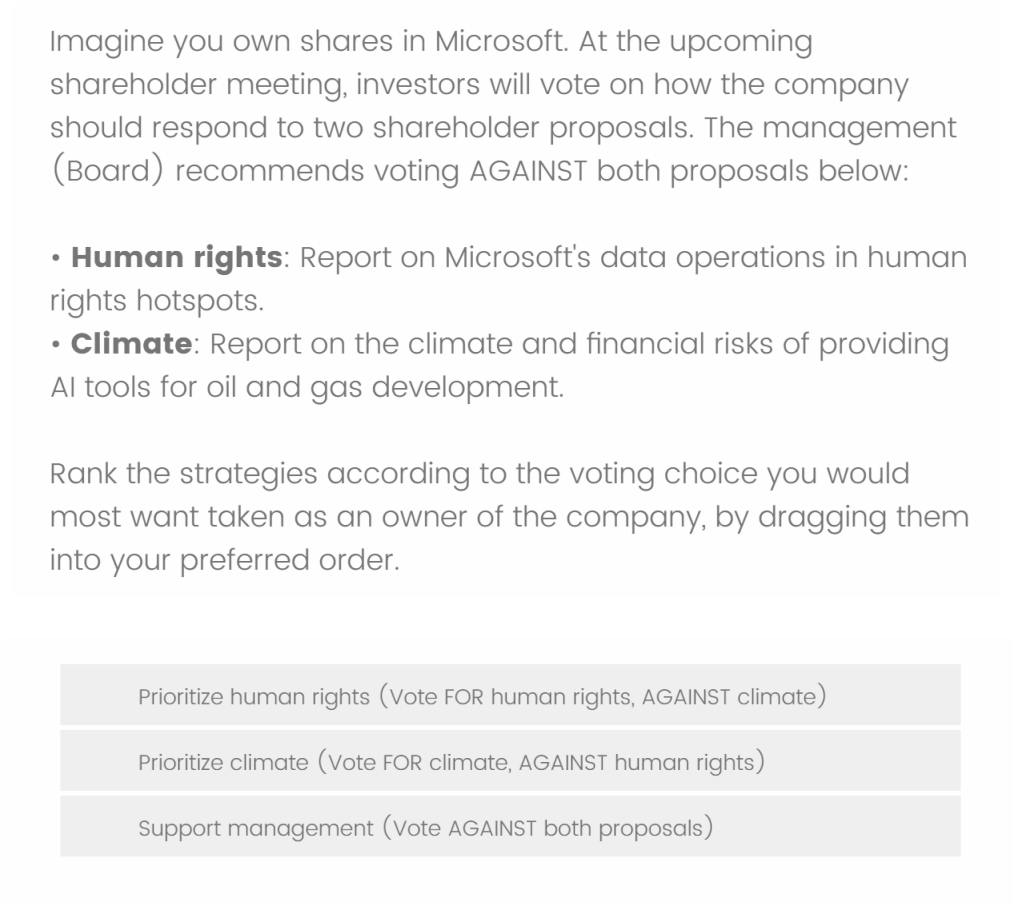


Figure A10: Microsoft scenario

Imagine you own shares in Amazon. At the upcoming shareholder meeting, investors will vote on how the company should respond to two shareholder proposals. The management (Board) recommends voting AGAINST both proposals below:

• **Worker conditions:** Conduct an independent audit on warehouse working conditions.
• **Climate**: Disclose how data center growth affects Amazon's climate goals.

Rank the strategies according to the voting choice you would most want taken as an owner of the company, by dragging them into your preferred order.

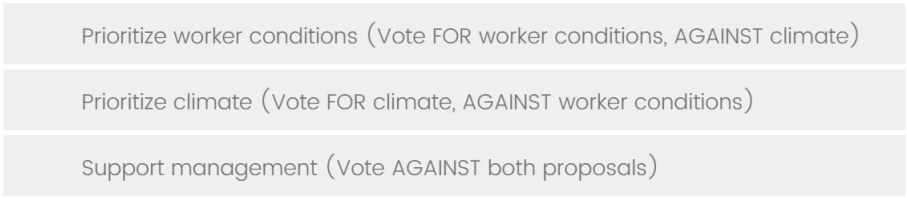


Figure A11: Amazon scenario

Imagine you own shares in Meta Platforms. At the upcoming shareholder meeting, investors will vote on how the company should respond to two shareholder proposals. The management (Board) recommends voting AGAINST both proposals below:

• **Anti-hate**: Issue a report on hate targeting marginalized communities on Meta's platforms.
• **Climate**: Disclose a transition plan to reduce greenhouse gas emissions despite data-center growth.

Rank the strategies according to the voting choice you would most want taken as an owner of the company, by dragging them into your preferred order.

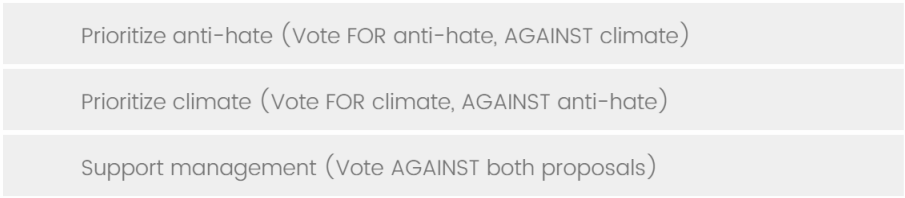


Figure A12: Meta scenario

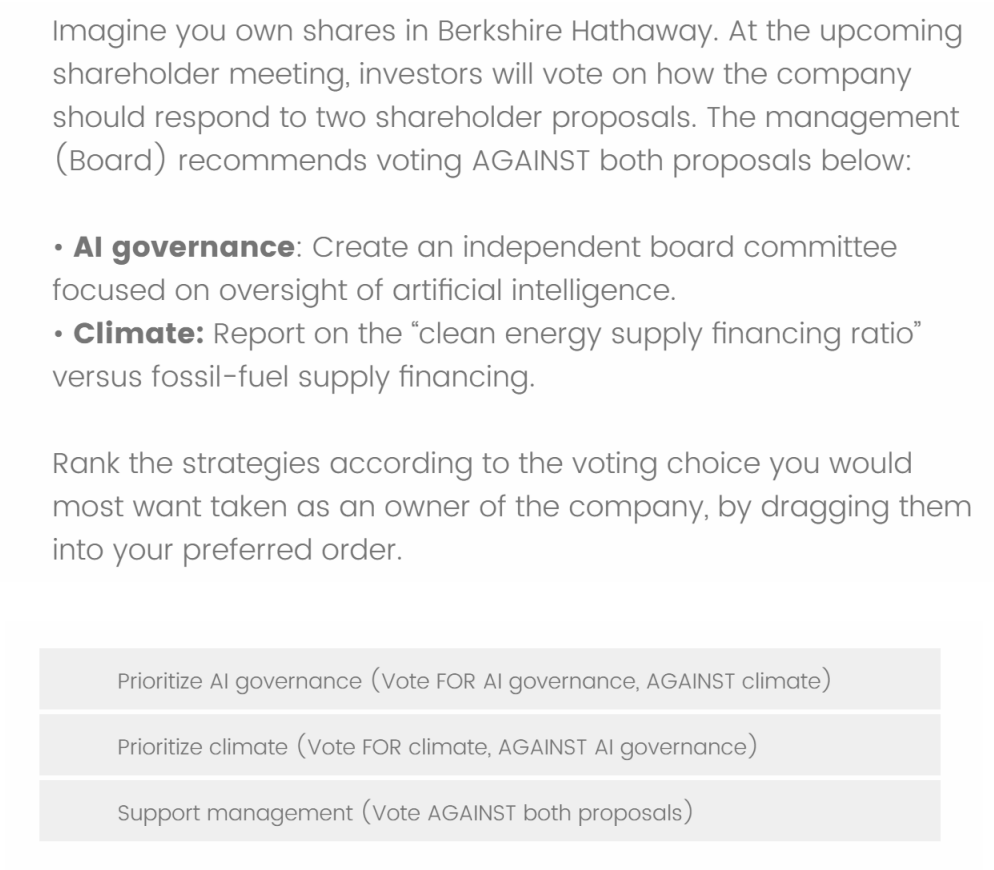


Figure A13: Berkshire Hathaway scenario

E.1.1 WTP elicitation

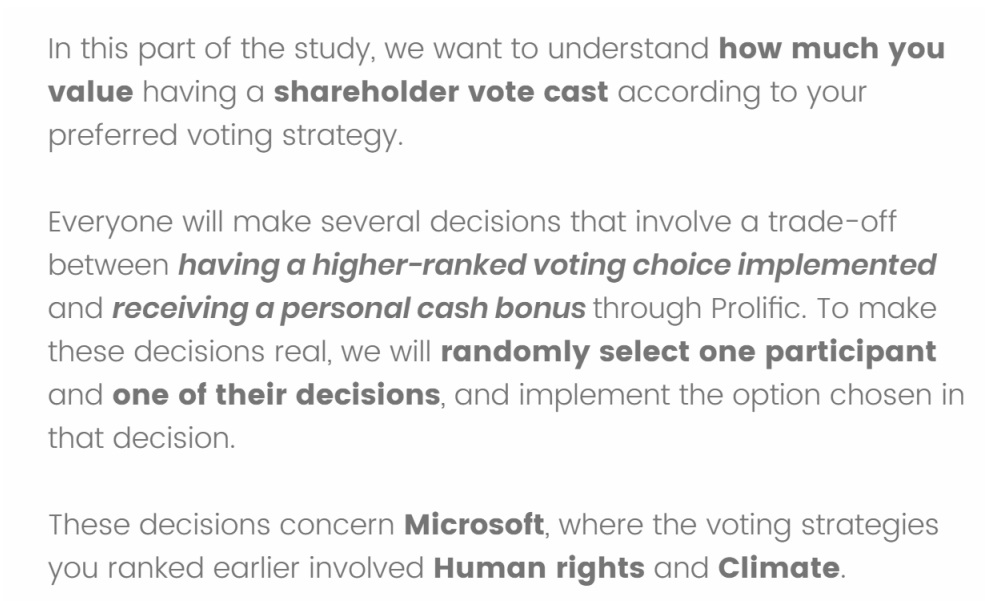


Figure A14: WTP introduction (respondents are randomly assigned to one of the five scenarios above, here Microsoft.

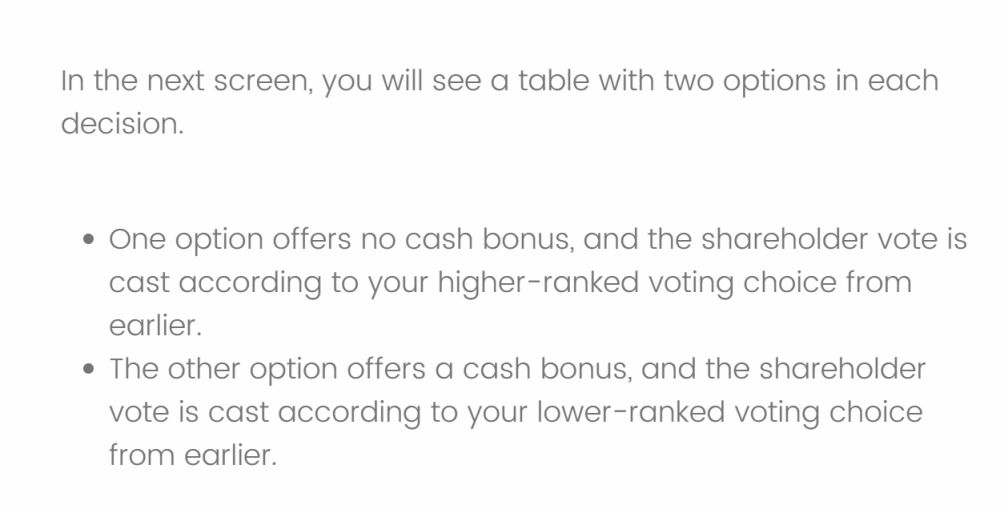


Figure A15: WTP introduction

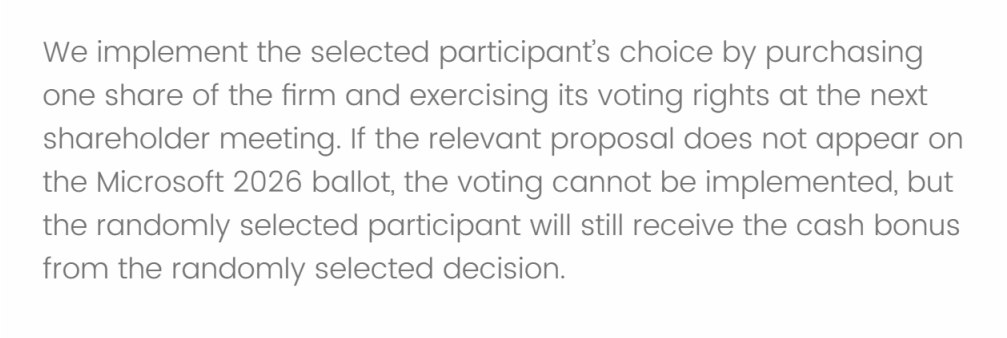


Figure A16: WTP explanation

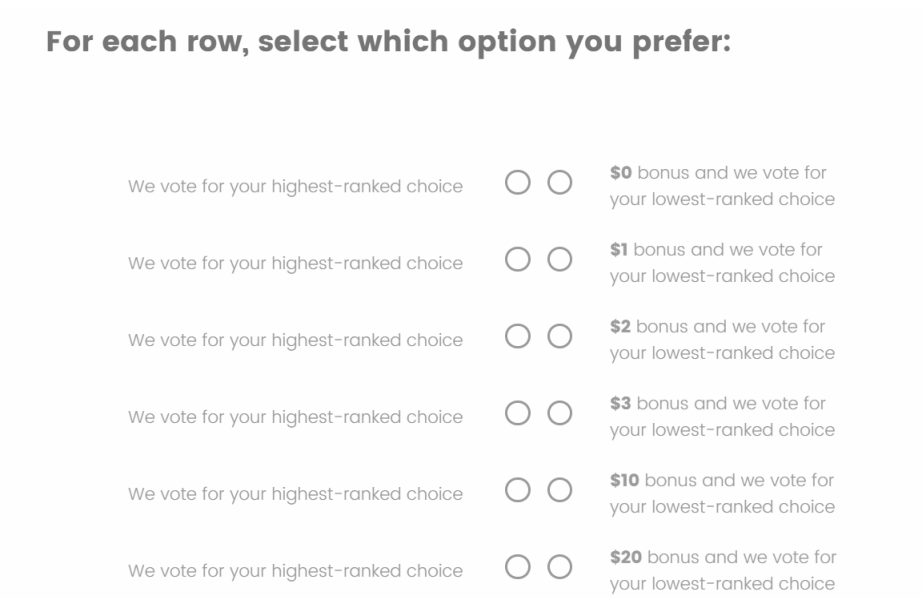


Figure A17: WTP elicitation for first versus third.

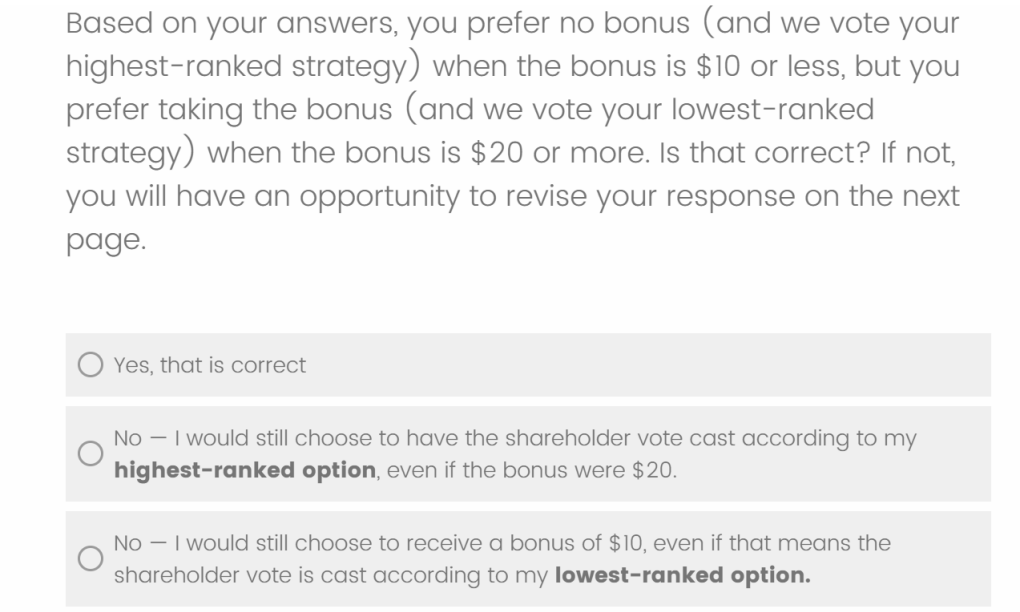


Figure A18: WTP verification (here using the example of a respondent switching between 10 and 20 dollar)

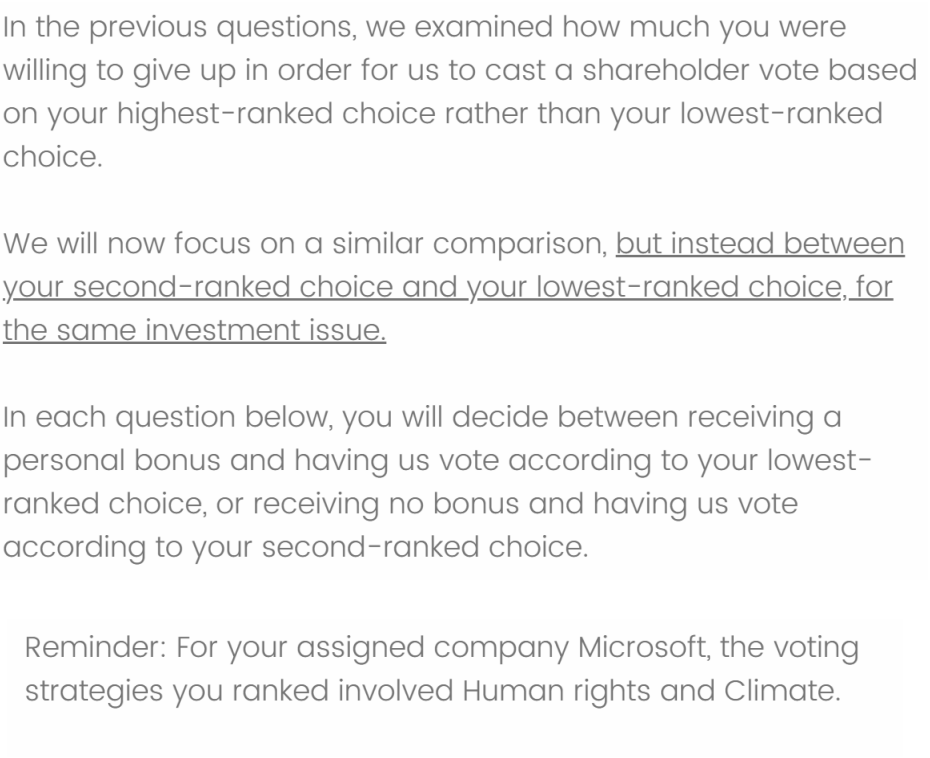


Figure A19: WTP for second vs third

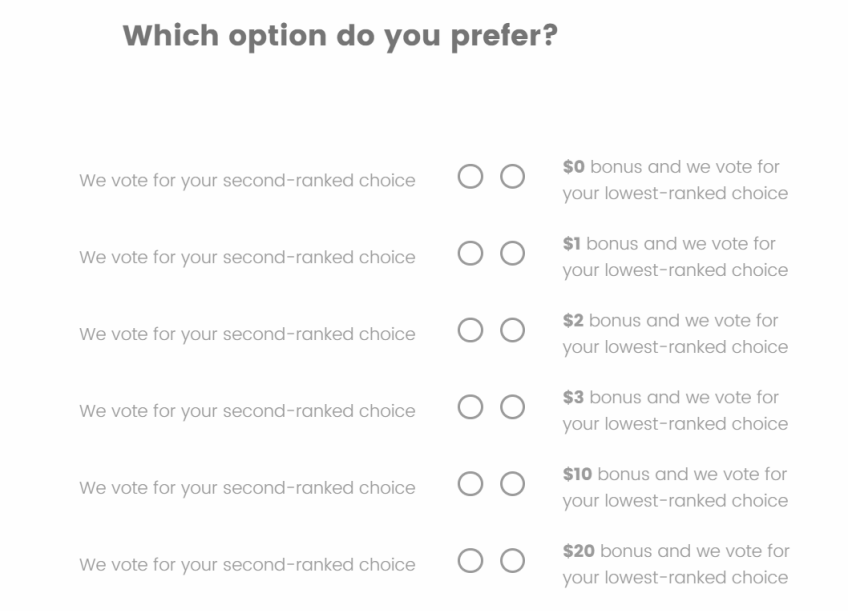


Figure A20: WTP elicitation for second vs third.

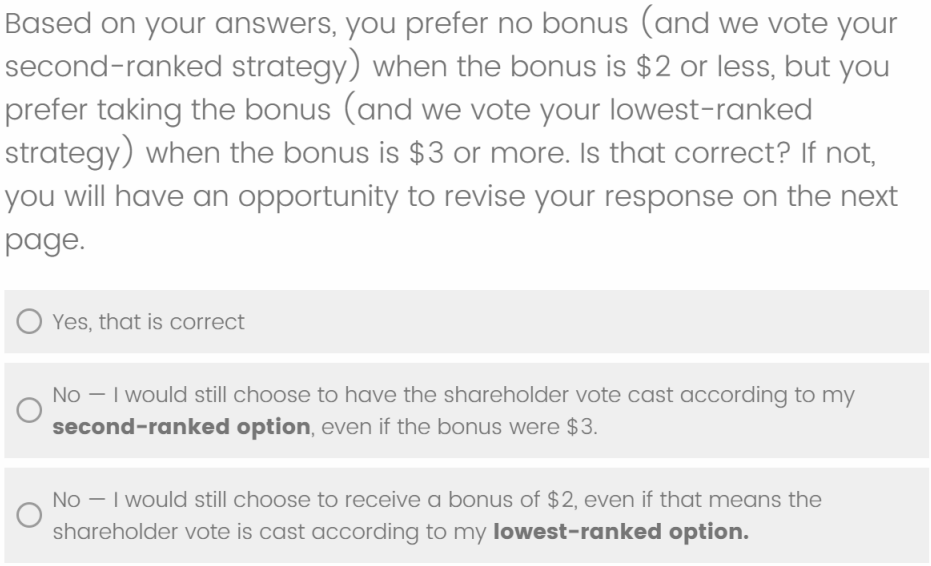


Figure A21: WTP verification (here using the example of a respondent switching between 2 and 3 dollar)

E.2 Experiment 3

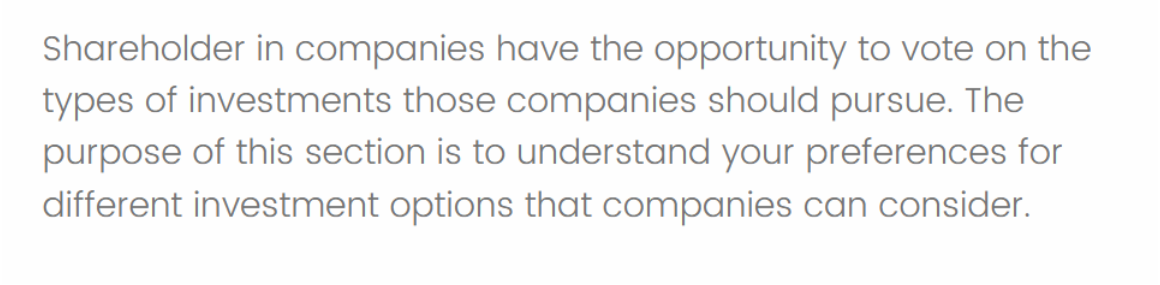


Figure A22: Introduction

E.2.1 Scenario H1 to H4

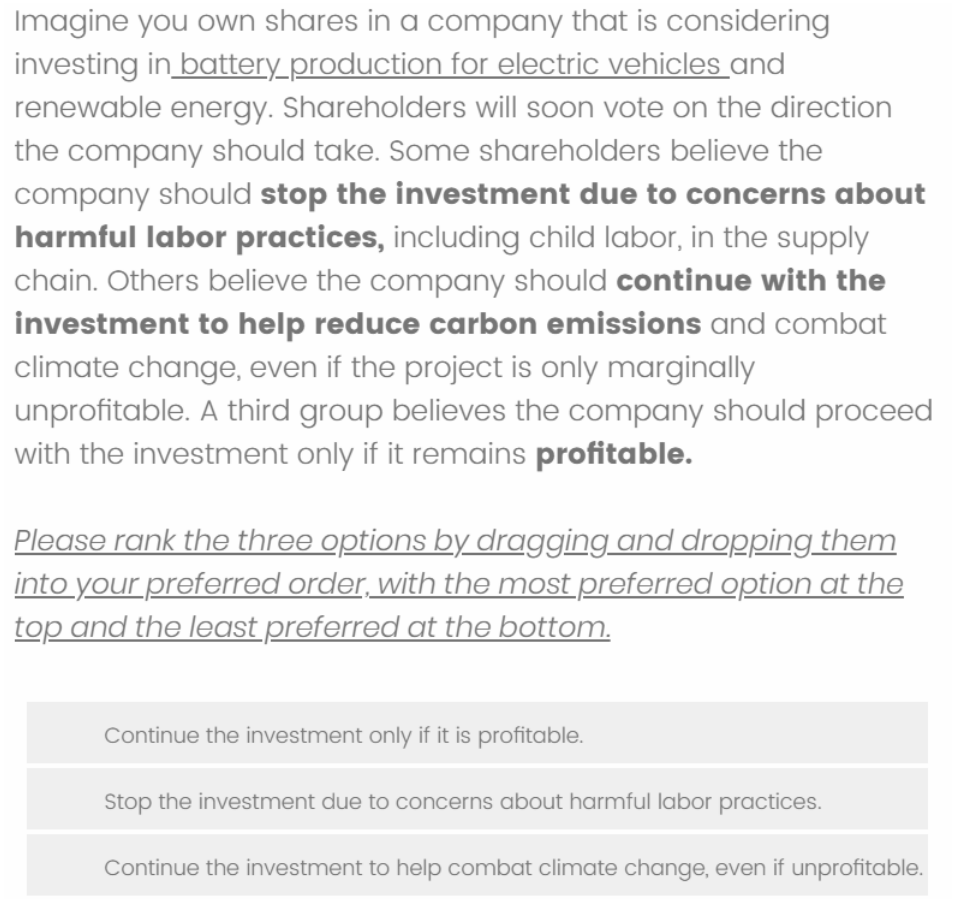


Figure A23: Scenario H1

Imagine you own shares in an agricultural company that is deciding on irrigation methods for large-scale farming in a water-scarce region. Shareholders will soon vote on how to proceed. Some shareholders advocate for implementing a **highly efficient irrigation system,** which conserves water but requires more energy and results in higher carbon emissions. Others argue that the company should use **traditional irrigation,** which consumes more water but requires less energy and produces lower emissions. A third group believes the company should move forward with the project only if is **profitable.**

Please rank the three options by dragging and dropping them into your preferred order, with the most preferred option at the top and the least preferred at the bottom.

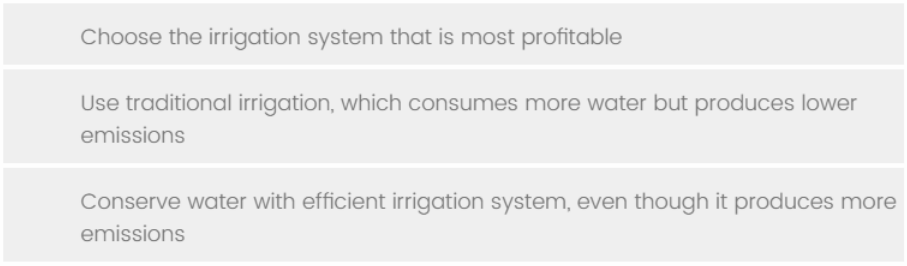


Figure A24: Scenario H2

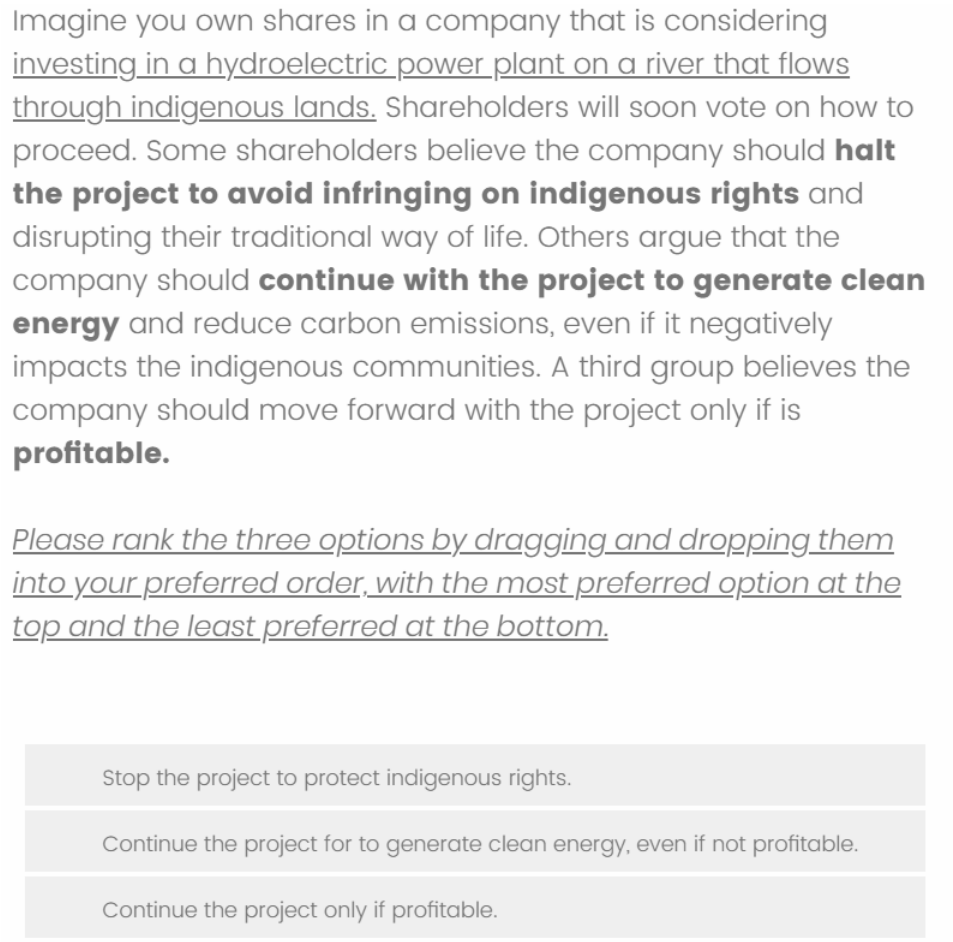


Figure A25: Scenario H3

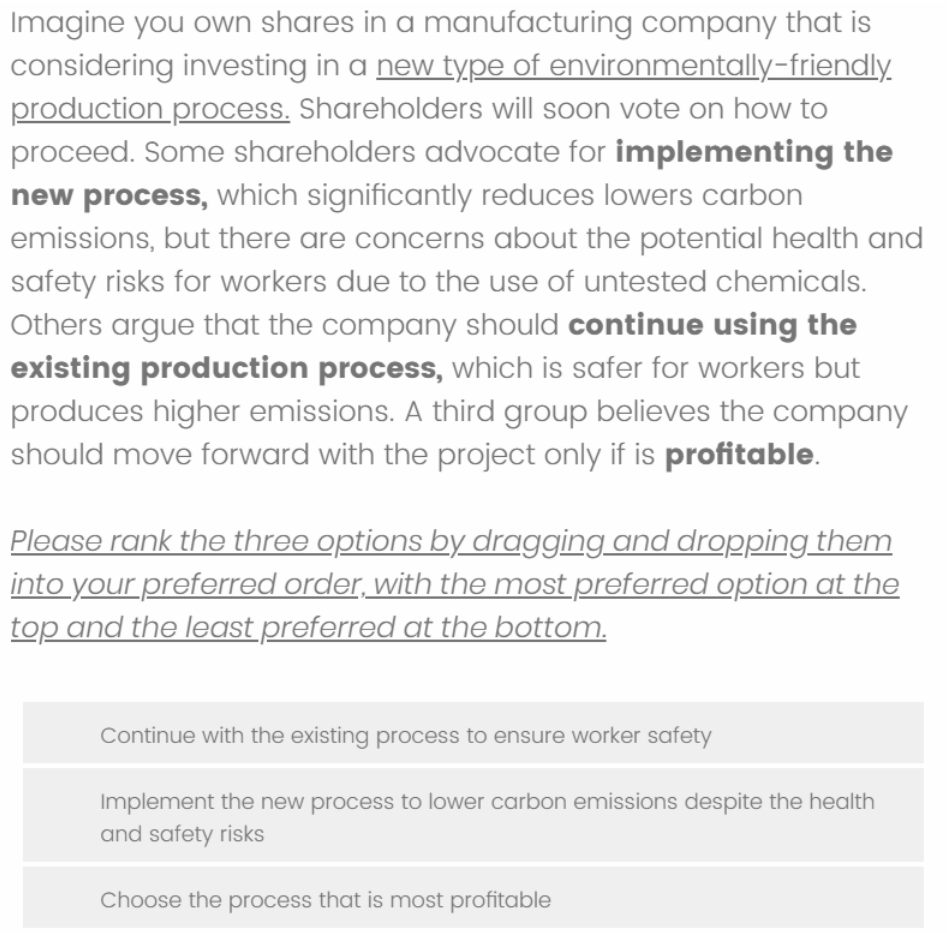


Figure A26: Scenario H4

E.2.2 Scenario E1 to E4

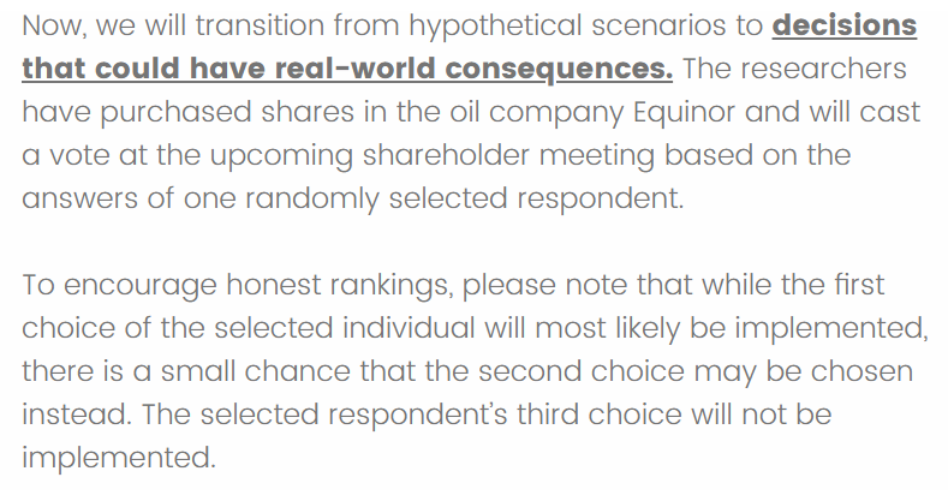


Figure A27: Introduction to Scenarios E1 to E4

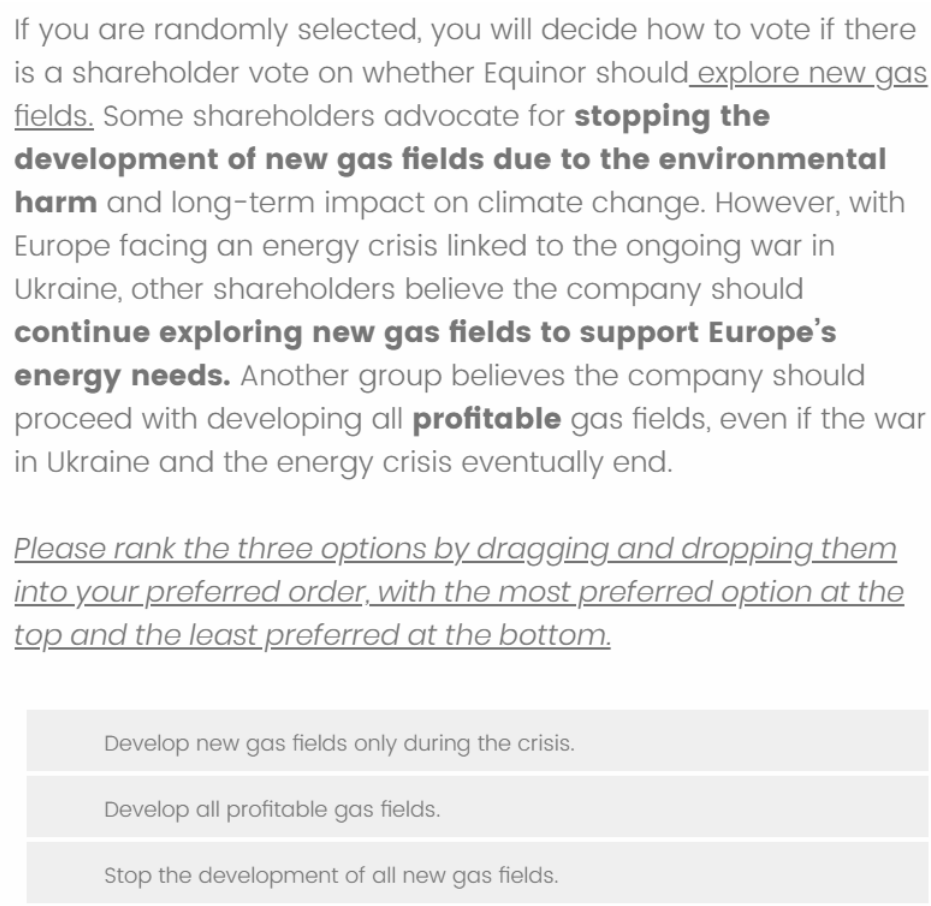


Figure A28: Scenario E1

If you are randomly selected, you will decide how to vote if there is a shareholder vote on whether Equinor should reduce its exports of oil to developing countries. Some shareholders advocate for **reducing oil exports to align with global climate goals** and reduce carbon emissions. However, other shareholders are concerned that **reducing exports could lead to increased poverty** and economic hardship in countries that rely heavily on affordable energy from oil imports, so they want to export even if this becomes less profitable than other projects. A third group believes the company should continue exporting oil to these countries as long as it remains **profitable.**

Please rank the three options by dragging and dropping them into your preferred order, with the most preferred option at the top and the least preferred at the bottom.



Figure A29: Scenario E2

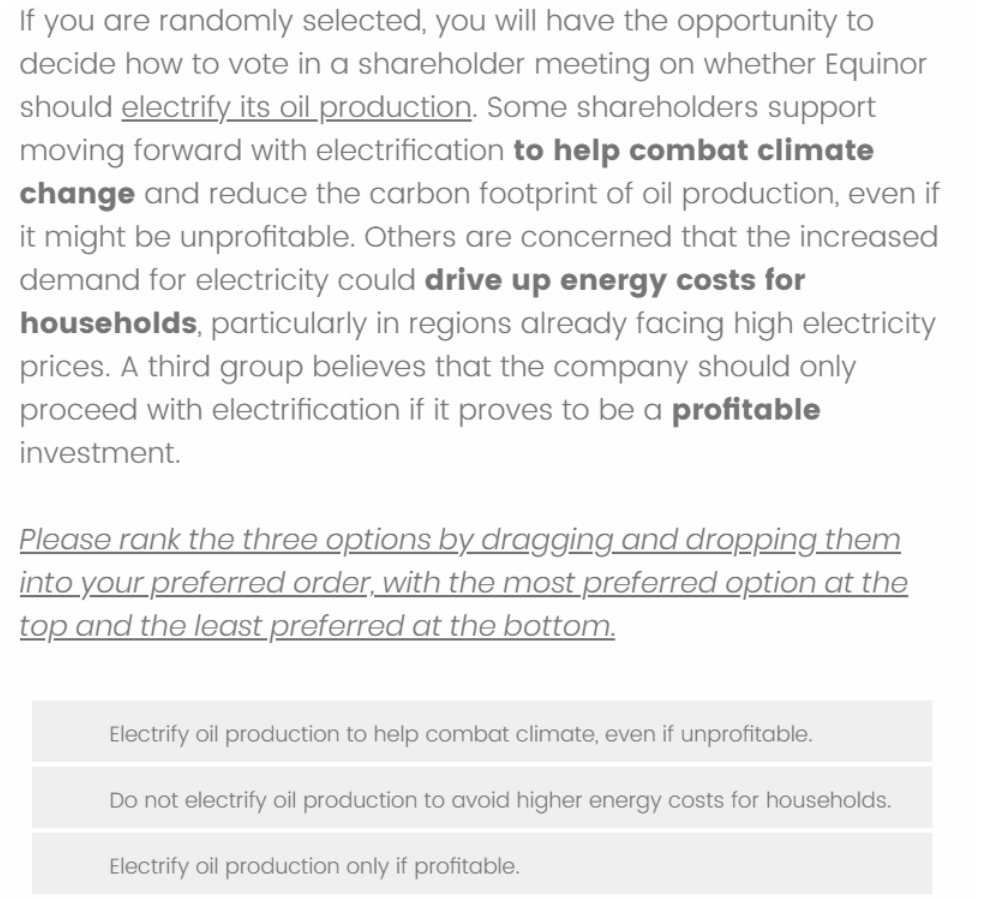


Figure A30: Scenario E3

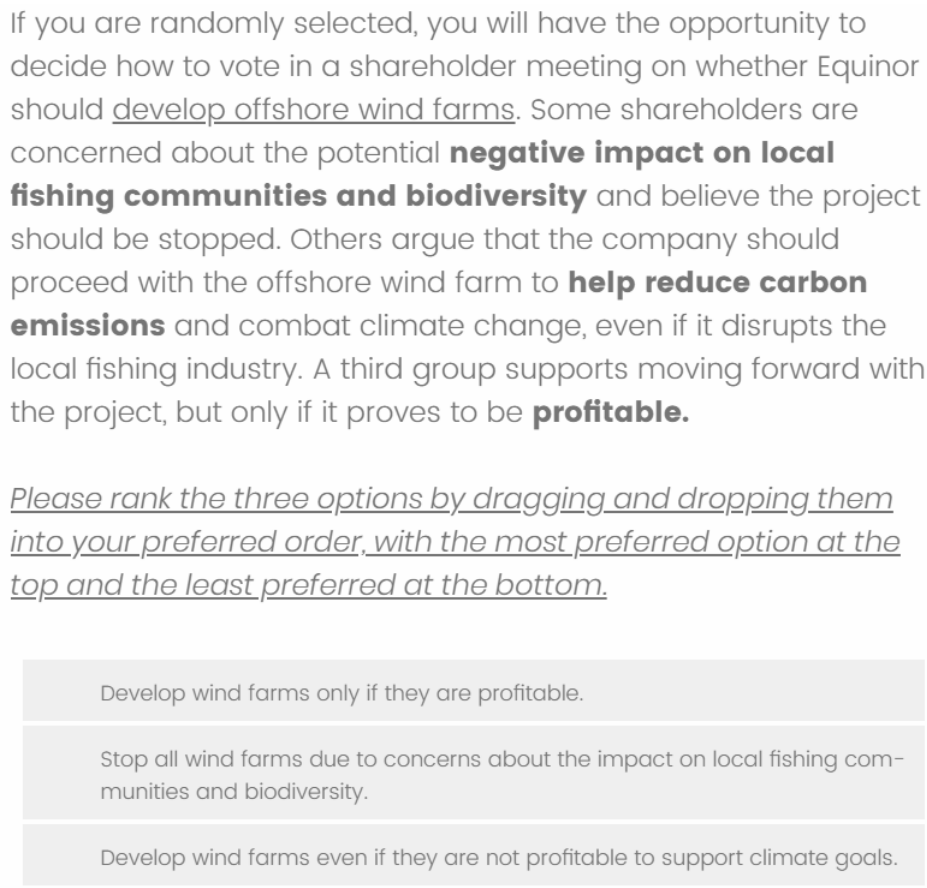


Figure A31: Scenario E4

E.2.3 WTP elicitation

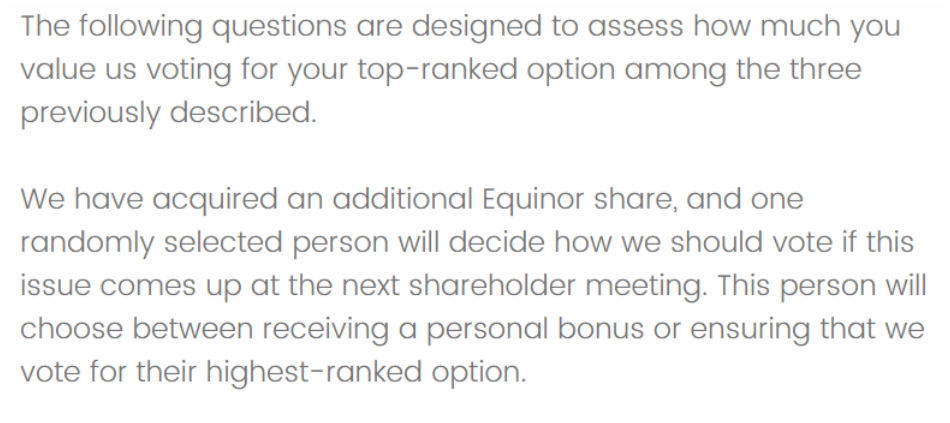


Figure A32: Introduction to WTP elicitation

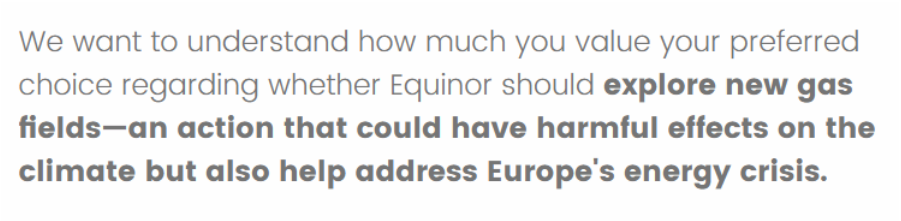


Figure A33: WTP E1

Note: This question was presented to the randomly selected quarter of respondents for whom WTP was elicited for E1

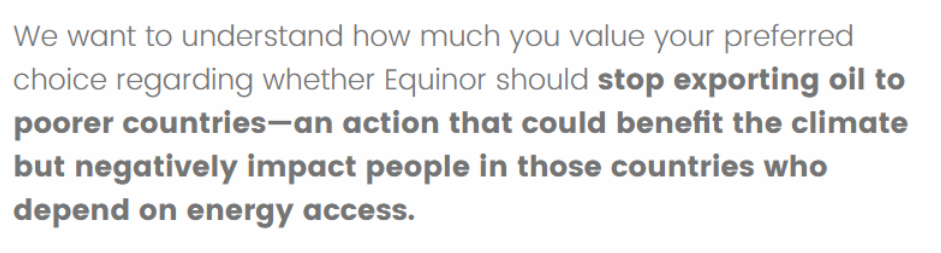


Figure A34: WTP E2

Note: This question was presented to the randomly selected quarter of respondents for whom WTP was elicited for E2

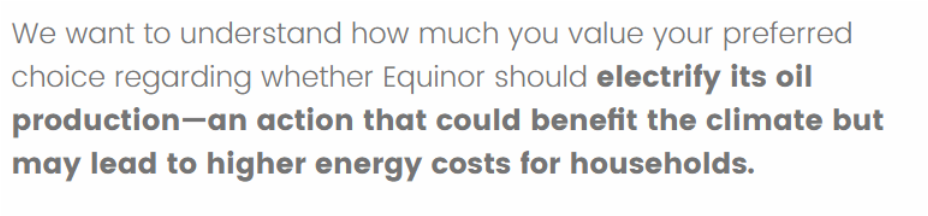


Figure A35: WTP E3

Note: This question was presented to the randomly selected quarter of respondents for whom WTP was elicited for E3

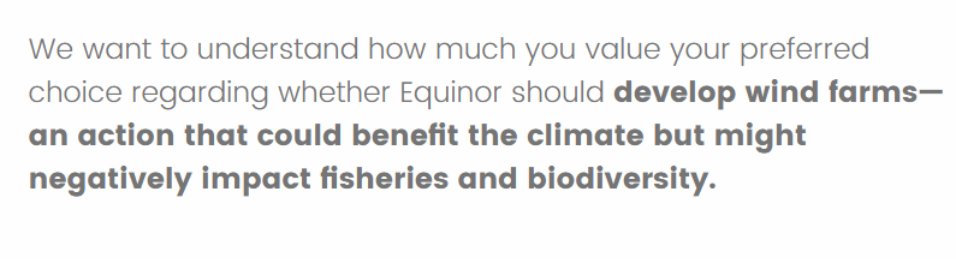


Figure A36: WTP E4

Note: This question was presented to the randomly selected quarter of respondents for whom WTP was elicited for E4

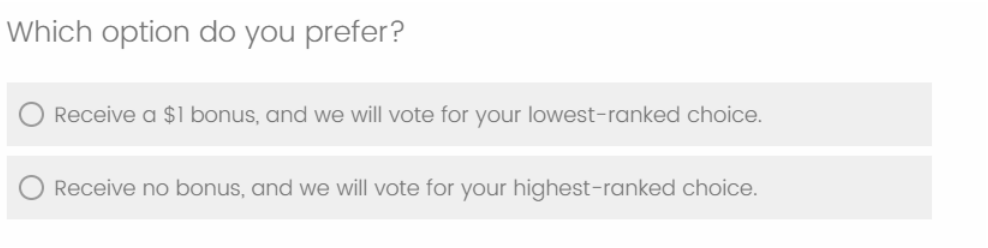


Figure A37: WTP for the First vs. Third Choice: Valuation Question 1

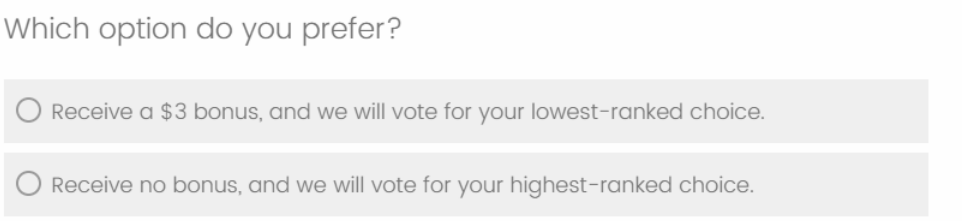


Figure A38: WTP for the First vs. Third Choice: Valuation Question 2

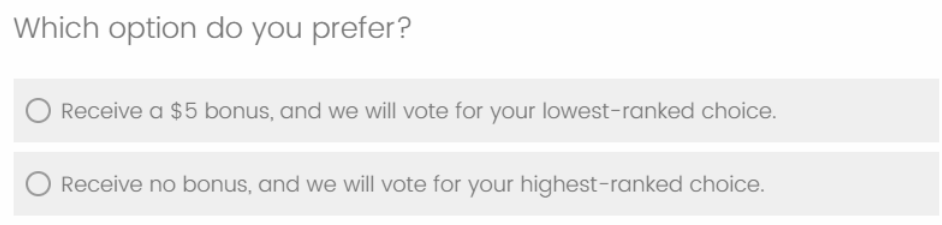


Figure A39: WTP for the First vs. Third Choice: Valuation Question 3

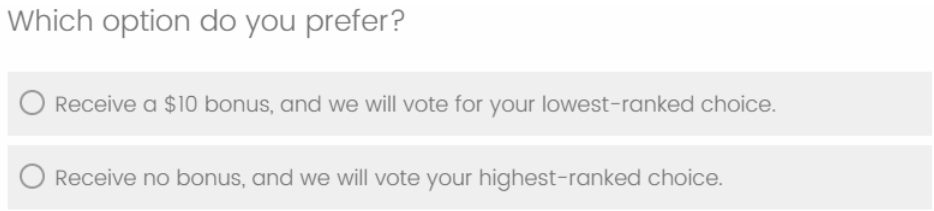


Figure A40: WTP for the First vs. Third Choice: Valuation Question 4

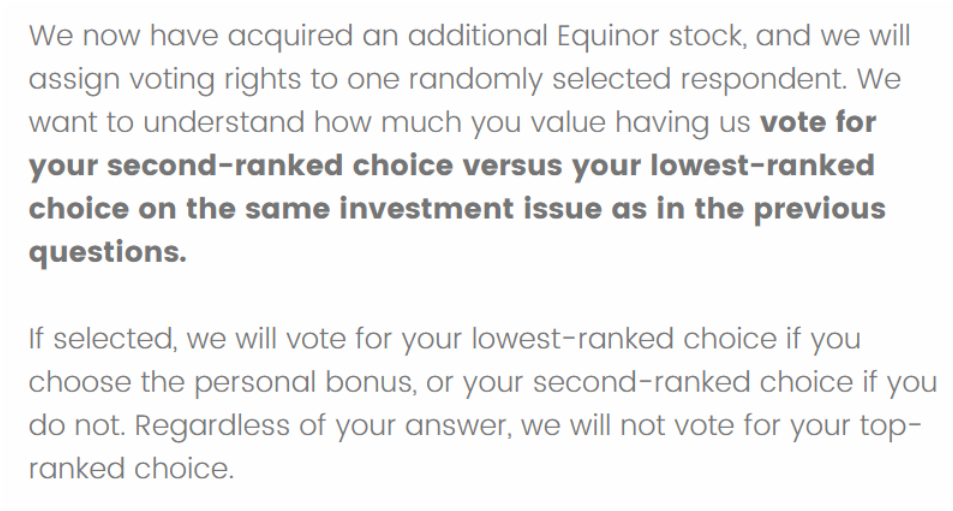


Figure A41: Introduction for WTP for Second vs Third Choice

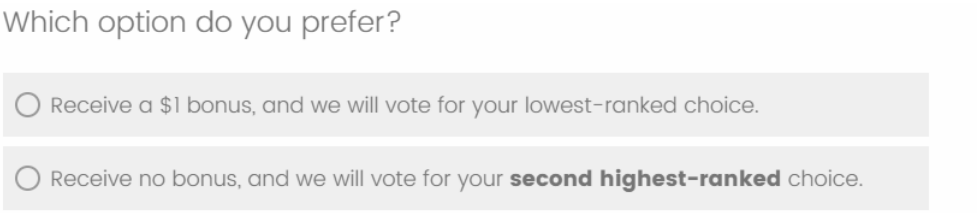


Figure A42: WTP for the Second vs. Third Choice: Valuation Question 1

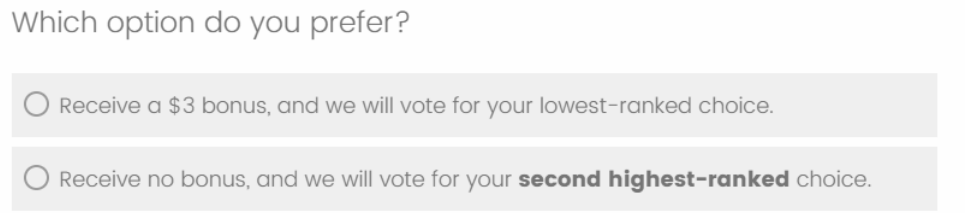


Figure A43: WTP for the Second vs. Third Choice: Valuation Question 2

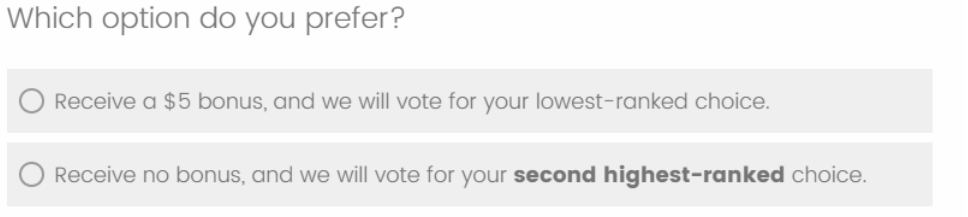


Figure A44: WTP for the Second vs. Third Choice: Valuation Question 3

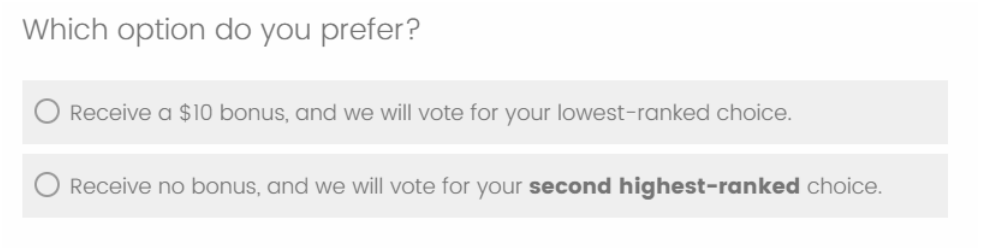


Figure A45: WTP for the Second vs. Third Choice: Valuation Question 4

E.2.4 Other questions



Figure A46: Political Ideology



Figure A47: Value of Investments

E.3 Experiment 4

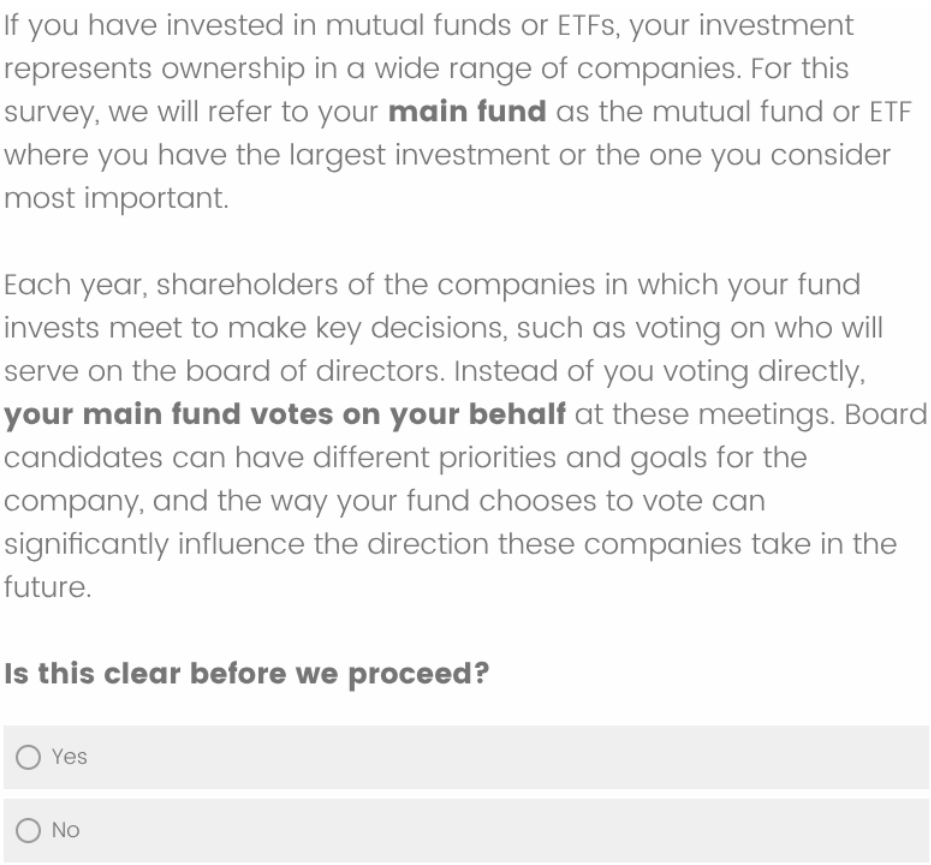


Figure A48: Introduction

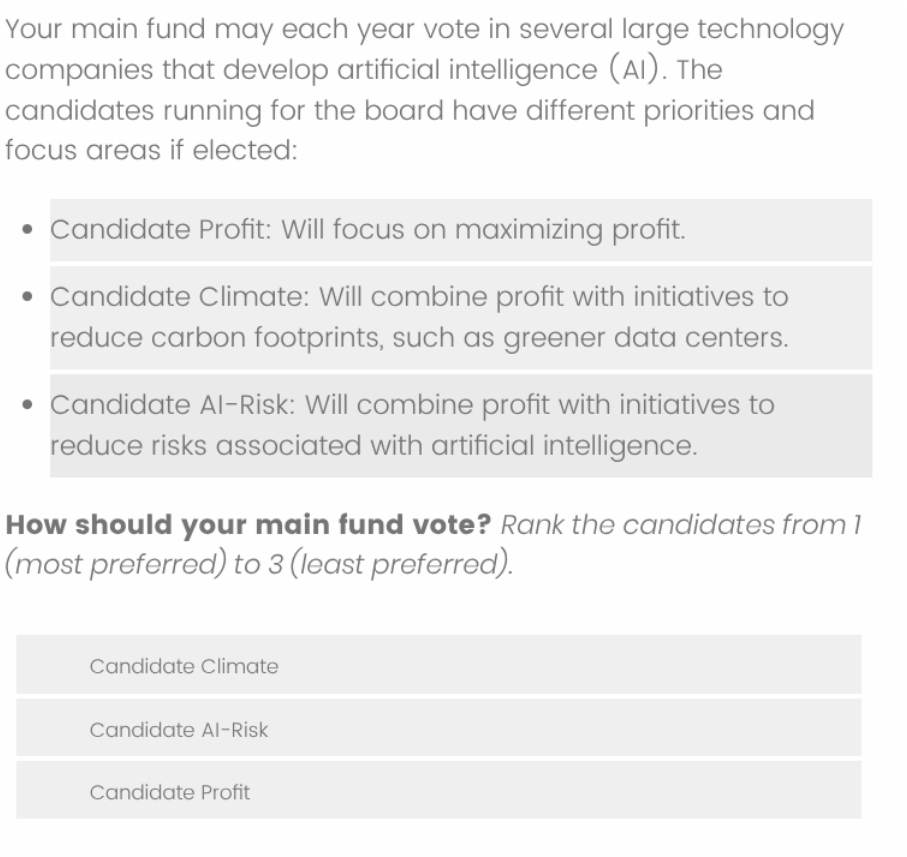


Figure A49: Scenario B1

Your main fund may each year vote in several large manufacturing companies operating factories in low-cost countries. The candidates running for the board have different priorities and focus areas if elected:

- Candidate Profit: Will focus on maximizing profit.
- Candidate Climate: Will combine profit with initiatives for carbon-neutral production.
- Candidate Union: Will combine profit with initiatives to improve opportunities for factory workers to organize through unions.

How should your main fund vote? *Rank the candidates from 1 (most preferred) to 3 (least preferred).*

Candidate Union

Candidate Profit

Candidate Climate

Figure A50: Scenario B2

Your main fund may each year vote in several large financial companies. The candidates running for the board have different priorities and focus areas if elected:

- Candidate Profit: Will focus on maximizing profit.
- Candidate Climate: Will combine profit with initiatives for green investments.
- Candidate Gender Quotas: Will combine profit with the introduction of policies to increase the proportion of women in the company's leadership.

How should your main fund vote? *Rank the candidates from 1 (most preferred) to 3 (least preferred).*

Candidate Climate

Candidate Gender Quotas

Candidate Profit

Figure A51: Scenario B3

Your main fund may each year vote in several large biotechnology companies. The candidates running for the board have different priorities and focus areas if elected:

- Candidate Profit: Will focus on maximizing profit.
- Candidate Climate: Will combine profit with initiatives to reduce the environmental impact of medical production.
- Candidate Medicine Access: Will combine profit with initiatives to improve global access to medicine, including in low-income countries.

How should your main fund vote? *Rank the candidates from 1 (most preferred) to 3 (least preferred).*

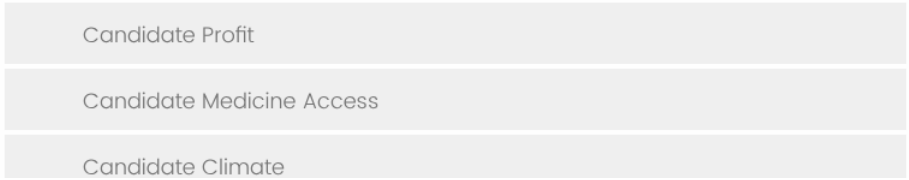


Figure A52: Scenario B4

Your main fund may each year vote in several large technology companies that face criticism for their handling of user data. The candidates running for the board have different priorities and focus areas if elected:

- Candidate Profit: Will focus on maximizing profit.
- Candidate Climate: Will combine profit with initiatives to reduce the company's carbon footprint, such as sustainable data centers.
- Candidate Data Security: Will combine profit with initiatives to secure user data through strict privacy policies and investments in security systems.

How should your main fund vote? *Rank the candidates from 1 (most preferred) to 3 (least preferred).*

Candidate Profit

Candidate Climate

Candidate Data Security

Figure A53: Scenario B5

Your main fund may each year vote in several large food companies that rely on intensive animal farming for meat and dairy production. The candidates running for the board have different priorities and focus areas if elected:

- Candidate Profit: Will focus on maximizing profit.
- Candidate Climate: Will combine profit with initiatives to reduce carbon emissions from agricultural production, such as transitioning to plant-based foods.
- Candidate Animal Welfare: Will combine profit with initiatives to improve animal welfare, including less intensive farming and humane slaughter practices.

How should your main fund vote? *Rank the candidates from 1 (most preferred) to 3 (least preferred).*

Candidate Animal Welfare

Candidate Climate

Candidate Profit

Figure A54: Scenario B6

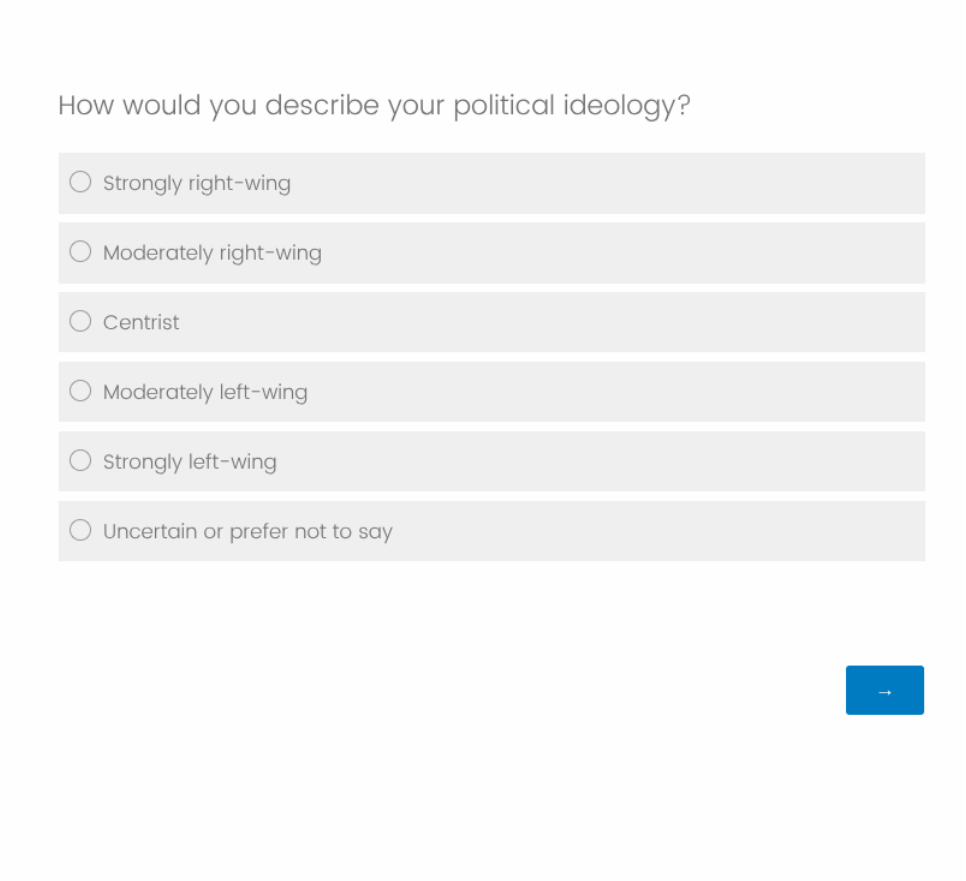


Figure A55: Political Ideology

E.4 Experiment 5

E.4.1 Norwegian language



Figure A56: Political Ideology



Figure A57: Introduction (Note: Proceeding to the next question was only possible after a set time had passed)

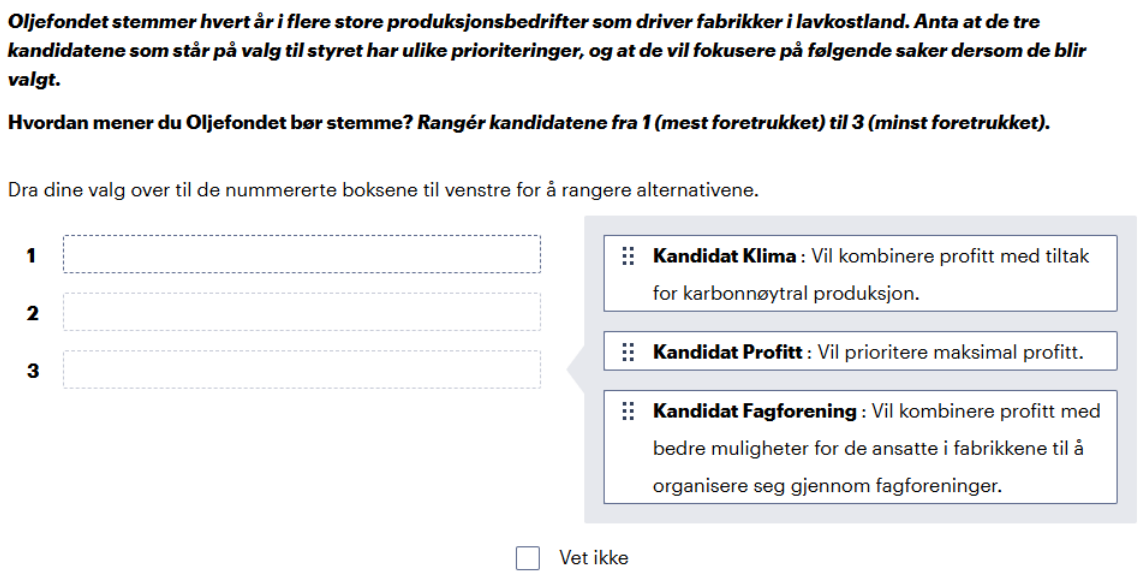


Figure A58: Scenario N1 (Appearance Before Any Rankings Were Made)



Figure A59: Scenario N1 (Example of Appearance After Ranking Two Candidates as First and Second)

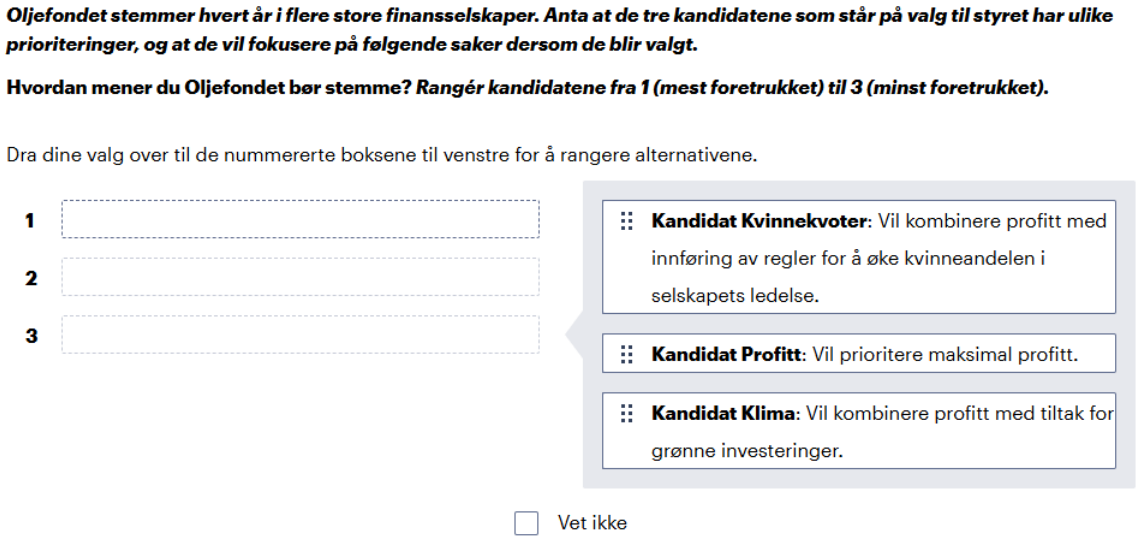


Figure A60: Scenario N2

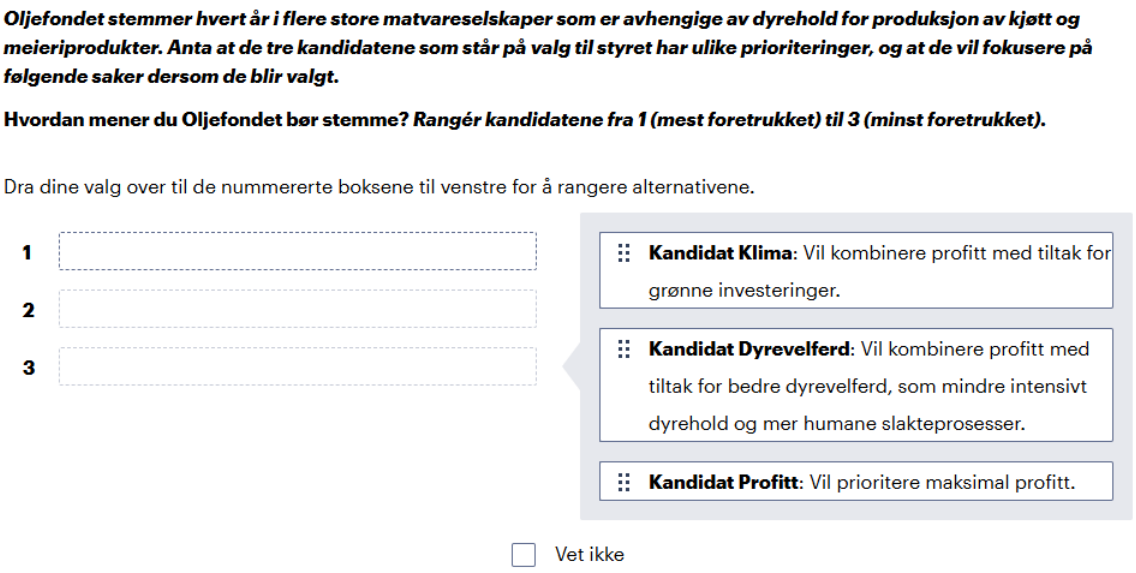


Figure A61: Scenario N3

E.4.2 Translated version



Figure A62: Political Ideology

Note: A few minor adjustments were made to the wording of the political ideology question, ensuring consistency with YouGov Norway's standard phrasing for Norwegian audiences, compared to Experiments 3 and 4.



Figure A63: Introduction

Note: I use the term 'Norwegian Sovereign Wealth Fund' in the translation, whereas the direct translation of the fund's name would be 'the Oil Fund', which is the term typically used in Norway, as the capital originates from oil revenues.

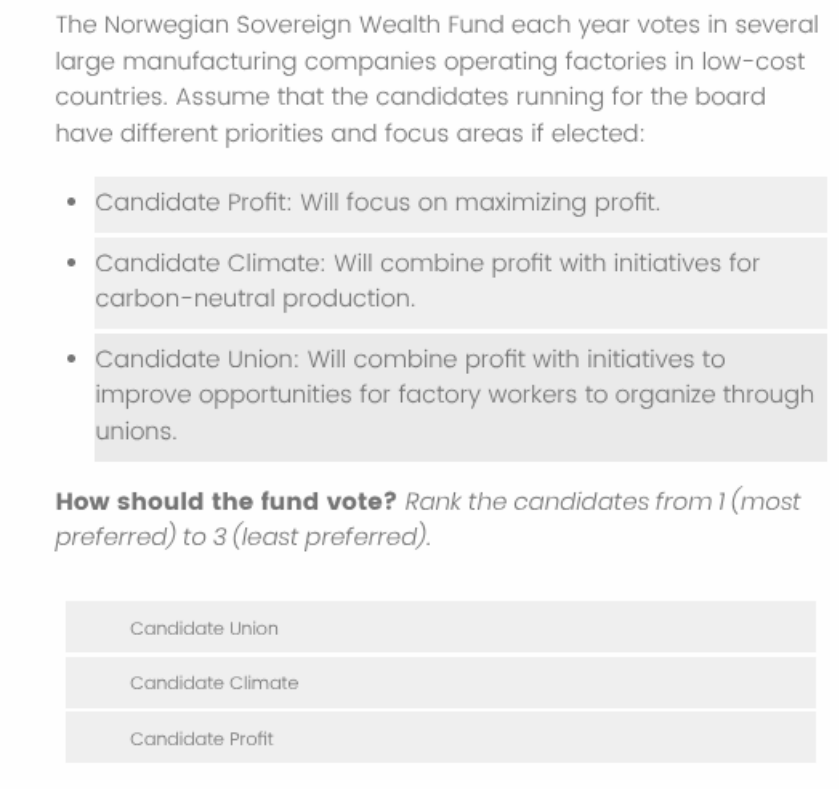


Figure A64: Scenario N1

The Norwegian Sovereign Wealth Fund each year votes in several large financial companies. Assume that the candidates running for the board have different priorities and focus areas if elected:

- Candidate Profit: Will focus on maximizing profit.
- Candidate Climate: Will combine profit with initiatives for green investments.
- Candidate Gender Quotas: Will combine profit with the introduction of policies to increase the proportion of women in the company's leadership.

How should the fund vote? *Rank the candidates from 1 (most preferred) to 3 (least preferred).*

Candidate Profit

Candidate Gender Quotas

Candidate Climate

Figure A65: Scenario N2

The Norwegian Sovereign Wealth Fund each year votes in several large food companies that rely on intensive animal farming for meat and dairy production. Assume that the candidates running for the board have different priorities and focus areas if elected:

- Candidate Profit: Will focus on maximizing profit.
- Candidate Climate: Will combine profit with initiatives to reduce carbon emissions from agricultural production.
- Candidate Animal Welfare: Will combine profit with initiatives to improve animal welfare, including less intensive farming and humane slaughter practices.

How should the fund vote? *Rank the candidates from 1 (most preferred) to 3 (least preferred).*

Candidate Animal Welfare

Candidate Climate

Candidate Profit

Figure A66: Scenario N3

E.5 Experiment 6

In the following questions, you will be asked about shareholder voting at five major U.S. companies. At shareholder meetings, investors can vote on proposals submitted by other shareholders. The company's management, including the board of directors, typically recommends how shareholders should vote on each proposal. For each company, we will describe two shareholder proposals and ask you to rank four **voting strategies.**

Figure A67: Introduction

Imagine you own shares in Alphabet (Google). At the upcoming shareholder meeting, investors will vote on how the company should respond to two shareholder proposals. The management (Board) recommends voting AGAINST both proposals below:

• **Human rights**: Conduct a human rights impact assessment of Alphabet's AI-driven targeted advertising policies.
• **Climate**: Provide enhanced disclosure on Alphabet's climate goals.

Rank the strategies according to the voting choice you would most want taken as an owner of the company, by dragging them into your preferred order.

Support management (Vote AGAINST both proposals)

Prioritize both issues (Vote FOR human rights, FOR climate)

Prioritize climate (Vote FOR climate, AGAINST human rights)

Prioritize human rights (Vote FOR human rights, AGAINST climate)

Figure A68: Alphabet scenario

Imagine you own shares in Microsoft. At the upcoming shareholder meeting, investors will vote on how the company should respond to two shareholder proposals. The management (Board) recommends voting AGAINST both proposals below:

• **Human rights**: Report on Microsoft's data operations in human rights hotspots.
• **Climate**: Report on the climate and financial risks of providing AI tools for oil and gas development.

Rank the strategies according to the voting choice you would most want taken as an owner of the company, by dragging them into your preferred order.



Figure A69: Microsoft scenario

Imagine you own shares in Amazon. At the upcoming shareholder meeting, investors will vote on how the company should respond to two shareholder proposals. The management (Board) recommends voting AGAINST both proposals below:

• **Worker conditions:** Conduct an independent audit on warehouse working conditions.
• **Climate**: Disclose how data center growth affects Amazon's climate goals.

Rank the strategies according to the voting choice you would most want taken as an owner of the company, by dragging them into your preferred order.

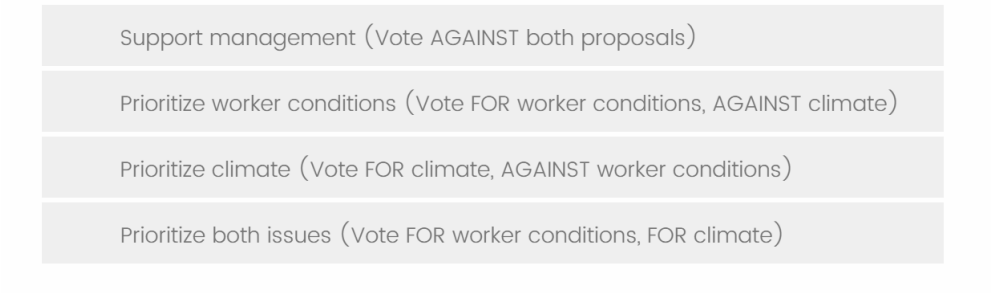


Figure A70: Amazon scenario

Imagine you own shares in Meta Platforms. At the upcoming shareholder meeting, investors will vote on how the company should respond to two shareholder proposals. The management (Board) recommends voting AGAINST both proposals below:

• **Anti-hate**: Issue a report on hate targeting marginalized communities on Meta's platforms.
• **Climate**: Disclose a transition plan to reduce greenhouse gas emissions despite data-center growth.

Rank the strategies according to the voting choice you would most want taken as an owner of the company, by dragging them into your preferred order.



Figure A71: Meta scenario

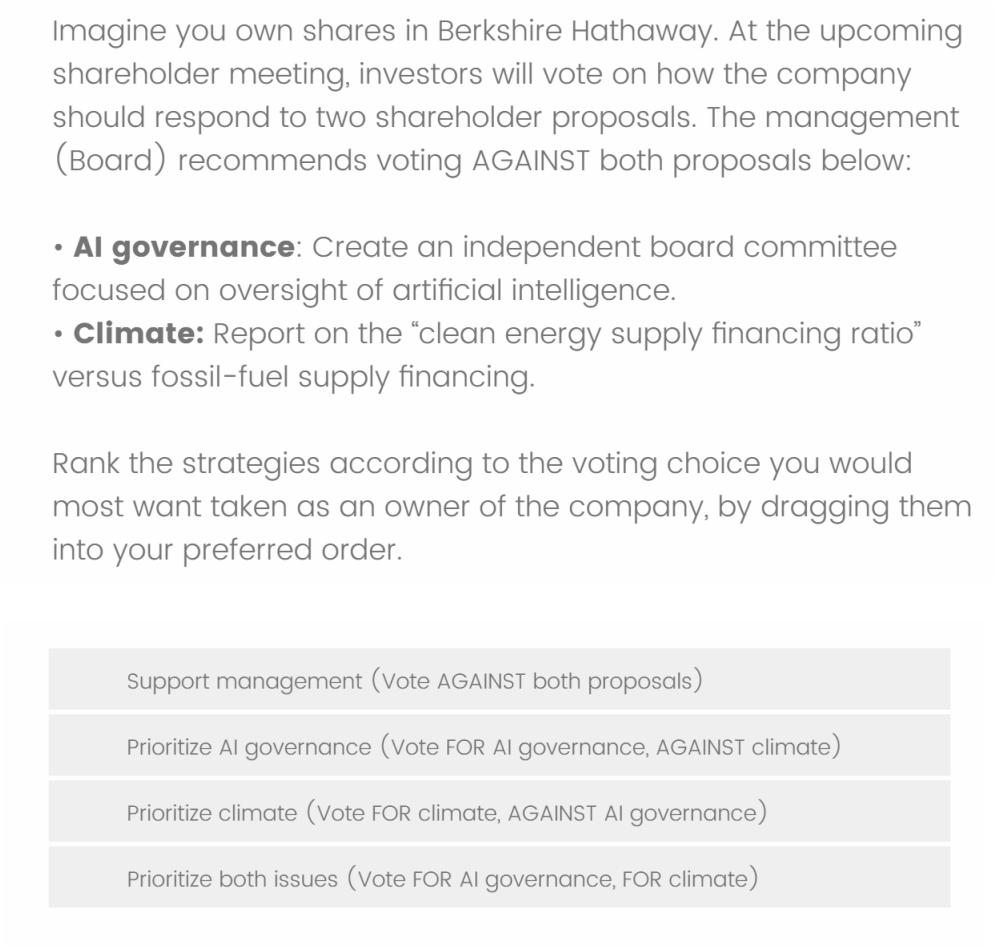


Figure A72: Berkshire Hathaway scenario

E.6 Experiment 8

E.6.1 Norwegian language

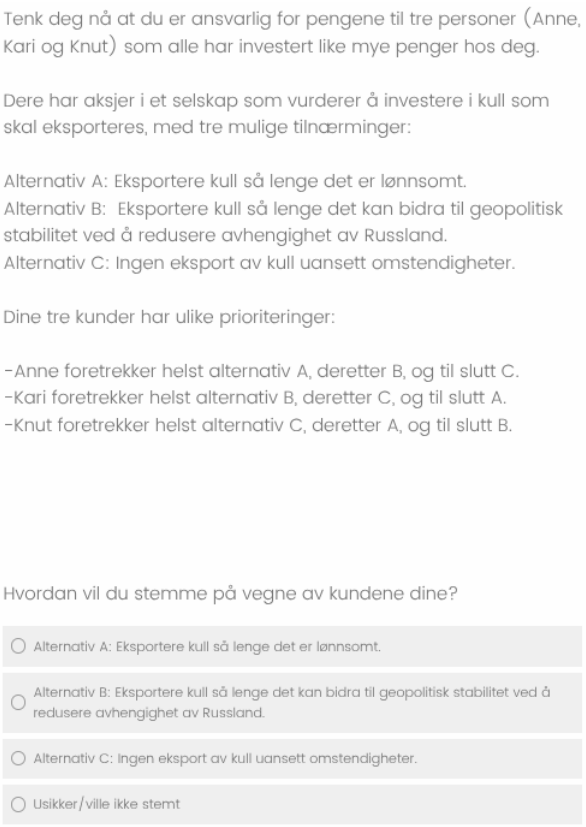


Figure A73: Description of Scenario



Figure A74: Approaches for Resolving Cyclical Preferences

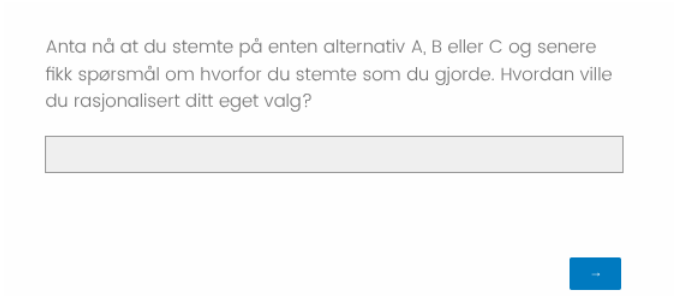


Figure A75: Open-Ended Question on Justifying the Aggregation Rule

E.6.2 Translated version

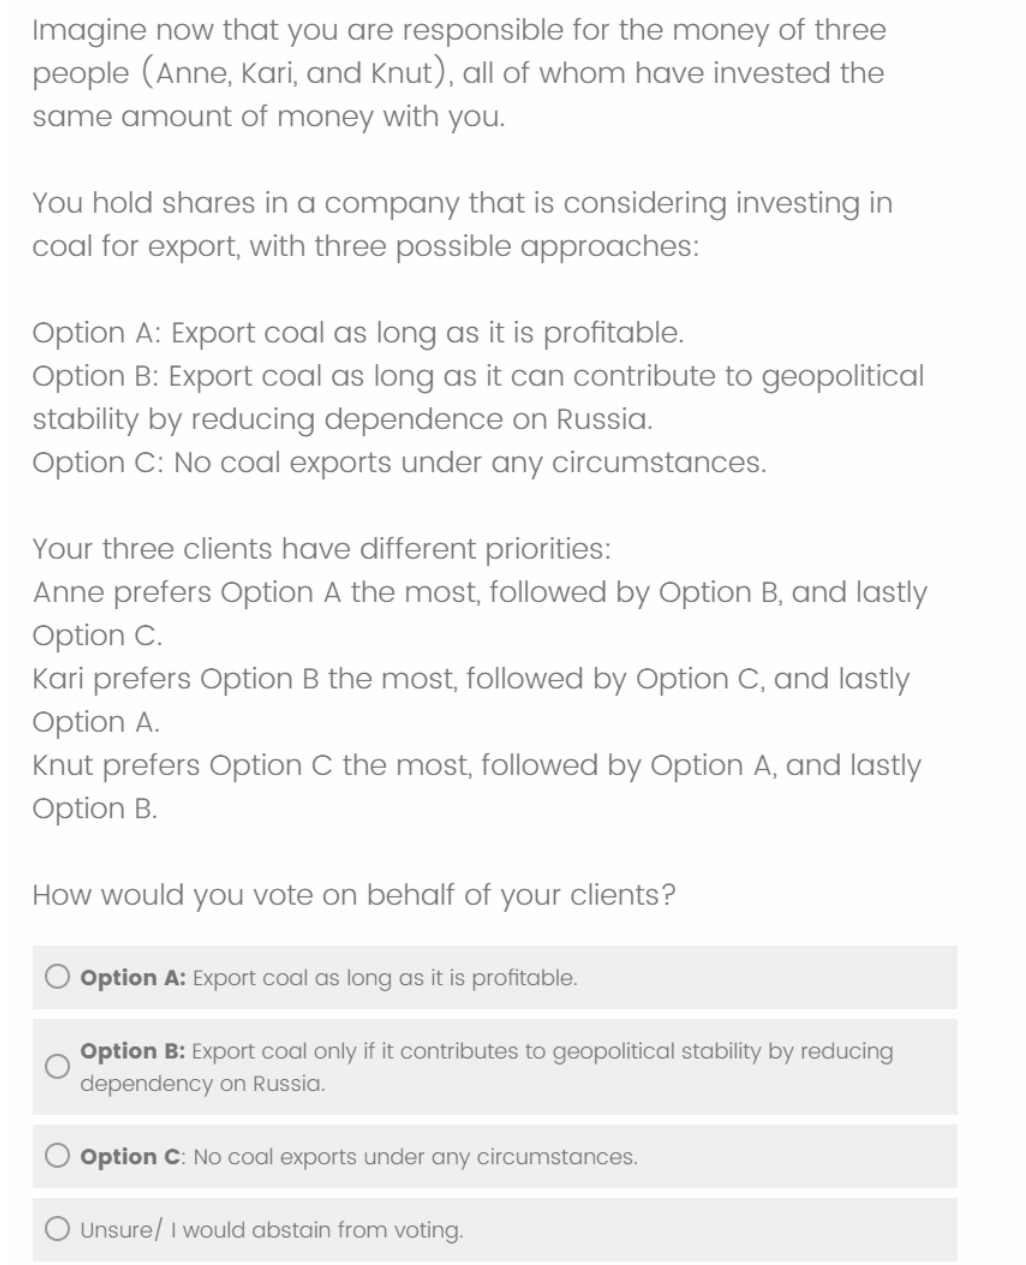


Figure A76: Description of Scenario

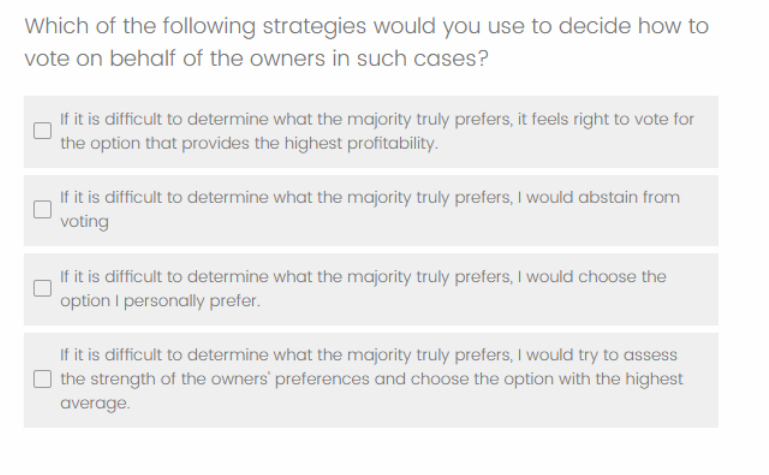


Figure A77: Approaches for Resolving Cyclical Preferences

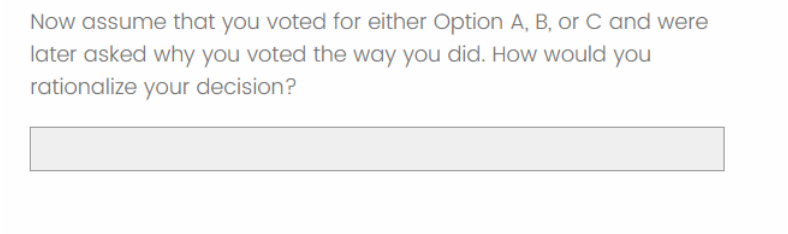


Figure A78: Open-Ended Question on Justifying the Aggregation Rule

E.7 Experiment 10

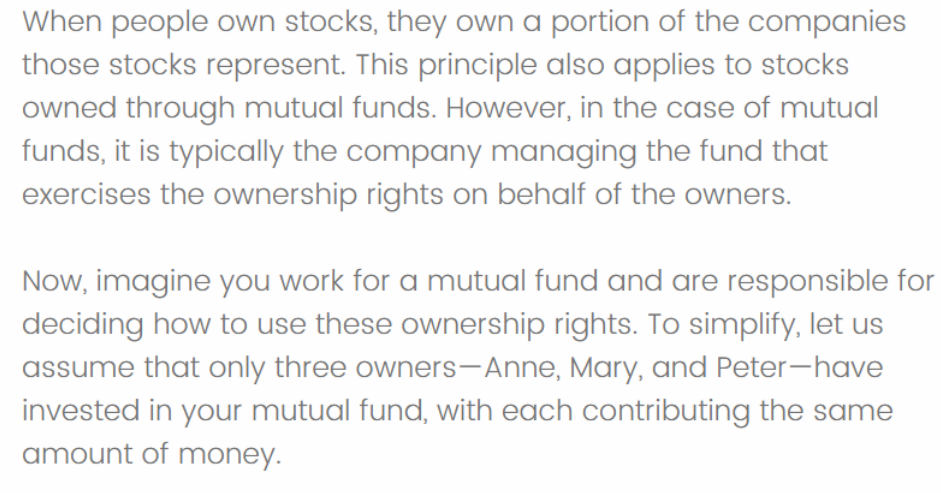


Figure A79: Introduction

The mutual fund you work for has purchased shares in a company that is evaluating whether to invest in coal production for export. The company is considering three possible approaches:

- **Option A:** Export coal as long as it is profitable.
- **Option B:** Export coal only if it contributes to geopolitical stability in Europe by reducing dependency on Russia.
- **Option C:** No coal exports under any circumstances.

The three owners who have invested equally in your mutual fund —Anne, Mary, and Peter—have different preferences regarding these options:

- **Anne** prefers **Option A** the most, followed by **Option B**, and lastly **Option C**.
- **Mary** prefers **Option B** the most, followed by **Option C**, and lastly **Option A**.
- **Peter** prefers **Option C** the most, followed by **Option A**, and lastly **Option B**.

At the upcoming shareholder meeting, the owners will vote on which approach the company should implement. Based on the preferences of the three investors, how would you cast the vote on their behalf?

- ○ **Option A:** Export coal as long as it is profitable.
- ○ **Option B:** Export coal only if it contributes to geopolitical stability by reducing dependency on Russia.
- ○ **Option C**: No coal exports under any circumstances.
- ○ I would abstain from voting.

Figure A80: Description of Scenario

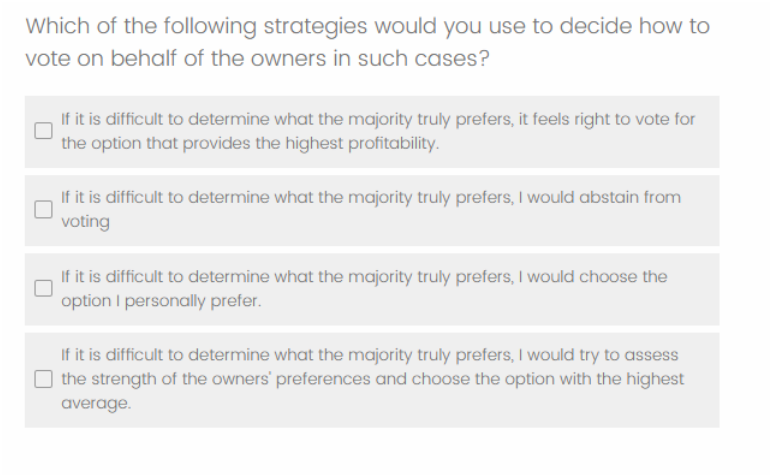


Figure A81: Approaches for Resolving Cyclical Preferences

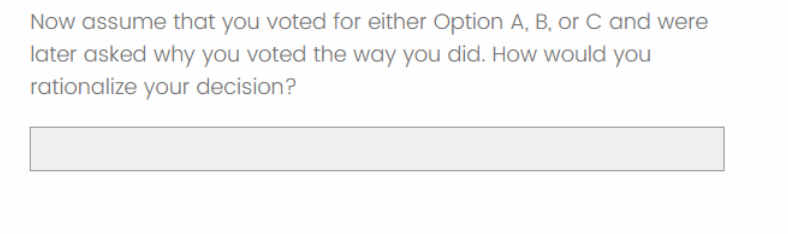


Figure A82: Open-Ended Question on Justifying the Aggregation Rule

F Pre-analysis plans

F.1 Experiment 1

ASPREDICTED

Wharton CREDIBILITY LAB
University of Pennsylvania

Algorithmic Shareholder Voting Experiment (#267097)

Author(s)
This pre-registration is currently anonymous to enable blind peer-review.
It has one author.

Pre-registered on:
2026/01/09 04:34 (PT)

1) Have any data been collected for this study already?
No, no data have been collected for this study yet.

2) What's the main question being asked or hypothesis being tested in this study?
This study tests two pre-specified questions:

H1 (Preference structure).
Do shareholder preferences over three algorithmically constructed outcomes (A: non-climate social priority, B: climate priority, C: board-aligned/profit priority) satisfy the single-crossing condition when respondents are classified into left, center, and right political groups?

H2 (Aggregation performance).
How do common ordinal voting rules (plurality, Borda, Condorcet, runoff) perform in aggregating these preferences relative to a cardinal benchmark based on willingness-to-pay (WTP)?

All scenarios used in this experiment are generated mechanically by a pre-specified algorithm applied to real shareholder ballots, with no discretionary scenario selection.

3) Describe the key dependent variable(s) specifying how they will be measured.
The primary dependent variables are:

1. Ordinal rankings of three alternatives (A, B, C) for each of five scenarios.

2. WTP measures elicited for one randomly selected scenario per respondent:

a)WTP to secure the first-ranked alternative over the third-ranked alternative.

b)WTP to secure the second-ranked alternative over the third-ranked alternative.

WTP is elicited using a multiple price list consisting of six binary choices between having the researchers vote for a higher-ranked alternative or receiving a monetary bonus of $0, $1, $2, $3, $10, or $20.

4) How many and which conditions will participants be assigned to?
All respondents rank the three alternatives in five scenarios.

To limit respondent burden, each respondent is randomly assigned to provide WTP valuations for one of the five scenarios they ranked.

The order of demographic questions (including political ideology) is randomized such that approximately half of respondents answer demographics before the ranking tasks and half after.

No other treatments are used.

5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.
Respondents are classified into three political groups:

-Left (L): "left-wing" or "strongly left-wing"

-Middle (M): all other responses

-Right (R): "right-wing" or "strongly right-wing"

For each scenario, alternatives are ordered on a left–right dimension as B–A–C, where:

-B (climate priority) is defined as the most left-wing alternative,

-C (board-aligned/profit) is defined as the most right-wing alternative,

-A lies between B and C.

F.2 Experiment 2

ASPREDICTED

Wharton CREDIBILITY LAB
University of Pennsylvania

Norway Experiment (follow up to #267097). (#276348)

Author(s)
This pre-registration is currently anonymous to enable blind peer-review.
It has one author.

Pre-registered on:
2026/03/02 04:56 (PT)

1) Have any data been collected for this study already?
No, no data have been collected for this study yet.

2) What's the main question being asked or hypothesis being tested in this study?
This data collection replicates central parts of data collection #267097.
The main hypothesis is that individual preference rankings over shareholder voting strategies exhibit structured preferences consistent with single-crossing.

3) Describe the key dependent variable(s) specifying how they will be measured.
The dependent variables are respondents' complete rankings of three voting strategies in each of five company scenarios.

For each scenario, rankings will be coded as ordinal rankings (1 = most preferred, 3 = least preferred).

4) How many and which conditions will participants be assigned to?
There are no experimental treatment conditions.

5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.
For each scenario separately and pooled across scenarios, we will:

Describe the distribution of rankings.

Test whether observed rankings satisfy single-crossing restrictions using the same procedure as in data collection #267097.

Compare aggregate patterns to those observed in #267097.

No additional hypothesis tests will be conducted.

6) Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.
No outliers will be excluded. Data collection is conducted by Ipsos using standard quality controls.

7) How many observations will be collected or what will determine sample size? No need to justify decision, but be precise about exactly how the number will be determined.
The target sample size is 3,000 completed responses.

Data collection stops when 3,000 complete responses have been obtained.

8) Anything else you would like to pre-register? (e.g., secondary analyses, variables collected for exploratory purposes, unusual analyses planned?)
No

BUNDLE
This pre-registration is part of a bundle which includes:
#276,378 - https://aspredicted.org/vi9i7w.pdf - Title: 'Strategic Voting Aggregation of Information and Preferences' #276,414 - https://aspredicted.org/ge5mh9.pdf - Title: 'Shareholder Power (Norway)'

F.3 Experiment 3

ASPREDICTED

Wharton CREDIBILITY LAB
UNIVERSITY of PENNSYLVANIA

ESGpreferences (#188348)

Author(s)
This pre-registration is currently anonymous to enable blind peer-review.
It has one author.

Pre-registered on: 08/31/2024 11:55 AM (PT)

1) Have any data been collected for this study already?
No, no data have been collected for this study yet.

2) What's the main question being asked or hypothesis being tested in this study?
The main goal is to investigate whether people's preferences are single-peaked in the ESG domain, and if there are natural ways of aggregating preferences (such as through a Condorcet winner)

Specifically, respondents are asked to rank three outcomes: prioritizing climate, prioritizing a social dimension (across 8 different social issues), and prioritizing profit.

In the second part, I will explore respondents' willingness to pay (WTP) for us to vote for their first and second-ranked choices over their third-ranked choice.

3) Describe the key dependent variable(s) specifying how they will be measured.
The key dependent variable is the individual ranking of the three outcomes.

In the second part, the focus is on measuring WTP (Willingness to Pay) using a multiple price list. The largest personal bonus a respondent is willing to forgo for us to vote for one policy over another is used to quantify their WTP.

4) How many and which conditions will participants be assigned to?
There are no treatments, but participants complete 4 hypothetical rankings followed by 4 rankings with real stakes. Each respondent is randomly assigned 2 issues for the hypothetical rankings and 2 issues for the rankings with real stakes.

5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.
The analysis is primarily exploratory, aiming to determine if respondents' preferences exhibit single-peakedness and to identify any Condorcet winners in each of the 8 cases. Subsequently, I will compare these Condorcet winners (if they exist) to the second stage involving WTP to assess how well the WTP (which seeks to measure preferences cardinally) aligns with the ranking-based approach.

6) Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.
I will not exclude any outliers that pass the attention check.

7) How many observations will be collected or what will determine sample size? No need to justify decision, but be precise about exactly how the number will be determined.
I aim to have 2,500 respondents. If there are fewer than 2,500 respondents by September 10th, 23:59 CET, I will close the data collection.

8) Anything else you would like to pre-register? (e.g., secondary analyses, variables collected for exploratory purposes, unusual analyses planned?)
Nothing else to pre-register.

F.4 Experiment 4

ASPREDICTED

Wharton CREDIBILITY LAB
UNIVERSITY of PENNSYLVANIA

Board member ranking (#208601)

Author(s)
This pre-registration is currently anonymous to enable blind peer-review.
It has one author.

Pre-registered on: 01/20/2025 03:42 AM (PT)

1) Have any data been collected for this study already?
No, no data have been collected for this study yet.

2) What's the main question being asked or hypothesis being tested in this study?
This data collection aims to evaluate whether the single-crossing property holds when aggregating shareholder preferences. Respondents will be presented with six scenarios in which they rank three distinct outcomes: a Profit outcome, a Climate outcome, and a third outcome unrelated to profit or climate. All six scenarios pertain to preferences regarding how the fund in which the respondent is investing should vote on board candidates with differing priorities.

In addition to this primary focus, the analysis will include supplementary examinations of preference aggregation methods, comparing Borda, Condorcet, plurality, and runoff voting systems. Furthermore, the study will explore the structure of preferences, specifically investigating the presence of single-peakedness.

3) Describe the key dependent variable(s) specifying how they will be measured.
The ranking of the three alternatives results in six distinct values based on their relative positions.

4) How many and which conditions will participants be assigned to?
No conditions.

5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.
I will then test whether the single-crossing property holds when respondents are classified into three political groups: left, middle, and right. To assign respondents to these groups, I will use a Likert scale question. Respondents who select "right-wing" or "strongly right-wing" will be placed in the "right" category, those selecting "left-wing" or "strongly left-wing" will be placed in the "left" category, and others will be categorized as "middle."

To analyze single-crossing, I will establish a ranking of the three alternatives from left to right. "Profit" will be designated as the most right-leaning alternative. "Climate" will be considered the most left-leaning alternative if more respondents with left political preferences rank "climate" as their first choice over the third alternative. Otherwise, "climate" will be considered the middle alternative, and the third alternative will be deemed the most left-leaning.

6) Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.
All respondents who pass the attention check will be included in the analysis, and no special treatment will be applied to outliers. Since all primary questions involve rankings, no outliers are expected.

7) How many observations will be collected or what will determine sample size? No need to justify decision, but be precise about exactly how the number will be determined.
I plan to collect responses from 1,500 US participants on Prolific before excluding those who fail the attention check. Participants will be selected from a pool of individuals who have made investments in mutual funds or ETFs. Data collection will conclude on January 27th at 23:59 CET, unless the target of 1,500 respondents is reached earlier.

8) Anything else you would like to pre-register? (e.g., secondary analyses, variables collected for exploratory purposes, unusual analyses planned?)
No

F.5 Experiment 5





BoardNorway (#210314)

Author(s)
This pre-registration is currently anonymous to enable blind peer-review.
It has one author.

Pre-registered on: 01/30/2025 07:10 AM (PT)

1) Have any data been collected for this study already?
No, no data have been collected for this study yet.

2) What's the main question being asked or hypothesis being tested in this study?
This data collection serves as a follow-up to studies #188348 and #208601. Specifically, it replicates the question from study #208601 in a Norwegian context, where respondents are indirect owners of the world's largest sovereign wealth fund.

3) Describe the key dependent variable(s) specifying how they will be measured.
As in study #208601, this data collection examines the same question but with a reduced set of ranking scenarios-three instead of six. These three scenarios are a subset of those in #208601, focusing on climate, profit, and the themes of 1) labor unions, 2) gender quotas, and 3) animal welfare.

4) How many and which conditions will participants be assigned to?
No conditions.

5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.
Same as in #208601

6) Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.
The data will be collected through YouGov Norway, which will handle all respondent communication and provide a representative sample of the Norwegian population. No outliers will be excluded or treated differently.

7) How many observations will be collected or what will determine sample size? No need to justify decision, but be precise about exactly how the number will be determined.
I aim for 1500 respondents.

8) Anything else you would like to pre-register? (e.g., secondary analyses, variables collected for exploratory purposes, unusual analyses planned?)
No

F.6 Experiment 6

ASPREDICTED

Wharton CREDIBILITY LAB
University of Pennsylvania

Experiment with four options (#275061)

Author(s)
This pre-registration is currently anonymous to enable blind peer-review.
It has one author.

Pre-registered on:
2026/02/23 06:17 (PT)

1) Have any data been collected for this study already?
No, no data have been collected for this study yet.

2) What's the main question being asked or hypothesis being tested in this study?
This study is a follow-up to data collection #267097. Participants are assigned to the same five scenarios as in that study. The only change is that each scenario now includes four alternatives to be ranked rather than three. The additional alternative allows participants to vote in favor of both the climate proposal and the non-climate social proposal simultaneously.

Unlike #267097, this data collection does not include a willingness-to-pay (WTP) elicitation task following the scenarios.

3) Describe the key dependent variable(s) specifying how they will be measured.
The key dependent variables are the same as in #267097. Specifically, the primary outcome is the ranking assigned to each alternative within each scenario.

4) How many and which conditions will participants be assigned to?
There are no substantive treatment conditions. Participants are randomly assigned to one of two survey-order conditions:
-Demographic questions before the main scenarios
-Demographic questions after the main scenarios

5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.
The analyses will follow the same procedures specified in #267097.

6) Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.
All respondents who pass the predefined strict attention and bot checks will be included in the analysis.

No additional outlier exclusions will be applied. Because the primary outcome is a ranking variable, extreme-value outliers are not applicable.

7) How many observations will be collected or what will determine sample size? No need to justify decision, but be precise about exactly how the number will be determined.
The target sample size is 1,000 respondents. Participants will be recruited from Prolific and restricted to U.S.-based individuals who self-report being equity market investors.

8) Anything else you would like to pre-register? (e.g., secondary analyses, variables collected for exploratory purposes, unusual analyses planned?)
No

F.7 Experiment 7

ASPREDICTED

Wharton CREDIBILITY LAB
University of Pennsylvania

Shareholder preferences January 2026 (#269073)

Author(s)
This pre-registration is currently anonymous to enable blind peer-review.
It has one author.

Pre-registered on:
2026/01/21 04:16 (PT)

1) Have any data been collected for this study already?
No, no data have been collected for this study yet.

2) What's the main question being asked or hypothesis being tested in this study?
This study examines how institutional design affects shareholder voting. The preregistered analyses address three questions:

-How making it easier to use voting choice (delegation to predefined voting profiles) affects voting outcomes relative to costly issue-by-issue voting, and how this effect can be decomposed into participation (turnout/selection) and delegation (coarsening) components.

-How expanding the choice set from three to six policy alternatives affects preference structure and aggregate voting outcomes.

-Whether voting in a contested board-election setting concentrates at binary extremes or instead exhibits interior, mixed choices, and whether preferences over board bundles remain sufficiently structured to permit aggregation.

3) Describe the key dependent variable(s) specifying how they will be measured.
The key dependent variables include:

-Binary FOR/AGAINST votes on individual shareholder proposals and implied votes based on delegation to voting profiles.
-Ordinal rankings of policy objectives, policy alternatives, and board proposals.
-Preference-intensity measures (measured on scale 0-10 and by allocating points across alternatives)

4) How many and which conditions will participants be assigned to?
Participants are randomly assigned to one of four between-subject conditions:

-Voting Choice (delegation to voting profiles).
-Costly issue-by-issue voting with low participation cost.
-Costly issue-by-issue voting with high participation cost.
-Forced issue-by-issue voting (full-participation benchmark).

Furthermore, the presentation order is randomized across tasks: the three- versus six-alternative rankings and the corresponding point-allocation questions are counterbalanced across participants.

5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.
The first part of the analysis examines how participation costs and voting choice affect voting outcomes, and decomposes differences between Voting Choice and costly issue-by-issue voting into participation/selection and delegation (coarsening) effects. All treatment effects will be estimated using ordinary least squares (OLS) regressions with robust standard errors.

In addition, I examine how expanding the choice set from three to six alternatives affects aggregation, including how different aggregation rules perform as the choice set expands, and analyze voting behavior and preference structure in a contested board-election setting, including whether choices exhibit regularities consistent with structured preferences

6) Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.
All respondents who pass attention checks and bot-detection checks are included in the analysis.

Bot-detection checks include reCAPTCHA, an audio-based task, and a short writing task designed to identify automated responses.

No trimming or winsorization of responses is performed. The main outcome variables are rankings, binary variables, or variables that are naturally bounded.

7) How many observations will be collected or what will determine sample size? No need to justify decision, but be precise about exactly how the number will be determined.
I will collect 2,000 completed observations.

8) Anything else you would like to pre-register? (e.g., secondary analyses, variables collected for exploratory purposes, unusual analyses planned?)
Nothing else to pre-register.

F.8 Experiment 9

ASPREDICTED

Wharton CREDIBILITY LAB University of Pennsylvania

Norway Aggregation Rule (#276586)

Author(s)
This pre-registration is currently anonymous to enable blind peer-review.
It has one author.

Pre-registered on:
2026/03/03 03:50 (PT)

1) Have any data been collected for this study already?
No, no data have been collected for this study yet.

2) What's the main question being asked or hypothesis being tested in this study?
This study is a follow-up data collection intended to replicate data collection #206840 in a Norwegian setting.

The main objective is to examine how respondents choose to vote when acting as a fund manager on behalf of investors with cyclical preference profiles, and which decision principles they report using in such situations.

3) Describe the key dependent variable(s) specifying how they will be measured.
The primary dependent variables are:

The selected voting alternative (A, B, C, or abstain).

The selected decision principle.

4) How many and which conditions will participants be assigned to?
There are no experimental conditions. All respondents receive the same scenario and questions.

5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.
I will report simple descriptive statistics, such as frequency distributions for the chosen voting alternative and for the selected decision principle.

6) Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.
No outliers will be excluded.

7) How many observations will be collected or what will determine sample size? No need to justify decision, but be precise about exactly how the number will be determined.
The target sample size is 500 respondents, collected by Norstat Norway. Data collection stops once 500 completed responses have been obtained.

8) Anything else you would like to pre-register? (e.g., secondary analyses, variables collected for exploratory purposes, unusual analyses planned?)
Nothing else to pre-register.

F.9 Experiment 10

ASPREDICTED

Wharton CREDIBILITY LAB
University of Pennsylvania

Voting under cyclical preferences (#206840)

Author(s)
This pre-registration is currently anonymous to enable blind peer-review.
It has one author.

Pre-registered on: 01/07/2025 12:24 PM (PT)

1) Have any data been collected for this study already?
No, no data have been collected for this study yet.

2) What's the main question being asked or hypothesis being tested in this study?
This study examines how individuals approach the aggregation of shareholder preferences in situations where cyclical preferences arise, resulting in the absence of a Condorcet winner.

3) Describe the key dependent variable(s) specifying how they will be measured.
The first dependent variable is a voting variable with four possible outcomes: voting for Option A, Option B, Option C, or abstaining. After casting their vote, respondents are asked to justify their choice by indicating their level of agreement with four statements and providing a response to an open-ended question.

4) How many and which conditions will participants be assigned to?
No treatments.

5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.
The analysis is exploratory. An open-ended question is included to capture how individuals justify their mapping of preferences into voting decisions. Responses to this question will be analyzed using large language models (LLMs) to classify respondents into distinct categories.

6) Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.
No outliers will be excluded from the analysis. All responses that pass the attention check will be included.

7) How many observations will be collected or what will determine sample size? No need to justify decision, but be precise about exactly how the number will be determined.
The sample size is set at 1,000 participants, drawn from Prolific. Participants who do not pass the attention check will be excluded. The sample is designed to be representative of the U.S. population in terms of age, gender, and political affiliation.

8) Anything else you would like to pre-register? (e.g., secondary analyses, variables collected for exploratory purposes, unusual analyses planned?)
No