Voting for Socially Responsible Corporate Policies

Adam Meirowitz
Yale University

Shaoting Pi
Iowa State University

Matthew C. Ringgenberg
University of Utah

First Draft: November 2021
This Draft: February 2023*

ABSTRACT

We examine voting for corporate policies when voters face a trade-off between maximizing firm value and implementing social objectives (like minimizing pollution). We show voting can aggregate preferences and lead to stable policies when a firm cares about one social objective in addition to firm value. But when voters care about two or more social objectives it can lead to governance indeterminacy or volatility in firm policies. We show an agenda setter (e.g., activist investor or CEO) can influence policy choice, but this power is not absolute. Our findings have implications for the design and regulation of proxy voting rules.

Keywords: Corporate governance, ESG, Proxy Voting, Shareholder Voting, Social Choice Theory, Socially Responsible Investing, Sustainability

JEL Classification Numbers: G12, G14

*Adam Meirowitz (corresponding author, adam.meirowitz@yale.edu) is at Yale University. Matthew C. Ringgenberg (matthew.ringgenberg@eccles.utah.edu) is at the University of Utah. Shaoting Pi (shaoting@iastate.edu) is Iowa State University. We thank Hui Chen, Shaun Davies, Kornelia Fabisik, Chong Shu, Alexander Wagner, and seminar participants at Iowa State University and the University of Zurich. All errors are our own.

I. Introduction

Voting plays an important role in corporate governance. Shareholders vote to elect members of the board of directors and they vote on proposals that may directly affect the actions of the firm. Similarly, members of the board of directors vote to appoint the chief executive officer and vote on a variety of firm policies. Traditionally, these stakeholders broadly agreed on the objective of the firm – maximize firm value – and this one-dimensional objective often simplifies voting.[1] More recently, a number of academics, practitioners, and regulators have argued that firms ought to care about more than just value, yet there has been little conceptual work evaluating whether this shift in the scope of what matters to investors, shareholders, and board members impacts the quality of firm governance.[2]

Since at least Fisher (1930), the predominant view in corporate finance states that firms should maximize firm value, without regard to other dimensions (i.e., the shareholder view). Indeed, Jensen (2001) argues it is, "logically impossible to maximize in more than one dimension." As a result, he argues managers should focus *solely* on maximizing firm value. We show this view is incomplete. In the traditional view, there is no natural way for corporate managers to specify a trade-off between competing objectives – as a result, there is no natural way to maximize an objective function that depends on multiple dimensions.[3] How-

[1] DeMarzo (1993) shows an exception: if markets are incomplete, then investors may also disagree on how to maximize firm value, which complicates the public choice problem.

[2] A large literature examines the challenges of corporate governance when the goal is to maximize firm value. For example, there is extensive work on the agency conflict that arises between investors and managers when ownership and management are separate. See, for example, Berle and Means (1932); M. Jensen and Meckling (1976); Demsetz (1983); Admati, Pfleiderer, and Zechner (1994); Burkart, Gromb, and Panunzi (1997); Maug (1998).

[3] Jensen (2001) writes, "...it is not logically possible to speak of maximizing both market share and profits...there is no purposeful way to decide where to be in the area where the firm can obtain more of one only by giving up some of the other." Similarly, in defense of the shareholder view Denis (2016) writes, "Once we leave market prices behind, what alternative decision rules should be used to determine which interested parties receive above-market rewards, and at which other parties' expense? There will be at least as many different opinions about this as there are types of interested parties. Whose opinion will prevail?"

ever, we follow the perspective of social choice theorists and political scientists and frame the question from the perspective of a collection of individual investors who each have well-defined preferences over multiple dimensions. We then take their purposeful voting behavior (which reflects their preferences) as an input into the process of corporate governance instead of using an approach that seeks to combine multiple utility functions into one social welfare objective function. Specifically, our focus is on the conditions under which individual preferences can be naturally aggregated via voting. Viewed through this lens, we draw on particular features of this problem domain and show it is possible to maximize in more than one dimension. Specifically, we build a theory of corporate voting over policies that impact firm value as well as other dimensions (i.e., the stakeholder view) which may include preferences for environmental, social, and governance policies ("ESG"). In our analysis, we flesh out how the movement from concern over value to concern over value plus ESG impacts governance.

Our approach builds on the literature on social choice theory and political economy. Arrow (1951) famously shows that no method of aggregating preferences will satisfy a small set of naturally satisfying axioms. However, subsequent literature shows that stable choices can emerge under various restrictions on voter preferences. A well-studied restriction is the case of single-peaked preferences which is satisfied if the feasible policies can be arranged in a single dimension and on this dimension all agents' preferences are quasi-concave (informally, monotone, or tent-shaped). When this restriction is satisfied, it is possible to identify choices that reflect the will of a majority or aggregate preferences. Moreover, it has been shown (e.g., Banks and Duggan (2000)) that fairly natural models of bargaining make tight predictions about which policies will emerge under strategic proposing and voting when this restriction on preferences is satisfied. Building on this literature, we examine voting for corporate policies when voters face a trade-off between maximizing firm value and one or more social policies (for example, reducing pollution and increasing employee satisfaction). In contrast to the prevailing view that it is impossible for firms to maximize a multi-dimensional objective function (e.g., Jensen (2001), Denis (2016), etc.), we show that adding one social dimension does not lead to additional problems relative to the pure shareholder view of the firm. While, in general, voting does not tend work well with even two dimensions, because there are natural trade-offs between maximizing firm value and incorporating social objectives we show the problem is essentially one dimension smaller. However, we also find that a number of challenges arise when the firm objective function is expanded to incorporate more than one social dimension.

A simple example helps to illustrate our key findings and also points to a key limitation in the treatment of voting over broader objectives as in Hart and Zingales (2017). Imagine a firm with three possible policy choices. The firm can implement a policy, P, that maximizes firm value or a policy G that is less profitable but environmentally sustainable or a policy E that is less profitable still but mandates ethical treatment of workers. Consider three investors (or three board members) who are charged with making the policy choice, denoted as investors 1, 2, and 3. Suppose that investor 1 cares only about firm value and thus orders the alternatives by expected firm value: P, G, then E. Suppose that investor 2 most prefers to protect the environment, but would still rather support the ethical treatment of workers over just maximizing firm value and so ranks the alternatives, G, E, then P. Finally, suppose investor 3 cares about the ethical treatment of workers but is not willing to sacrifice returns for environmental policies and thus ranks the alternatives E, P, and then G. If any *two* of these policies are offered, a stable choice will emerge. In particular: given a choice between E and P, E wins. Given a choice between P and G, P wins. And given a choice between

G and E, E wins.[4] However, if all three policies are offered, none of the alternatives beats the other two alternatives. While E beats P, G beats E, yet P beats G. As a result, when the three policies are offered, none of them is naturally preferred or stable under majority rule. The ultimate policy choice may thus depend on additional and less obvious features of the institution. This creates additional challenges for an investor who might face serious uncertainty about the firm's likely policy choice.

What are the implications of this finding? First, contrary to popular belief, it is possible to identify natural policy choices using voting if there is more than one dimension of interest. But, if the shareholders care about more than two dimensions, there is generally no natural policy choice. As a consequence, the choices that emerge will depend on the process by which policies are proposed, and the volatility of firm choices will tend to increase. This simple example illustrates a deep and practical concern about preference aggregation. By moving away from the assumption that there are just two alternatives (e.g., Hart and Zingales (2017)) and in particular, by constructing preferences over three alternatives we see the possibility for there to be no natural choice.

Although majority rule (and other stronger super-majority rules) are not immune from Arrow's impossibility theorem, there are still compelling reasons to use them. In particular, when preferences are single-peaked (or more general satisfy order-restriction) majority rule is known to be well-behaved and many systems that involve fairly decentralized proposal rights and majority voting will tend to select policies that are quite responsive to the preferences of the so-called median voter. But when preferences do not satisfy these types of restrictions, the outcomes can depend heavily on seemingly subtle institutional features. Whether a policy-making domain exhibits enough preference diversity for this problem to

[4] Furthermore, assuming that the utility scale across investors 1, 2, 3 is comparable then selecting the majority winning alternative from a binary comparison is justified on utility maximizing grounds.

become important is an applied question, that to date, has not been extensively studied in finance contexts (such as choosing socially responsible corporate policies). Our paper fills this void.

We also analyze how the process of submitting a voting proposal affects voting results. In reality, both insiders of a firm, such as the CEO, and outsiders of a firm, such as a hedge fund investor, can propose a policy at shareholder meetings. Inside a firm, the power of proposing is typically concentrated. For example, a firm's corporate charter may entitle the proposal right to a specific person, such as the chairperson of the board. On the other hand, outside a firm, the power of proposing is usually diffused. Under Securities and Exchange Commission (SEC) rules, any shareholder who has held $2000 of stock for at least three years (or higher amounts for shorter periods) is eligible to submit a shareholder proposal. To capture these two possible situations, we consider two kinds of agenda-setting environments, monopoly-agenda control and diffused-agenda control. With the monopoly agenda setting, we consider that a privileged agent (e.g., the CEO of a firm or the chairperson of a firm's board) has monopoly control of the agenda. She can select a policy and that will be pitted against the status quo. Then, all shareholders vote to decide whether to implement the proposed policy. In the diffused-agenda setting case, agenda-setting power is far less concentrated. We assume that a firm is owned by several hedge funds, and each of them has a chance to propose a new policy. The process is dynamic. By considering the process of setting a voting agenda, we establish two additional conclusions: (1) the ability to control the voting agenda is important and allows the proposer to strongly influence policy choice, but (2) this power is not absolute, as the majority's will operates as a constraint on what the proposer can achieve.

Our findings have important practical implications. Investors and board members are increasingly asked to vote on policies that cover a number of dimensions. In 1999, there were approximately 200 shareholder proposals on environmental and social issues; by 2018, that number had grown to nearly 500 proposals (Papadopoulos, 2019). In 2020 alone, Russell 3000 stocks had at least 288 shareholder proposals on governance issues, 174 on environmental and social issues, 76 on civic engagement issues, and 56 on executive compensation issues.[5] Yet despite the increasing prevalence of such diverse shareholder proposals, little is known about the impact of expanding the domain of governance to a multi-dimensional choice space. Moreover, increasingly academics and regulators have scrutinized the voting power of the Big Three index fund providers (BlackRock, State Street, and Vanguard).[6] In response, BlackRock and Vanguard have recently announced plans to cede voting authority to their individual investors (BlackRock (2022), Vanguard (2023)). Viewed in light of our findings, this may increase the heterogeneity of investor preferences and may increase the salience of these other dimensions, possibly leading to less stable policies.

The rest of this paper proceeds as follows. Section II explains our contribution relative to the existing literature. Section III provides the setup for our model and discusses key assumptions. In Section IV, we present our main findings regarding the possibility of identifying natural or stable policies from just preferences, the approach of social choice theory. We examine policy stability and the conditions under which a selected policy may not reflect the will of the majority as investors care about more and more issues. We then build off of these results to include a role for other institutional features to determine what policies are chosen. In Section V, we examine how the power to set an agenda and propose policies affects policy choice. In VI, we explore the entry and exit channels. Finally, Section VII

[5]See Gibson Dunn (2020). The sample includes all Russell 3000 companies in the Institutional Shareholder Services shareholder proposals and voting analytics database in the 2020 proxy season.

[6]See, for example, Azar, Tecu, and Schmalz (2018), Anton, Ederer, Gine, and Schmalz (2023).

concludes.

II. Background

Our paper makes contributions to several strands of literature. First, we contribute to the literature on corporate governance and agency conflicts. The theoretical literature on corporate governance argues that investors can affect firm policies through two main channels: voice and exit (Edmans (2014)). That is, investors can vote to change firm policies, or they can sell their shares and exit the firm. Several papers argue that exiting as an investor can discipline managerial behavior (e.g., Edmans (2009); Dasgupta and Piacentino (2015); Levit (2019)). Alternatively, a large literature examines how investors vote (e.g., N. Malenko and Shen (2016), Levit, Malenko, and Maug (2019), A. Malenko and Malenko (2019), Brav, Jiang, and Li (2018), Bolton, Li, Ravina, and Rosenthal (2020), Bubb and Catan (2018)). In particular, McCahery, Sautner, and Starks (2016) show that voting against policies recommended by corporate managers is an important governance mechanism. However, to date, the literature has not examined the complications that arise in social choice problems when voters care about multi-dimensional goals, as we do in this paper. DeMarzo (1993) is one of the few papers to examine corporate voting in the presence of diverse investor preferences. In his setting, markets are incomplete, and as a consequence, shareholders disagree over how to achieve the goal of maximizing firm value. He shows that in a simple majority rule voting system, the firm will make production decisions that are optimal from the perspective of the largest shareholder. He then shows that when a board of directors controls agenda setting, shareholders will no longer be able to affect firm policies.

Our paper also contributes to the growing literature on socially responsible investing.

Friedman (1970) helped establish the idea that firms should have a single objective function – they should maximize firm value. In contrast, Hart and Zingales (2017) argue that firms should work to maximize investor utility, instead of investor profits.[7] While this debate remains unsettled, it is increasingly clear that investors are attracted to the idea of socially responsible investing. For example, Hartzmark and Sussman (2019) show that investors flow into funds that receive a high sustainability rating from Morningstar. Similarly, Gantchev, Giannetti, and Li (2020) show that mutual funds alter their holdings to improve their Morningstar sustainability ratings. What is less clear is whether these funds are able to successfully alter firm policies. Oehmke and Opp (2020) examine the conditions under which socially responsible investing could theoretically impact firm behavior. They argue that socially responsible investors can affect firm behavior via a financing channel. However, Berk and van Binsbergen (2021) argue that, in their current state, socially responsible funds are not large enough to significantly impact most firms' cost of capital by exiting (or entering) a stock. As such, their results suggest socially responsible investing is unlikely to impact firm behavior through the exit mechanism. Empirically, Dikolli, Frank, Guo, and Lynch (2021) examine voting by socially responsible investors and find they are more likely than other investors to vote in favor of socially responsible policies. However, Michaely, Ordonez-Calafi, and Rubio (2021) show this finding is not stable: they find that the voting behavior of socially responsible investors varies, depending on the context. Our paper provides a possible explanation for this finding: when investors care about more than two dimensions, it is likely that policy choices will become unstable and depend on the manner in which they are proposed. Finally, Heath, Macciocchi, Michaely, and Ringgenberg (2021) examine the relation between socially responsible investing and real-world measures of en-

[7]It is important to note that because they focus on a binary choice problem, they side-step the possibility of a non-trivial social choice problem, as we study in this paper.

vironmental and social policy outcomes, such as pollution and employee satisfaction; they find that socially responsible investing does not improve firm behavior.

Our paper offers an application of social choice theory and public choice more generally. A theoretical and applied literature focuses on settings in which the underlying choice space is multi-dimensional and agent preferences are spatial. By this, we mean preferences admit a utility representation in which each agent has a finite ideal policy and utility tends to fall off as one moves away from this ideal. Plott (1967) and Saari (1997) characterize conditions under which majority rule admits a stable policy and show how increasing the dimensionality of the policy space tends to make stable policies less likely. Applied work in politics tends to focus on low-dimensional spatial models. Meirowitz (2004) shows how to relate the case of monotone preferences and a feasibility constraint to the earlier dimensionality results for spatial preferences. Eraslan (2016a) studies a multilateral barging model in which different risk-averse players may be chosen as the proposer with different probability or have different discount rates. Here we apply these abstract tools to the domain of corporate governance with social objectives focusing on a case with monotone preferences over a multiple-dimensional space and use the ability to translate this into a problem with spatial preferences over a lower-dimensional space.

III. Model Setup

A. Modeling Considerations

We develop a framework for thinking about the collective choice by a group of investors, shareholders, or a board. Our goal is to parsimoniously explore how the functioning of preference aggregation depends on the number of issues the shareholders consider. We first follow the social choice theoretic approach examining when policy choices can be justified by examining only the underlying preferences of the decision-makers. We will focus primarily on the aggregation of preferences by majority rule and ask whether the existence of natural or stable policies under majority voting depends on the scope of the decision problem. In various contexts, it has been shown that when such stable policies exist, they tend to emerge from a broad class of models. When such stable policies do not exist, it is not necessarily the case that policymaking will be erratic or unstable (although that is possible). But the outcome of these games tends to be more sensitive to other assumptions or features of the game. We build on the social choice theoretic results by then analyzing several canonical classes of agenda-setting models.[8]

B. Primitives

We consider a group of $n \geq 3$ decision-makers charged with voting on a corporate policy.[9] These decision-makers could be investors, like hedge funds, pension plans, or members of the board of directors. As our focus is on governance in the presence of trade-offs between profitability and other outcomes impacted by corporate behavior, we conceive of the choice space as multi-dimensional. Thus, a corporate policy is a vector $x = (x_1, ..., x_k)$ in R^k_+. For convenience, we refer to this choice set as X.

We are interested in dimensions like shareholder value as well as socially desirable attributes like fair treatment of workers or environmentally sound production processes. As such, it is natural to assume that all decision-makers agree that each dimension is good

[8]We strive for a self-contained exposition. For much more depth and historical context, we direct the interested reader to two volumes: Austen-Smith and Banks (2000a) for the social choice theory and Austen-Smith and Banks (2000b) for the institutionally richer game-theoretic approach.

[9]It is convenient to assume that n is odd. When the population is odd, a tie is not possible under majority rule. If n is an even number, the analysis is less tractable, but it does not generally provide greater intuition. Where appropriate we point out the relevant changes.

but decision-makers differ in how much they value each of these goods. Continuing with our example, some investors care more about stock returns, while other investors care more about the treatment of workers. For convenience, we work with the parsimonious and well-studied Cobb-Douglas family of utility functions.[10] The utility to individual i from policy x is given by $u_i(x) = \Pi_{j=1}^{k} x_j^{\alpha_{i,j}}$, where the strictly positive vector $\alpha_i = (\alpha_{i,1}, ...\alpha_{i,k})$ captures the importance of each dimension to individual i. Thus, the heterogeneity of decision-makers' preferences is captured by differences in their weight vectors. We assume that no two voters have the same utility parameters, $\alpha_i \neq \alpha_j$.[11]

To capture the fact that firms typically face tradeoffs in terms of these goods, we assume that the firm must select x from a constraint set $C \subset X$. It suffices to assume this set is compact and convex. For tractability, we will parameterize it as falling below a hyper-plane. In particular, for a strictly positive vector $\beta = (\beta_1, ..., \beta_k)$, we define the feasibility set as $C = \{x : \Sigma_{j=1}^{k} \beta_j x_j \leq 1\}$.

In the next subsection, we examine the problem of selecting $x \in C$ in a manner that somehow reflects the preferences of the group. One can proceed with various aggregation rules. Over the years, several arguments have been advanced in favor of majority rule. May (1952) is compelling when there are just two alternatives/candidates. In practice, majority rule or stronger supermajority rules are widely employed by firms in making decisions in a variety of settings. We assume that investors, or members of the board of directors, select x under simple majority rule. Subsequently, we impose additional institutional details and model collective choice as an extensive form game in which an agenda is endogenously

[10]The main results extend to any strictly monotone and strictly quasi-concave preferences. In fact, with only superficial changes, the exact arguments provided here can be extended to any strictly monotone and strictly quasi-concave preferences that admit a twice-continuously-differentiable utility representation.

[11]In the sequel, we discuss how the key results would be slightly modified to allow for individuals to have the same preferences, $\alpha_i = \alpha_j$ and we consider large blocks of voters controlled by one decision-maker in an extension.

selected, and any votes that occur are via majority rule.

IV. The Will of the Majority

To determine if there is a policy x that reflects the will of the majority in this minimalist institutional perspective, we need to determine whether there is a point that beats all others under majority rule. Such a point is called a majority rule core point. A majority rule core point, if one exists, is then stable under majority rule. We first introduce the idea of the majority rule preference ordering.

Definition 1 (Majority Rule Preference Ordering, xR_my) *For any two polices x and y, if at least $\frac{n+1}{2}$ individuals have $u_i(x) \geq u_i(y)$, then we say that policy x is weakly preferred to y by a majority, which is denoted xR_my.*

Then, we define the majority rule core.

Definition 2 (Majority Rule Core, $M(C)$) *The majority rule core $M(C)$ is defined as $M(C) = \{x \in C : \forall y \in C, xR_my\}$.*

When the majority rule core is small (or even a singleton), then we have a notion of policy choices that reflect the will of the majority. But when the core is empty, no policies can be justified on the grounds that they reflect the will of the majority. On the basis of Definition 1 and Definition 2, we can naturally define a majority core point.

Definition 3 (Majority Rule Core Point, $x \in M(C)$) *A policy x is a majority rule core point if $x \in M(C)$.*

We care about finding majority rule core points, because a majority rule core point is stable; the point cannot be beaten by any other point in C under pairwise majority rule voting. When a majority rule core point exists, it is known that a fairly large set of explicit games will lead to the selection of this point or one close to it.

It is not difficult to see that any point in C that is not on the boundary (a point is not on the boundary of C if it satisfies $\Sigma_{j=1}^{k}\beta_j x_j < 1$) is less attractive than some points on the boundary of C to every voter. That is to say, everything in the interior of C is Pareto dominated by some points on the boundary of C. Formally speaking,

Lemma 1 *$M(C) \subset \hat{C}$, where $\hat{C} = \{x : \Sigma_{j=1}^{k}\beta_j x_j = 1\}$*

This lemma helps us limit our search for majority rule points to the set $\hat{C} = \{x : \Sigma_{j=1}^{k}\beta_j x_j = 1\}$. As we will see later, this fact generally reduces the dimensionality of the problem by 1.

A. Voting over Just Two Dimensions

We begin by focusing on choice problems when there are two dimensions, $k = 2$. A natural example is the case of a firm that faces a trade-off between maximizing firm value and reducing pollution. In this example, we can conceive of dimension 1 as firm value and dimension 2 as reducing pollution. While all investors want higher firm value and lower pollution, they disagree on the optimal way to structure their trade-off between the two dimensions. Accordingly, the vector $\alpha_i = (\alpha_{i,1}, \alpha_{i,2})$ captures the marginal value of each dimension to investor i. Without the loss of generality, we can assume that $\alpha_{i,1} + \alpha_{i,2} = 1$ and $\alpha_{i,j} \in (0, 1)$ for all i, j. For convenience, we label the players so that $\alpha_{i,1} \leq \alpha_{i',1}$ if $i < i'$.

A.1. Example: Two Dimensions and Three Voters

We first use an example with three voters to illustrate the main points. Suppose we have three voters, and they are different in their preferences over the two dimensions, maximizing firm value and reducing pollution. The grey triangle in Figure 1 represents the constraint of all feasible policies. The utility curves in different colors represent different voters' preferences. The three tangency points between their utility curves and the frontier of the constraint, x_1^*, x_2^*, and x_3^*, represent each voter's optimal policy.

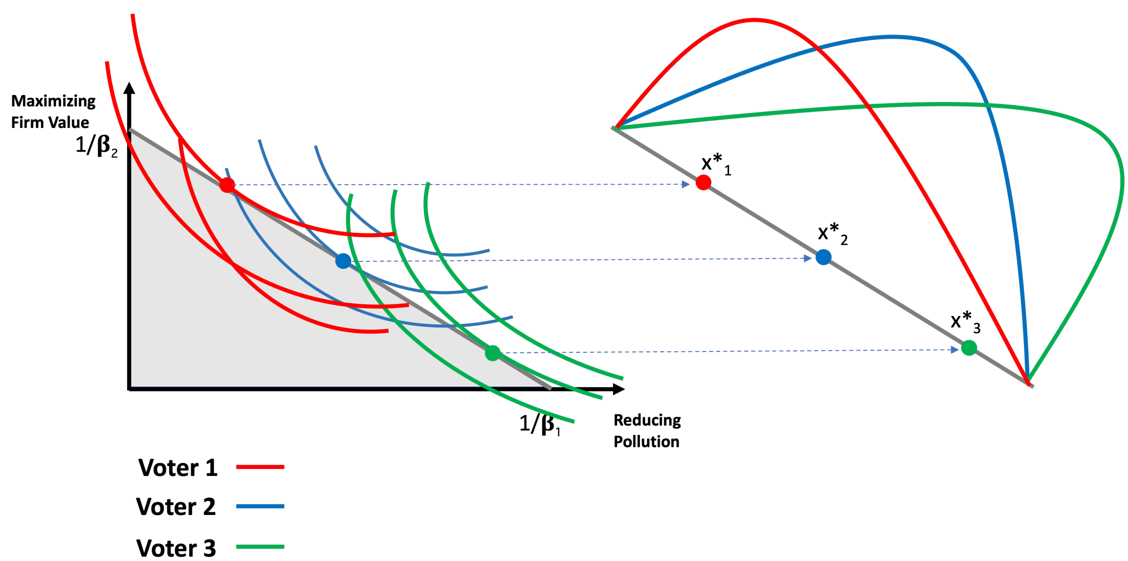

Figure 1. Two Dimensions and Three Voters

According to Lemma 1, we can simply focus on $\hat{C}$ to study the majority core rule core. As can be seen from the projection, any alternative policy on the left of x_2^* makes voter 2 and voter 3 less happy, while any alternative policy on the right of x_2^* is less satisfactory than x_2^* to voter 2 and voter 1. Therefore, the median voter's favorite policy x_2^* is the unique majority rule core point.

A.2. Two Dimensions and n Voters

Now we consider the general case that there are n voters. Recall from Lemma 1 that we can restrict consideration to the frontier, $\hat{C}$. With just two policy dimensions, $\hat{C}$ is a line segment (1-dimensional). In particular, $\hat{C} = \{x : \beta_1 x_1 + \beta_2 x_2 = 1\}$. It is easy to characterize each person's optimal point on $\hat{C}$ by solving the constrained optimization problem.

$$\max_{x_1,x_2} \; u_i(x) = x_1^{\alpha_{i,1}} x_2^{\alpha_{i,2}}$$

$$s.t. \quad \beta_1 x_1 + \beta_2 x_2 = 1$$

$$x_1, x_2 \geq 0$$

So, the optimal policy for player i is the vector

$$x_i^* = (\frac{\alpha_{i,1}}{\beta_1}, \frac{1-\alpha_{i,1}}{\beta_2})$$

Obviously, a movement away from i's optimal point x_i^* makes i less happy.

In the case of two dimensions ($k = 2$), we can directly show that the point $x^*_{\frac{n+1}{2}}$ is stable. More precisely, $x^*_{\frac{n+1}{2}}$ is the unique majority rule core point.

Proposition 1 *(Median Voter Theorem) In the 2 dimensional problem, the majority rule core is the favorite feasible policy of the decision maker with the median preference parameter,* $M(C) = x^*_{\frac{n+1}{2}} = (\frac{\alpha_{\frac{n+1}{2},1}}{\beta_1}, \frac{1-\alpha_{\frac{n+1}{2}1}}{\beta_2})$. *We call this voter the median voter.*

The proof appears in the appendix.

Note that if n were even, the core would consist of all policies on $\hat{C}$ that lie between the induced ideal of voter $\frac{n}{2}$ and voter $\frac{n}{2} + 1$.

In sum, contrary to commonly held views (e.g., Jensen (2001), Denis (2016)), when voters care about maximizing firm value and one additional dimension (such as minimizing pollution), corporate voting *can* successfully aggregate preferences and lead to stable policy choices. Even though we are starting with a two-dimensional problem, the dimension reduction afforded by the assumption that each dimension is good and the constraint set allows us to obtain a version of the well-known median voter theorem from 1-dimensional problems. Put differently, because decision-makers prefer more to less of each dimension, and each dimension is costly, we show the two-dimension firm objective can be reduced to a one-dimensional objective leading to stable policy choices. For example, if governance is over profitability and one aspect of socially responsible behavior, there is a natural notion of choice or policy selection. In the next subsection, we examine how the addition of more dimensions affects the problem.

B. Voting over More than Two Dimensions

We now examine whether the qualitative finding that the favorite policy of one of the decision-makers is stable under majority rule extends to the case where there are more than 2 dimensions (for example, maximizing firm value, minimizing pollution, and maximizing employee satisfaction). The answer is a resounding *no*.

B.1. Example: Three Voters and Three Dimensions

We first study the sample situation in which three decision-makers vote to select a policy from a choice set X and each policy has three dimensions, x_1, x_2, and x_3. For convenience, we assume that $\beta_1 = \beta_2 = \beta_3 = \frac{1}{2}$. So, the feasible set is $C = \{x : \Sigma_{j=1}^{3} \frac{1}{2} x_i \leq 1\}$. Decision-maker 1's utility function is $U_1 = x_1^{0.34} x_2^{0.33} x_3^{0.33}$. Decision-maker 2's utility function is $U_2 = x_1^{0.33} x_2^{0.43} x_3^{0.24}$. Decision-maker 3's utility function is $U_3 = x_1^{0.35} x_2^{0.25} x_3^{0.4}$.

A special case may help understand these settings better. For instance, we can consider the first dimension x_1 as firm value, the second dimension x_2 as reducing pollution, and the third dimension x_3 as employee satisfaction. All three decision-makers agree that these three dimensions are valuable but differ in how valuable each dimension is compared to other dimensions. In particular, decision-maker 1 thinks firm value is more important than reducing pollution and employee satisfaction, which are equally important to her. Decision-maker 2 thinks reducing pollution is the most important, and firm value is more important than employee satisfaction. Decision-maker 3 believes that employee satisfaction is more important than firm value, which is more important than reducing pollution. A policy that receives two or three votes from the three decision-makers will be selected.

By Lemma 1, it is sufficient to focus on the boundary of the feasible set C and consider the induced or restricted preferences of the players on this triangle. Figure 2 exhibits the constraint surface $\hat{C} = \{x : \Sigma_{j=1}^{3} \frac{1}{2} x_i = 1\}$ in a yellow triangle and one indifference curve for each player.

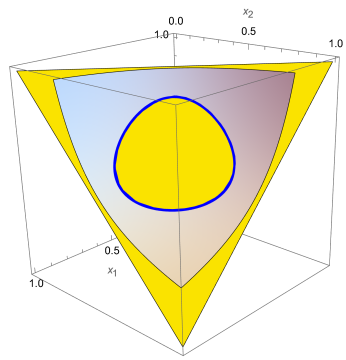

(a) $U_1 = x_1^{0.34}x_2^{0.33}x_3^{0.33} \equiv 0.65$

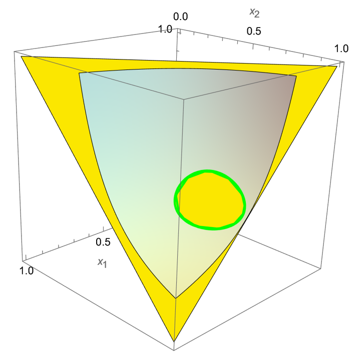

(b) $U_2 = x_1^{0.33}x_2^{0.43}x_3^{0.24} \equiv 0.68$

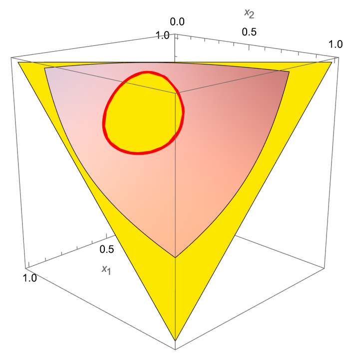

(c) $U_3 = x_1^{0.35}x_2^{0.25}x_3^{0.4} \equiv 0.67$

Figure 2. $\hat{C}$ and One Indifference Curve For Each Player

As we can see from Figure 2, the restriction of each indifference curve to the boundary is an oval. We may conceive of preferences for individual i over points in $\hat{C}$ as described by a series of ovals radiating out from x_i^*. Points on smaller ovals are preferred to points on larger ovals.

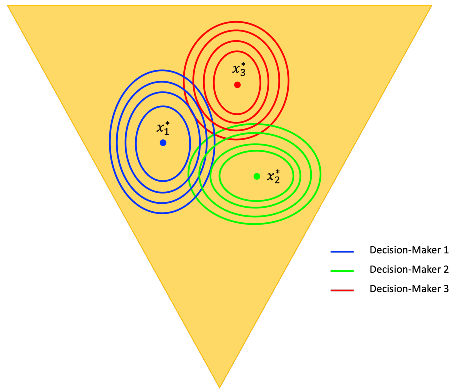

Figure 3. Three Decision-Makers' Indifference Curves

An important geometric feature is worth highlighting: for any point y in $\hat{C}$, when we draw the indifference curves for our three individuals (see Figure 3) that contain the point y, we can find petal-shaped regions (see Figure 4).

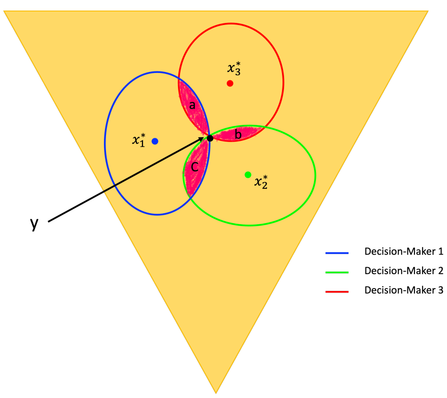

Figure 4. Win-set of y on $\hat{C}$

The pink petal-shaped regions include policies that at least 2 of the individuals strictly prefer to y. Specifically speaking, decision-maker 1 and decision-maker 3 prefer any policies in area a to y, because any policies in area a are on a smaller blue oval and a smaller red oval. Decision-maker 2 and decision-maker 3 prefer any policies in area b to y, because any policies in area b are on a smaller green oval and a smaller red oval; Decision-maker 1 and decision-maker 2 prefer any policies in area c to y, because any policies in area c are on a smaller green oval and a smaller blue oval. The union of these three petals (regions a, b, and c) is called the win-set of y.[12] That is to say, any policy y in $\hat{C}$ will be beaten by any policies in its petal-shaped region. Therefore, in this example, there is not a policy that can beat all other policies under majority rule, and thus no policy choice can reflect the will of the three decision-makers.

The key feature in our example is that the three investors' ideal points are not on a line. If, instead, the three investors' preferences happen to satisfy the condition that their induced ideal points are co-linear, then things would be different; in this case, the ideal point of the middle voter is stable. Figure 5 shows this possibility.

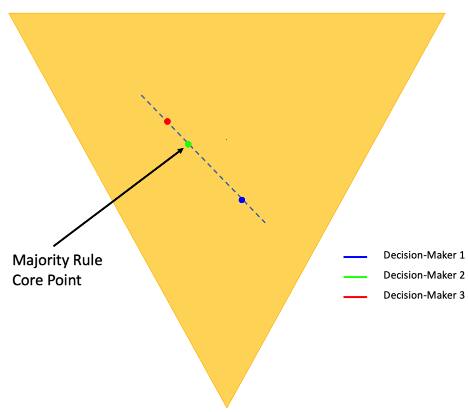

Figure 5. Collinear Optimal Points of Decision-makers

To see that in this example decision-maker 2's optimal point is the majority rule core point, consider a movement from her ideal point. Suppose that the movement is on the line connecting the three ideal points. Then, the movement makes both the middle voter and one

[12]The presence of points that beat y can also be seen by appeal to a different mathematical tool. Recall that the gradient vectors of a utility function evaluated at a point y indicate directions of movement that result in the highest utility gains. If we focus just on the surface $\hat{C}$, from any point y in $\hat{C}$, the gradient vectors on the surface at y of at least two of our individuals' utility functions differ by less than 90 degrees. This means that y would be beaten by other points that obtain from a small movement in the direction of any convex combination of the two agent's gradient vectors in $\hat{C}$. Thus, y is not stable or an element of the majority rule core. This argument is true for any point y in $\hat{C}$ (or in C).

of the two extremal voters worse off, as the movement deviates from the middle voter's ideal point and simultaneously moves away from one of the extremal voters. Thus, two voters object to the move. On the other hand, suppose the movement from the middle voter's ideal point is off the line connecting the three ideal points. Then, if the move is in a direction normal (perpendicular) to the line, the move makes all three voters worse off. If the move is not a direction normal to the line, it makes two (or possibly all three) voters worse off.

When there are more than three voters, it is possible for points to be stable (the majority rule core is non-empty), even if the ideal points are not on a line. However, the possibility is extremely rare and requires a very knife-edge condition. Below we analyze this in general cases of $k \geq 3$.

B.2. More than Three voters and/or More than Three Dimensions

In our three-voters and three-dimensional example, we have seen that the majority rule core is generally empty. Now, we derive geometric conditions on the primitives, $\alpha_1, \alpha_2, ..., \alpha_n$, that are necessary and sufficient for the majority rule core to be empty when $k \geq 3$. We then examine how robust a model satisfying these conditions can be.

Our derivations are less tedious if we work with an increasing monotone transformation of the Cobb-Douglas functions ($U_i(x) = ln(u_i(x))$. Namely, voter i's optimal policy in C solves

$$\max_{x_1, x_2, x_3 \ldots x_j} \quad ln(u_i(x)) = \Sigma_{j=1}^{k} \alpha_{i,j} ln(x_j)$$

$$s.t. \quad \Sigma_{j=1}^{k} \beta_j x_j \leq 1$$

$$x_j \geq 0$$

Solving this optimization program, we obtain

$$x_i^* = (\frac{\alpha_{i,1}}{\beta_1}, \frac{\alpha_{i,2}}{\beta_2}, ..., \frac{1-\Sigma_{j=1}^{k-1}\alpha_{i,j}}{\beta_k})$$

To understand preferences for local changes in policy on $\hat{C}$, we derive the gradient of u_i on the surface $\hat{C}$ at a point $x \in \hat{C}$. We obtain this gradient, by directly substituting the equality constrain $x_k = \frac{1-\Sigma_{j=1}^{k-1}\beta_j x_j}{\beta_k}$ and then differentiating with respect to the first $k-1$ dimensions of x. This new axis system results in a dimension reduction and ensures that we are moving along the surface $\hat{C}$ as we change a coordinate of x. With this axis system, the gradient vector for voter i and point $x \in \hat{C}$ is the $k-1$ dimensional vector

$$\nabla u_i(x) = (\frac{\alpha_{i1}}{x_1} - \frac{\beta_k \alpha_{ik}}{\beta_1 x_k}, ..., \frac{\alpha_{i(k-1)}}{x_{k-1}} - \frac{\beta_k \alpha_{ik}}{\beta_{k-1} x_k})$$

Local Stability requires that any movement from x makes at least $\frac{n+1}{2}$ voters worse off. Accordingly, a point x is locally stable if and only if no direction of movement from x can make at least $\frac{n+1}{2}$ voters better off. Our first result draws on this intuition about stability and establishes geometric properties of the gradient vectors of the utility functions at a core point.

Lemma 2 *Assume $n \geq 3$ and $k \geq 3$, the point $x^* \in C$ is in the majority rule core if and only if the following conditions are true: (1) $y \in \hat{C}$ and (2) one voter, j, has $x_j^* = x^*$ or put alternatively, $\nabla u_j(x^*) = 0$ and (3) for voters $N_{-\{j\}}$ there exists a pairing $p : N_{-\{j\}} \to N_{-\{j\}}$ (i.e. it is is one-to-one and $p(p(i)) = i$) such that for each $i \in N_{-\{j\}}$, there exists a positive scalar, λ_i s.t. $\nabla u_i(x^*) = -\lambda_i \nabla u_{p(i)}(x^*)$.*

The proof appears in the appendix.

Note with an even number of voters, this condition is sufficient but not necessary. Moreover, if we relax the assumption that players do not have identical preferences, then the conditions in 2 and 3 can be extended so that n is partitioned into a set of voters with ideal points at the core and a set of voters whose gradients can be paired off in the manner of (3).

We can relate the pairing of gradients in condition (3) to the underlying primitives, $\alpha_i' s$ and β, and directly examine how common it is for a parameterization to admit a point x^* satisfying lemma 1.

Assume that x^* is a core point. Taking voter $i \in N_{-\{j\}}$ and voter $p(i)$ and writing the equality for dimension d from condition (3) we obtain the equality

$$\beta_d x_k^* \alpha_{id} - x_d^* \beta_k \alpha_{ik} = -\lambda_i (\beta_d x_k^* \alpha_{p(i)d} - x_d^* \beta_k \alpha_{p(i)k})$$

Using the fact that the core point coincides with j's constrained ideal, $x_d^* = \frac{\alpha_{jd}}{\beta_d}$, we obtain

$$\frac{\alpha_{jk}}{\alpha_{jd}} = \frac{\beta_k^2}{\beta_d^2} (\frac{\alpha_{ik} - \lambda_i \alpha_{p(i)k}}{\alpha_{id} - \lambda_i \alpha_{p(i)d}})$$

Finally, we say a statement is generically true if in the space of parameters, it is true on a set that is both open and dense. For any profile of technology constraints, β the set of possible profiles of voter preferences that are Cobb-Douglas over k dimensions is R_{++}^k (strictly positive α_{id}'s). It is not difficult to see then that a statement is generically true if the subset of this parameter space upon which the statement is true has full rank.

If x^* is a majority rule core point, then for each of the $n-1$ voters with induced ideal not equal to x^* each of their gradients is linearly dependent (α_{id} is related to $\alpha_{p(i)d}$) for dimensions 1 through $k-1$. In other words, once you select α_j you can then freely select $\frac{n-1}{2}$ values of $\nabla U_i(x^*)$ (each is of dimensionality $k-1$) and $\frac{n-1}{2}$ scalars λ_i (each of dimensionality

1).This means that the rank of the parameters $\{\alpha_i\}_{i\in N-\{j\}}$ is $(k-1+1)[\frac{n-1}{2}]$. Meanwhile, the parameters for voters other than j live in $R^{(k-1)(n-1)}$. Accordingly, The parameters satisfying lemma 1 are full rank if

$$(k)\frac{n-1}{2} \geq (k-1)(n-1)$$

We thus need

$$k(n-1) \geq 2(k-1)(n-1)$$

$$k \geq 2k-2$$

$$2 \geq k$$

This leads to the following result.

Proposition 2 *If $n \geq 3$ (odd), the majority rule core is generically empty if $k \geq 3$.*[13]

B.3. Extension to supermajority rules

Now we relax the focus from simple majority rule to a set of rules termed non-collegial counting rules. In particular, non-collegial counting rules require q ($n > q > \frac{n-1}{2}$) of n voters to support a policy change in order for the change to occur. Then, the following result can be obtained by modifying the argument in Saari (1997) to our setting

Proposition 3 *Assume $n \geq 5$. For any or $n > q > \frac{n-1}{2}$, the core of the q-rule is generically empty if*

$$k > 1 + 2q - n + max\{\frac{4q-3n-1}{2(n-q)}\}$$

[13]Note with n even it is easier for points to be stable and in fact the result changes to $k > 3$. See Saari (1997).

This extension shows that while supermajority requirements make it easier to find stable policies or to justify policy choices as the will of the voters, we cannot escape indeterminacy as the number of dimensions grows. The movement to supermajority rules introduces another challenge. These rules make it easier for policies to be stable and so determining which of the stable policies will be chosen may be difficult. Put slightly differently, predicting which stable policy will be chosen may be difficult.

V. Agenda Control

The above results are negative; they indicate that minimal heterogeneity of voter tastes leads to a situation where a choice cannot be justified as the will of a majority if there are more than two dimensions. When the firm limits consideration to only one aspect of socially responsible behavior, the take-away is more satisfying. The choice that corresponds to the will of the voter with the median preference parameter is natural (stable). One takeaway from the negative social choice results and, in particular, our proposition 2 is that a positive theory of firm choice requires more than just preferences as an input.

We now delve deeper into the process and show how the presence of agenda-setting power can partially drive choices and ask whether dimensionality is important once additional aspects are considered. For example, imagine that a member of the board of directors or a key activist shareholders (e.g., a hedge fund investing in the firm) has the power to determine the voting agenda. Would this change the nature of the problem? In this section, we establish two additional conclusions: (1) the ability to control the agenda is important and allows an agent or agents to strongly influence policy choice, but (2) this power is not absolute, as the will of the majority operates as a constraint on what an agenda setter can achieve. We find that when there is not a natural agenda setter, so investors face uncertainty about who will control the agenda, the case of more than 2 dimensions may be a source of volatility in decision-making.

The economics and political science literature contains many treatments of agenda-setting. For our purposes, we highlight two particularly extreme institutional structures. Each has been employed in several key papers. In the first, we assume that a privileged agent (e.g., the CEO of a firm or the chairperson of a firm's board) has monopoly control of the agenda. She can select a policy that will be pitted against the status quo in a pairwise comparison under majority rule. Romer and Rosenthal (1978) applied this model to the study of school funding. In the second, other extreme case, we assume that agenda-setting power is far less concentrated. For example, consider a firm is invested by several activist hedge funds, and every manager of these activist hedge funds possibly presents a shareholder proposal. The approach dates back to Baron and Ferejohn (1989) in the case of distributive politics and has been used by Banks and Duggan (2000) and others in the context of spatial preferences. We assume that policy is selected in a dynamic game. In each period, a proposer is selected by "random recognition." We assume that each agent has a chance of being recognized to propose a policy. Her proposal is then pitted against the status quo. If her proposal passes, the game ends. If her proposal fails, a new proposer is randomly selected, and this process continues until a choice is accepted.

A. Monopoly Agenda Control: Proposing by CEO or Chairperson

While the classic treatment of proposal and voting rights is Romer and Rosenthal (1978) in the case of public policy budgeting, applications to firm decision-making are natural. Marino and Matsusaka (2005) have looked at the domain of corporate budget setting, and

Matsusaka and Ozbas (2017) develop a model and trace out the relation between the concentration of proposal power and firm value in the presence of activism. These applications focus on the case of a one-dimensional choice. Here we move to higher dimensional problems.

A natural starting point is to posit that one of the voters has agenda-setting power(e.g, a CEO or the chairperson of a board). We call this voter m. Voter m (and she alone) can propose a policy alternative. Then, all voters must collectively decide whether to accept this alternative or not. Our model is only closed if voters are explicitly (or implicitly) comparing the proposal to some other source of utility. For convenience, we assume that there is a default or status quo policy, $s \in C$, and that voting is between enacting x or maintaining s following the proposal x by the agenda setter. Appeal to sub-game perfect equilibrium in which voters do not use weakly dominated strategies is natural in this game. It is straightforward to see that m will select the policy that maximizes her utility from the set of policies that can beat s.

To characterize the set of policies that can beat x under majority rule, we return to the idea of the win set of s depicted earlier as the intersection of the pink petals in our three-voters example (Figure 4). Formally define D to be the set of coalitions in N containing at least $\frac{n+1}{2}$ voters. For a voter $i \in N$, let $P_i(s)$ denote the set of policies that i weakly prefers to s. The win set is then defined as $W(s) = \cup_{d \in D} \cap_{i \in d} \{P_i(s) \cap \hat{C}\}$. The optimal policy for m to propose then solves

$$Max_{x \in W(s)} u_m(x)$$

Standard arguments can be used to show that the above problem has at least one solution when the preferences are continuous and the feasibility set is compact (conditions satisfied in our Cobb-Douglas specification with finite vector β). But, as can be seen in Figure 7, the win-set is not generally convex. This is not just a technical observation. When the win-set is not convex, multiple solutions to the proposer's problem can exist, and small changes in the preferences or location of s can lead to large changes in the final policy. However, building on the main theme of the previous comparisons, we know that things are different in the case of $k = 2$. When $k = 2$, the median voter's preferences are decisive in the sense that a policy x beats s under majority rule if and only if the voter with the median preference parameter prefers x to s. This means that the win-set is convex when $k = 2$, and so the proposer's constrained optimization problem has a unique solution, which is continuous in s and the other parameters.

We illustrate the above points with two examples. The first example is based on example A.1 (two dimensions and three voters). But now we consider that voter 1 is the proposer who can propose a policy x. All three voters then vote to determine whether to implement the proposal x or to keep the status quo s. Suppose that the current policy s is on the right of the median voter 2's favorite policy, as shown in Figure 6. Obviously, voter 1 does not benefit from proposing any policies on the right of s, as those policies are even further to her optimal policy x_1^* than the current policy s. Although voter 1 prefers policies on the left of p, a proposal of a policy on the left of p will be voted against by voter 2 and voter 3, and thus must fail. Therefore, in equilibrium, voter 1 wants to propose policy p, which will be accepted and replace s.

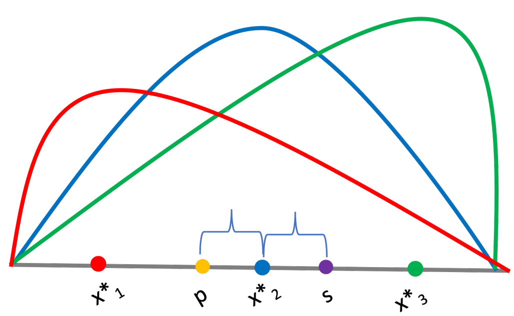

Figure 6. Two Dimensions and Three Voters

Note that although the proposer does have the power to move the policy towards her optimal one, the proposer is constrained by the status quo and other voters; she has to win the median voter by ensuring the median voter's payoff from accepting her proposal to be at least the same as from keeping the status quo.

In the second example, we consider there are four players the pass of a proposal requires the support of at least three players in order to move from the status quo s depicted as the black dot in Figure 7. Figure 7 exhibits the win-set for three players and the indifference curve for the proposer through her favorite policy that can beat s. At one parameterization of the proposer's preferences, the black curve represents her indifference curve through an optimal policy that is tangent with the blue contour. Then, the proposer can replace s with policy 1, because her proposal of policy 1 can win support from the blue voter and the red voter. But if we parameterize the proposer's preferences slightly differently so that her indifference curve becomes the dotted curve, then the proposer can now have two optional coalitions. In particular, if she proposes policy 1, she will have support from the blue and the red voter, and thus policy 1 will be implemented. If she proposes policy 2, she will win over the green and the red voter, and thus policy 2 will be chosen. Because both policy 1 and policy 2 lie on the dotted indifferent curve, the proposer is indifferent with these two policies and thus these two coalitions. Either policy 1 or policy 2 can be the equilibrium voting result, and which policy the proposer proposes will be chosen. Similarly, if we do another nearby parameterization so that the dotted and dash curve represents the proposer's indifference curve, then coalition with the blue and the red voter is no longer optimal. The proposer will propose policy 2, which will have the support of the green voter and the red voter. Thus, the resulting policy becomes 2.

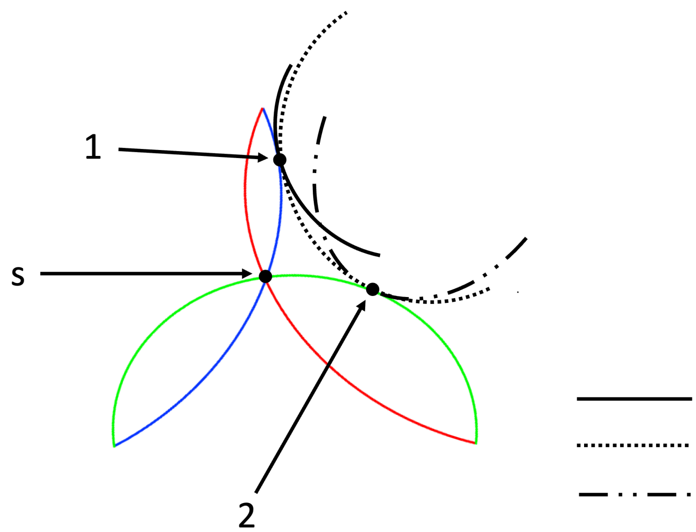

Figure 7. Optimal proposal by Agenda Setter

This example illustrates that small changes in the proposer's preferences result in a large jump in the resulting policy. The above analysis and the two examples lead to the proposition below.

Proposition 4 *Consider the monopoly proposing model. With $k = 2$, there is a unique equilibrium, and the final policy x is continuous in the parameters, s, β, α. With $k > 2$, there is an equilibrium, but it may not be unique, and small changes in parameters may yield non-negligible changes in the equilibrium.*

The proof appears in the appendix.

A few qualifications to the bad news of the above result are in order. The pathology of multiple solutions is rare. Multiplicity occurs when the proposer is indifferent between two or more policies, each of which is supported by different coalitions. From any parameterization that leads to this, a small perturbation can be found which makes the proposer strictly prefer one form of coalition to the other(s). Thus, the multiplicity occurs on a non-generic set of parameter values. The discontinuities are intimately related to multiplicity. In order to find cases where small changes in parameters do not lead to small changes in outcomes, we must pass through a pathological example with multiple equilibria. This violation in continuity is not a violation of upper or lower hemi-continuity of the equilibrium correspondence. It is tied to jumps from one segment of the equilibria correspondence to another.

Moving from this technical discussion to a substantive interpretation, the point is this. With monopoly agenda setting, if one knows parameters, a unique and stable prediction is typically available. But sometimes, the set of equilibria can be larger, and the final policies that are possible will be quite different. Moreover, occasionally a little change in the parameters or perhaps a little uncertainty on the part of the investor about what the parameters are can lead to a sense that two (or more) very different policies might emerge.

B. Diffuse Agenda Control: Proposing by Several Activist Shareholders

Although the idea that the generation of policies to be voted upon within the firm may fall to a single agent or a small set of homogenous agents can be natural, other perspectives are also defensible. Banks and Duggan (2000) develop and analyze an infinite horizon model in which each voter is endowed with preferences (our Cobb-Douglas preferences over a constraint set fall within the class they focus on) and each voter has a probability ρ_i of being recognized as proposer in any period. The case of monopoly agenda setting can be viewed as one in which $\rho_i = 1$ for the monopoly agenda setter and 0 for all other voters. The case of symmetric proposal power has $\rho_i = \rho$ for all players.

In corporate governance settings, we consider a firm owned by n activist shareholders (e.g, hedge funds or active mutual funds). Each activist shareholder has one share. In period one, an activist shareholder is randomly chosen as the proposer with a probability of ρ_i. In reality, an activist shareholder with higher ρ_i can be considered as an activist shareholder who owns more shares or has more experience and reputation and thus is more likely to present a shareholder proposal. The proposer makes a policy proposal, and then everyone votes. If the policy gains majority support, it is implemented, and the game ends. If not, then we move to the next period. In the next period, a new randomization occurs to select the proposer. Again, each activist shareholder becomes the proposer with a probability of ρ_i, and the new proposer makes an offer. This new offer is voted up or down. If the new offer is accepted by a majority, then it is implemented. If not, the game continues. All activist shareholders are assumed to discount policy payoffs so that acceptance of policy x in period t is discounted by rate δ compared to acceptance in period $t-1$.

Banks and Duggan (2000) focus on a parsimonious class of equilibria, stationary equilibria, that exhibit no delay and involve weakly undominated voting strategies. Stationarity requires that voting and proposing strategies are time and history-independent. No delay requires that each possible proposer makes an offer that is accepted, and thus the game ends with acceptance of the first proposal. Note that because all players discount, any equilibrium that exhibits delay is Pareto Dominated by some (possibly random) strategy profile that does not exhibit delay. In this sense, no-delay is normatively attractive. Banks and Duggan (2000) show that there are no-delay stationary equilibria to this game and provide deep connections between equilibrium outcomes of the game and the structure of the core.

We focus on the volatility of firm policy. Consider a setting with n odd and $k = 2$. As we saw above, in Section IV.A, $\hat{C}$ is one-dimensional, and the majority rule core coincides with the constrained optimal policy of voter $i^* = \frac{n+1}{2}$. In particular, the core is $x^* = (\frac{\alpha^1_{i^*}}{\beta_1}, \frac{\alpha^2_{i^*}}{\beta_2})$. Theorem 5 of Banks and Duggan (2000) yields the following characterization for this class of examples

Proposition 5 *With $k = 2$, all no-delay stationary equilibria involves each voter proposing x^* for some x^* in the majority rule core of $\hat{C}$. In particular, with odd n, all proposals are equal to x_i^*, and there is no uncertainty about the policy which will be chosen in any no-delay stationary equilibrium.*

The proof appears in the appendix.

In contrast, when the core is empty (generically the case when $k > 2$), different proposers will have strong incentives to make different proposals. The logic behind this claim can be seen by returning to Figure 4, which exhibits the case of three voters bargaining over a triangle. The emptiness of the core means that from any offer, it is possible to find an alternative offer that two players like better. This means that in a three-player example, if all players were supposed to make the same proposal, then at least two of those players would recognize there is a profitable deviation (namely, to propose something they both like better than the point they are supposed to propose). Lemma 2 from Banks and Duggan (2000) builds on this logic and establishes that if the core is empty, then every equilibrium involves randomness over what policy will be chosen.

Proposition 6 *With $k > 2$, generically, every no-delay stationary equilibrium puts a positive probability on achieving at least two distinct policies.*

The proof appears in the appendix.

It is worth noting that there are two sources of this volatility when the core is empty. First, equilibrium may require that some (or all) proposers are randomizing over what policy they propose. This behavioral uncertainty propagates into policy uncertainty. The second source of volatility strikes us as conceptually more important. When the core is empty, different proposers will be able to make different proposals and have their proposals passed. Each extracts some rents from their privileged status as proposers (as in the gains from moving first in bargaining models like Rubinstein). A key fundamental in this modeling approach is uncertainty about who might gain control of the agenda at any given moment. This uncertainty reflects uncertainty about fine details of the likely behavior of key decision-makers. The equilibrium analysis shows that this uncertainty about the process or institution propagates into uncertainty about the policy outcomes from governance. Comparing the last two propositions allows us to relate policy volatility to the dimensionality of the policy space.

Proposition 7 *If we conceive of firm policy-making as the infinite horizon bargaining game in which each voter has probability ρ_i of being recognized, then moving from the case of $k = 1$ to $k = 2$ does not increase policy volatility, but moving to $k > 2$ necessarily does increase policy volatility.*

It is worth taking stock of the insights from the monopolistic and diffuse agenda power models. What happens when we move from $k = 2$ to $k > 2$? If we are in the monopolistic agenda power model, then volatility generally does not increase, but any observer/investor's uncertainty over what will happen may increase. If we are in the diffuse agenda power model, then increasing dimensionality from 2 generically increases volatility, because we move from an equilibrium in which every proposer proposes the same policy to an equilibrium in which each possible proposal proposes possibly different policies which will be passed.

VI. Who Votes?

Finally, we note that our theoretical treatment of corporate governance by necessity abstracts away from important nuances about who is involved in governance. For example, we have ignored the entry and exit channels. As previously discussed, investors can affect firm policies through two main channels: voice and exit (Edmans (2014)). A natural difference between corporate governance and voting by political elites is that shareholders may decide to sell and thus exit if they are unsatisfied with the choices made. It is natural to ask how this feature might alter our conclusions. One answer is that this possibility can create interesting dynamic incentives. Policy choice today can influence who votes the next time a policy choice is made and players that exhibit this patience will have to anticipate the effects of their proposals on the make-up of the voting body in the future. Gieczewski and Kosterina (2021) consider this in a model that is particularly suited to areas of R&D. Second, we ignored the possibility of entry. Several papers examine stock lending and empty voting (e.g., Christoffersen, Geczy, Musto, and Reed (2007)). These practices could make the pool of who votes and the number of votes each decision-maker casts endogenous as well.

In settings where membership is endogenous, there is room for certain actors to strategically influence membership. To the extent that an investor believes firm choices will be more congruent with her preferences, if she believes that a large enough portion of other shareholders has similar preferences, there is scope to consider public announcements by large investors or blocks that may impact who decides to buy or keep their shares.

To flesh out the point, consider a setting in which there are several similar firms. Further imagine that for each of these firms, a large block has announced their preferences, α_i. An investor shopping between these firms may have a preference for the one that has a large block of voters with preferences that are more congruent to her own. Thus, public declarations by the large block can help drive sorting by other investors and create niches. This comment is by necessity speculative, as the literature on dynamic policy bargaining possesses several forms of non-monotonicity results, and a full analysis of the dynamic model is beyond our scope.[14] That caution aside, Eraslan (2016b) shows that in infinite horizon multiplayer divide the dollar games payoffs are non-decreasing in one's recognition probability and so if we associate the size of a block with the likelihood of placing a proposal on the agenda, payoffs would likely be positively related to the number of other voters with the same preferences.[15]

We can, however, exhibit a strong result if in fact blocks are sufficiently large. A convenient assumption in our analysis has been that the preference profiles were distinct ($\alpha_i \neq \alpha_j$ for any two voters i and j). In the case of block voting, however, we may think that there are a number of votes controlled by a single decision-maker. For example, in 2020, an investment firm named Engine No. 1 controlled approximately 0.02% of the voting shares in Exxon Mobil however it was able to replace three members on the Exxon Mobile board of directors by convincing a block of voters to elect directors who acknowledged climate change. In such situations, a very strong result obtains.

Proposition 8 *In settings where m agents control n votes, if there exists an agent i with at least $\frac{n+1}{2}$ votes then $M(C) = \{x_i^*\}$.*

The proof is a so-called one-liner. Because voter i prefers x_i^* to every other feasible policy x' and i casts at least a majority of votes, definition 1 implies that the majority preference ordering must rank x_i^* ahead of every other policy x' and thus by definition 2 the majority rule core must coincide with x_i^* This is true regardless of the number of dimensions. Thus,

[14] At some level, it is sufficient to return to our discussion of the non-convexity of the win sets. This fact allows one to create counter-examples to many seemingly true conjectures.

[15] Of course, an additional cause of caution stems from the fact that Eraslan studies a model of purely distributive bargaining and not a spatial bargaining model and folk-knowledge that infinite horizon models counter-examples to this monotonicity exist.

in cases where an investor sees firms controlled by large enough blocks, she will have an unambiguous preference for the firm whose controlling block has preferences closest to her own, all else equal. Moreover, we may conceive of the strategy of amassing sufficiently large blocks as serving the purpose of protecting oneself from the volatility highlighted in the previous section. Of course, not all agents face this degree of liquidity, and there are many other factors that may speak against amassing this much stake in a firm.

VII. Conclusion

The amount of capital allocated to socially responsible investments has increased dramatically over the last decade. As a result, firm managers are increasingly being pressured to consider multi-dimensional objective functions. We examine whether and how changes to the number of objectives of managers or investors affect voting outcomes and the choice of firm policies. While the traditional objective of maximizing firm value can be expressed as a one-dimensional voting problem, the proliferation of socially responsible investment goals leads to a multi-dimensional choice problem in which different voters have different preferences about how to trade-off these dimensions. Building on the large literature on social choice theory, we study voting for corporate policies and highlight several key intuitions.

In contrast to the view that it is "logically impossible to maximize in more than one dimension" (Jensen (2001)), we show that it is possible for a voting system like majority rule to lead to a stable choice that reflects investor preferences when voters face a trade-off between maximizing firm value and *one* social policy. While, in general, voting does not tend work well with two dimensions, because there are natural trade-offs between maximizing firm value and incorporating social objectives we show the problem is essentially one dimension smaller. However, when voters face a trade-off between maximizing firm value and more than one social policy, this result generally no longer holds. In particular, we show that a number of challenges arise: (1) the will of shareholders need not be well-defined, (2) it will not be possible to identify stable choices by thinking only about voting (3) the choices that emerge may depend on who has agenda-setting power, and (4) the volatility of firm choices will tend to increase, especially if agenda-setting power is diffuse. Overall, our findings have important implications for the stakeholder theory of the firm. Namely, when decision-makers care about more than two objectives, it may adversely affect the quality of corporate governance.

References

Admati, A. R., Pfleiderer, P., & Zechner, J. (1994). Large shareholder activism, risk sharing, and financial market equilibrium. *Journal of Political Economy*, *102*(6), 1097–1130.

Anton, M., Ederer, F., Gine, M., & Schmalz, M. C. (2023). Common ownership, competition, and top management incentives. *Journal of Political Economy*, *forthcoming*.

Arrow, K. (1951). A difficulty in the concept of social welfare. *Journal of Political Economy*, *58*, 328–346.

Austen-Smith, D., & Banks, J. S. (2000a). *Positive Political Theory I: Collective Preference.* University of Michigan Press.

Austen-Smith, D., & Banks, J. S. (2000b). *Positive Political Theory II: Strategy and Structure.* University of Michigan Press.

Azar, J., Tecu, I., & Schmalz, M. C. (2018). Anti-competitive effects of common ownership. *Journal of Finance*, *73*(4), 1513-1565.

Banks, J. S., & Duggan, J. (2000). A bargaining model of collective choice. *The American Political Science Review*, *94*(1), 73–88.

Baron, D. P., & Ferejohn, J. A. (1989). Bargaining in legislatures. *The American Political Science Review*, *83*(4), 1181–1206.

Berk, J. B., & van Binsbergen, J. H. (2021). The impact of impact investing. *Working Paper*.

Berle, A. A., & Means, G. C. (1932). *The Modern Corporation and Private Property.* New York: Macmillan Publishing Co.

BlackRock. (2022). *Blackrock expands voting choice to additional clients.* (`https://www.blackrock.com/corporate/newsroom/press-releases/article/corporate-one/press-releases/2022-blackrock-voting-choice`)

Bolton, P., Li, T., Ravina, E., & Rosenthal, H. (2020). Investor ideology. *Journal of Financial Economics*.

Brav, A., Jiang, W., & Li, T. (2018). Picking friends before picking (proxy) fights: How mutual fund voting shapes proxy contests. *Working Paper*.

Bubb, R., & Catan, E. (2018). The party structure of mutual funds. *Working Paper*.

Burkart, M., Gromb, D., & Panunzi, F. (1997). Large shareholders, monitoring, and the value of the firm. *The Quarterly Journal of Economics*, *112*(3), 693–728.

Christoffersen, S., Geczy, C., Musto, D., & Reed, A. (2007). Vote trading and information aggregation. *Journal of Finance*, *62*(6), 2897–2929.

Dasgupta, A., & Piacentino, G. (2015). The wall street walk when blockholders compete for flows. *The Journal of Finance*, *70*(6), 2853–2896.

DeMarzo, P. M. (1993). Majority voting and corporate control: The rule of the dominant shareholder. *Review of Economic Studies*, *60*(3), 713–734.

Demsetz, H. (1983). The structure of ownership and the theory of the firm. *The Journal of Law and Economics*, *26*(2), 375–390.

Denis, D. (2016). Corporate governance and the goal of the firm: In defense of shareholder wealth maximization. *The Financial Review*, *51*, 467–480.

Dikolli, S. S., Frank, M. M., Guo, M. Z., & Lynch, L. J. (2021). Walk the talk: Esg mutual fund voting on shareholder proposals. *Working Paper*.

Edmans, A. (2009). Blockholder trading, market efficiency, and managerial myopia. *The Journal of Finance*, *64*(6), 2481–2513.

Edmans, A. (2014). Blockholders and corporate governance. *Annual Review of Financial Economics*, *6*(1), 23-50.

Eraslan, H. (2016a). Uniqueness of stationary equilibrium payoffs in the baron–ferejohn model with risk-averse players. *International Journal of Economic Theory*, *12*(1), 29–40.

Eraslan, H. (2016b). Uniqueness of stationary equilibrium payoffs in the baron-ferejohn model with risk averse players. *International Journal of Economic Theory*, *12*(1).

Fisher, I. (1930). *The theory of interest: As determined by impatience to spend income and opportunity to invest it*. New York: Kelley and Millman.

Friedman, M. (1970, 01). The social responsibility of business is to increase its profits. *New York Times Magazine*, *32*, 173-178. doi: 10.1007/978-3-540-70818-6_14

Gantchev, N., Giannetti, M., & Li, R. (2020). Sustainability or performance? ratings and fund managers' incentives. *Ratings and Fund Managers' Incentives (November 15, 2020)*.

Gibson Dunn. (2020, August 4). Shareholder proposal developments during the 2020 proxy season. *White Paper*.

Gieczewski, G., & Kosterina, S. (2021). Endogenous experimentation in organizations. *Working Paper*.

Hart, O., & Zingales, L. (2017). Companies should maximize shareholder welfare not market value. *Journal of Law, Finance, and Accounting*, *2*, 247–274.

Hartzmark, S. M., & Sussman, A. B. (2019). Do investors value sustainability? a natural experiment examining ranking and fund flows. *Journal of Finance*, *74*, 2789–2837.

Heath, D., Macciocchi, D., Michaely, R., & Ringgenberg, M. C. (2021). Does socially responsible investing change firm behavior? *Working Paper*.

Jensen, M., & Meckling, W. (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, *3*(4), 305-360.

Jensen, M. C. (2001). Value maximisation, stakeholder theory, and the corporate objective function. *European Financial Management*, *7*, 297–317.

Levit, D. (2019). Soft shareholder activism. *The Review of Financial Studies*, *32*(7), 2775–

2808.

Levit, D., Malenko, N., & Maug, E. G. (2019). Trading and shareholder democracy. *Journal of Finance, Forthcoming, European Corporate Governance Institute–Finance Working Paper*(631).

Malenko, A., & Malenko, N. (2019). Proxy advisory firms: The economics of selling information to voters. *The Journal of Finance*, *74*(5), 2441–2490.

Malenko, N., & Shen, Y. (2016). The role of proxy advisory firms: Evidence from a regression-discontinuity design. *Review of Financial Studies*, *29*(12), 3394–3427.

Marino, A., & Matsusaka, J. (2005). Decision process, agency problems, and information: An economic analysis of capital budgeting procedures. *The Review of Financial Studies*, *18*(1).

Matsusaka, J., & Ozbas, O. (2017). A theory of shareholder approval and proposal rights. *Journal of Law Economics and Organization*, *33*(2).

Maug, E. (1998). Large shareholders as monitors: is there a trade-off between liquidity and control? *Journal of Finance*, *53*(1), 65–98.

May, K. (1952). A set of independent necessary and sufficient conditions for simple majority decisions. *Econometrica*, *20*(4), 680–684.

McCahery, J. A., Sautner, Z., & Starks, L. T. (2016). Behind the scenes: The corporate governance preferences of institutional investors. *Journal of Finance*, *71*, 2905–2932.

Meirowitz, A. (2004). The code of social choice problems with monotone preferences and a feasibility constraint. *Social Choice and Welfare*, *22*(3), 467–478.

Michaely, R., Ordonez-Calafi, G., & Rubio, S. (2021). Es votes that matter. *Working Paper*.

Oehmke, M., & Opp, M. (2020). A theory of socially responsible investment. *Working Paper*.

Papadopoulos, K. (2019, February 6). The long view: The role of shareholder proposals in shaping u.s. corporate governance (2000-2018). *Harvard Law School Forum on Corporate Governance*.

Plott, C. R. (1967). A notion of equilibrium and its possibility under majority rule. *American Economic Review*, *57*(4), 787–806.

Romer, T., & Rosenthal, H. (1978). Political resource allocation, controlled agendas, and the status quo. *Public Choice*, *33*(4), 27–43.

Saari, D. (1997). The generic existence of a core for q-rules. *Economic Theory*, *9*, 219–260.

Vanguard. (2023). *Empowering everyday investors through proxy voting choice.* (https://corporate.vanguard.com/content/corporatesite/us/en/corp/articles/empowering-everyday-investors-proxy-voting-choice.html)

A. Appendix

A. Proofs

A.1. Proof of Proposition 1

Proof. Suppose by contradiction that $x \neq x^*_{\frac{n+1}{2}}$ but $x \in M(C)$. Moreover, suppose that x is on the left of $x^*_{\frac{n+1}{2}}$. Then, the median voter and the other $\frac{n-1}{2}$ voters who are on the right of the median voters strictly prefer $x^*_{\frac{n+1}{2}}$ to x. So, any point on the left of $x^*_{\frac{n+1}{2}}$ will be beaten by $x^*_{\frac{n+1}{2}}$.

Following the same logic, any point on the right of $x^*_{\frac{n+1}{2}}$ will also be beaten by $x^*_{\frac{n+1}{2}}$, because the median voter and the other $\frac{n-1}{2}$ voters who are on the left of the median voters strictly prefer $x^*_{\frac{n+1}{2}}$ to x. ■

A.2. Proof of Lemma 2

Proof. The result is stated and proven as corollary 5.1 in Austen-Smith and Banks (2000a) (p. 148) for an interior point and smooth and strictly convex utility functions in some d dimensional space. On $\hat{C}$ our induced preferences are smooth and strictly convex. Moreover, given the Cobb-Douglas utility, any point that is optimal for one player (satisfying part 1) must be interior in the interior of $\hat{C}$. ■

A.3. Proof of Proposition 4

Proof. For the case of $k = 2$, recall that $W(s)$ corresponds to the set $P_{med}(s)$ where *med* is the median voter. This set is non-empty ($s \in P_m(s)$), compact, and convex, given that the preferences of m are continuous and strictly convex. It is actually just a closed interval. Moreover, the objective function of the proposer is also continuous and strictly concave. By Berge's Theorem of the Maximum, the solution is a continuous correspondence of the parameters. Moreover, because of the convexity conditions, at most one solution exists for each parameterization. The equilibrium is a continuous function of the parameters.

For the case of $k > 2$, $W(s)$ is no longer generally a convex set and so uniqueness of policies that are optimal for m can fail. The second example is sufficient to demonstrate how multiple policies may be optimal for the proposer and that we can perturb parameters by any amount $\epsilon > 0$ and ensure that the resulting equilibria differ by more than ϵ. ■

A.4. Proof of Proposition 5

Proof. By the analysis supporting our proposition 1, the case of $k = 2$ is equivalent to a problem of bargaining over the 1-dimensional surface $\hat{C}$ with strictly convex preferences on this set. Theorem 5 of Banks and Duggan (2000) establishes that, in this case, all no-delay stationary equilibria involve all proposers offering exactly the same point in the core. ■

A.5. Proof of Proposition 6

Proof. By Banks and Duggan (2000) lemma 1, there is a no-delay stationary equilibrium in which each player proposes the same policy with probability one only if that point is in the core. However, by proposition 2, the core is generically empty when $k > 2$. ■