DOI: 10.1017.S0003055407070451

Leading the Party: Coordination, Direction, and Communication

TORUN DEWAN *London School of Economics*
DAVID P. MYATT *University of Oxford*

Party activists face a coordination problem: a critical mass—a barrier to coordination—must advocate a single policy alternative if the party is to succeed. The need for direction is the degree to which the merits of the alternatives respond to the underlying fundamentals of the party's environment. An individual's ability to assess the fundamentals is his sense of direction. These three factors—the barriers to coordination, the need for direction, and an individual's sense of direction—combine to form an index of both the desirability and the feasibility of leadership. We offer insights into Michels' Iron Law: a sovereign party conference gives way to leadership by an individual or oligarchy if and only if the leadership index is sufficiently high. Leadership enhances the clarity of intraparty communication, but weakens the response of policy choices to the party's environment. Our model can also be applied to the coordination problems faced by instrumental voters in plurality-rule elections, and so relates to the psychological effect of Duverger's Law.

What is leadership? When is it desirable? And when is it feasible? In defining leadership some authors have highlighted its focal role (Calvert 1995; Myerson 2004). More recently, Levi (2006, 10) argued that leadership "provides the agency that coordinates the efforts of others." Following these precedents, we develop a formal model in which the direction provided by leadership can help to coordinate the actions of a mass. We ask: is such direction best provided by one, a few, or the many? These institutional forms correspond to *de facto* dictatorship, oligarchy, and pure democracy. We explore the relative desirability of these institutional forms and consider their feasibility: when will members of a democratic body voluntarily follow the lead taken by either an individual or an elite subset of their membership? Michels (1915) offered unequivocal answers. He argued that leadership by an elite was both desirable and feasible because an organization can only coordinate effectively when led by a small subset of its membership; the mass voluntarily renounces its democratic rights. This view was based on his study of the German Socialist Party (a forerunner of the Social Democratic Party) which, though nominally adhering to the principle of conference sovereignty, allowed key decisions to be taken in the closed surroundings of its *fraktion* meetings. Michels observed that the conference generally rallied behind the party leadership.[1]

To provide our own answers to our questions, we model a coordination problem faced by a mass of political actors. Our main focus is a world in which our actors are the activist members of a political party. Each activist must advocate one of two policies. A policy succeeds if and only if a critical fraction of the party (a *barrier to coordination*) supports it; an uncoordinated party splits and fails. Party members would like the best policy to succeed. A stylized interpretation of the mid-1990s reform of the British Labour party helps to fix ideas: one policy would be the adoption of Tony Blair's "New Labour" program, whereas the alternative would be the retention of "old" Labour ideals. Whatever the alternatives' merits, a unified party could challenge for power; a split would have relegated the party to the wilderness.

Interestingly, our coordination problem also arises in plurality-rule elections. Consider the well-known 1970 New York senatorial election (Table 1) in which two liberals, Ottinger and Goodell, competed against the conservative Buckley. The liberal vote split, and so Buckley won. However, because the total liberal vote exceeded Buckley's tally, sufficient coordination among liberal voters (a sufficient condition would be if two thirds of them had backed the same candidate) would have prevented the conservative win.

Given their common objective, it seems obvious that all activists should advocate the best policy. (Or, for the New York voting game, all liberals should vote for the best liberal candidate.) This is difficult when the identity of the best policy is uncertain. This is so when the merits of policies depend on the political world in which the party lives, because party members' assessments of their world may differ. Leadership may help; as Levi (2006) suggested, it "provides the learning environment that enables individuals to transform or

Torun Dewan is Lecturer, Department of Government, London School of Economics, London WC2A 2AE, UK (t.dewan@lse.ac.uk).

David P. Myatt is Senior Research Fellow, Department of Economics, University of Oxford, Oxford OX1 3UQ, UK (david.myatt@economics. ox.ac.uk).

We thank Steve Fisher, Justin Fox, Abhinay Muthoo, Barry Weingast, Lee Sigelman, Chris Wallace, and three anonymous referees, as well as seminar participants at Essex, LSE, and at the APSA annual meetings in Philadelphia 2006 for their helpful comments and suggestions.

[1] For an application of Michels' work to British political parties see McKenzie (1963), and for a critical review see Hands (1971). Michels claimed that elites develop to pursue goals that betray the original movement. Modern interpretations have highlighted the similarities with the process of agency drift, suggesting that *ex ante* control mechanisms can counteract this effect (Lupia and McCubbins 2000). Koelble (1996), for example, using the arguments from the literature on Congress, suggested that "party activists prefer strict monitoring of the actions of the representatives, mechanisms for screening and selection which force the agent to report and act as delegate rather than as a trustee for the party." We do not address the agency issue, focusing instead on the conditions under which a party conference gives way to a leadership elite.

TABLE 1. The 1970 New York Senatorial Election

Candidate	Votes	Share
James R. Buckley	2,288,190	39%
Charles E. Goodell	1,434,472	24%
Richard L. Ottinger	2,171,232	37%
Total	5,893,894	100%
Conservatives (Buckley)	2,288,190	39%
Liberals (Goodell + Ottinger)	3,605,704	61%

revise beliefs." Such an environment is useful when the merits of policy alternatives react strongly to underlying fundamentals that are unknown; when this is so there is a pressing *need for direction*.

It is instructive to consider first a world without a leader; after all, leaderless institutions might also provide a learning environment. We imagine a conference of party activists assessing policies. An activist attends meetings, listens to opinions, engages in conversations, and eavesdrops on others: he sees a private signal of the party's mood, which in turn reflects the underlying merits of the competing policy platforms, and where the signal's precision is his *sense direction*. An advocacy game is played in which each activist uses a threshold rule: he backs a policy if and only if his signal sufficiently favors it. If threshold rules are used, then differences in signal realizations (activists hear different things) can generate disagreement. Our formal analysis identifies a unique stable equilibrium in which all party members use the same threshold rule. Full coordination is prevented by a negative-feedback effect, similar to that seen in models of jury voting, which means that a bias by others toward a particular policy pushes an activist away from it. Interestingly, this effect can also bias the party toward the policy with the higher barrier to coordination.

So, does a leaderless party perform well? An advantage of conference is that policy responds to the fundamentals by aggregating the opinions of party members. A disadvantage is that the private nature of opinion formation hinders the emergence of a common goal, and so the party may fail to coordinate. If policies are evenly matched, so that there is a mood of indifference, then activists' signals may well point in different directions. If activists respond by advocating different policies, then the party splits. This risk of mis-coordination might generate what Michels (1915, 35–41) described as a "need for leadership felt by the mass." A leader could perhaps dictate the policy choice and so avoid miscoordination, although her decision, based on her own signal, may not reflect the true best interests of the party. This trade-off was recognized by Madison. His reading of ancient history made it clear that, in critical circumstances, decision-making power was given to individuals. In the 38th *Federalist Paper* (Hamilton, Madison, and Jay 1788), he mused:

> Whence could it have proceeded, that a people, jealous as the Greeks were of their liberty, should so far abandon the rules of caution as to place their destiny in the hands of a single citizen? Whence could it have proceeded, that the Athenians, a people who would not suffer an army to be commanded by fewer than ten generals, and who required no other proof of danger to their liberties than the illustrious merit of a fellow citizen, should consider one illustrious citizen as a more eligible depositary of the fortunes than themselves and their posterity, than a select body of citizens, from whose common deliberations more wisdom, as well as more safety, might have been expected?

Madison's insight was that "these questions cannot be fully answered, without supposing that the fears of discord and disunion ... exceeded the apprehension of treachery or incapacity in a single individual." In our context, the "capacity of a single individual" is his sense of direction, measured relative to the need for direction, and the "fear of discord" stems from the barriers to coordination faced by the party.

To assess further the desire for leadership, we ask: if a leading activist could choose between (1) dictating policy herself, and (2) allowing the policy to emerge from a party conference, then what would she do? Turning to the feasibility of leadership, we suppose that a leader stands up and makes a speech that is heard and commonly interpreted by all. Activists now draw on both a private signal (from conference) and a public signal (from the speech). When will they ignore the former and act only on the latter?

Our answers incorporate the barriers to coordination faced by the policy alternatives, the need for policy direction, and the sense of direction enjoyed by each activist. Our analysis reveals that it is natural to combine these three elements into a single index, labeled R, of the desirability and feasibility of leadership:

$$R = \frac{\text{Barriers to Coordination} \times \text{Sense of Direction}}{\text{Need for Direction}}.$$

We offer a brief preview of our results here. We find that a leader would ideally wish to bias the threshold rule used by others toward her own opinion; the size of this bias increases with R. Furthermore, the leader always prefers to dictate policy if and only if $R > 1$. This inequality also determines the feasibility of leadership: if the leader gives a public speech then party members discard their private views and rely entirely on the public recommendation of the leader if and only if $R > 1$. When $R < 1$, then only "extremist" activists, with signals that point decisively toward one of the policies, retain a desire to dictate; the "moderates" with equivocal assessments would rather refrain from such decisive leadership.

Extending our model, we also offer an analysis of oligarchy; this leads us to refer to the leadership index R as *Michels' Ratio*. We explore whether the party may prefer leadership by an elite subset of its membership. Specifically, we suppose that a k-strong elite can share their assessments of the fundamentals and so reach a common view, and that a representative is able to communicate perfectly this view. The criterion for leadership by the elite is $kR > 1$, and so $1/R$ is the minimum size of a Michelsian oligarchy.

RELATED LITERATURE

Before proceeding to construct the formal model of activism and leadership that underpins our arguments, we first review a selection of three themes from the literature on which we build.

Party activists play a "global game" in the sense of Morris and Shin (2003); a game "of incomplete information whose type space is determined by the players each observing a noisy signal of the underlying state." In our model, the "noisy signal" includes information gleaned from a party conference (a private signal) or from a leader's speech (a public signal).[2] Economists have used global games to model many phenomena, including currency crises (Morris and Shin 1998), bank runs (Goldstein and Pauzner 2005), and debt pricing (Morris and Shin 2004). However, the approach has seen little use in political science. Some of the papers in this literature are concerned with the welfare effects of public-information dissemination in the context of coordination games. In our world, the emergence of a leader corresponds to the introduction of public information. For this reason, our paper relates to the work of Morris and Shin (2002) and subsequent debates (Angeletos and Pavan 2004; Morris, Shin, and Tong 2006; Sevensson 2006).

Our work also builds upon Myatt's (2007) analysis of strategic voting in plurality-rule elections. His stable voting equilibrium with multicandidate support contrasts with the "Duvergerian" equilibria of previous studies (Cox 1994, 1997; Myerson and Weber 1993; Palfrey 1989). Unlike those earlier models voters do not enjoy perfect knowledge of the electoral situation and instead base their decisions on private signals. Because signals may point in different directions all candidates receive votes in equilibrium; supporters of two candidates who share a dislike of a third candidate fail to coordinate fully. We comment further on the relationship of our work with strategic-voting problems in the penultimate section of the paper.

Finally, some of our insights stem from the literature on information aggregation. The negative-feedback effect that features in some of our analysis was also central to the work of Feddersen (1992) and Feddersen and Pesendorfer (1997, 1998). For example, Feddersen and Pesendorfer (1998) showed that the unanimity of jury verdicts can increase the likelihood of a conviction irrespective of the true state of the world: when a juror is pivotal in securing a conviction under unanimity, he will put less emphasis on his own private information than on what must be true, given that others are making informed decisions.

COORDINATING ACTIVISM

The setting for our analysis is a simultaneous-move, binary-action, common-interest game in which party activists must decide which of two policy platforms to advocate. Because the players share common payoffs, the game lacks the tension between private and social interests that is central to the view of collective action popularized by Olson (1965). Nevertheless, the actors face a collective-action problem in a different sense: a critical mass must coordinate behind one of the policies if that platform is to be adopted successfully. Failure to coordinate results in a party split and subsequent electoral failure.

Formally, n party activists simultaneously decide to support either platform A or platform B. We imagine n to be large, so that the player set is either the entire party membership or a large and representative subset of them. We write x for the number of activists who advocate platform A, so that the other $n-x$ activists back B. There are three possible outcomes: (1) policy A succeeds; (2) policy B succeeds; or (3) the party is split and fails to move forward. For two fractions p_A and p_B satisfying $1 > p_A > p_B > 0$,

$$\text{Outcome} = \begin{cases} \text{Platform } A & \text{if } \frac{x}{n} > p_A, \\ \text{Failure} & \text{if } p_A \geq \frac{x}{n} \geq p_B, \text{ and} \\ \text{Platform } B & \text{if } p_B > \frac{x}{n}. \end{cases}$$

Thus p_A is the coordination required for the success of A, and $1-p_B$ is the coordination required for B. We think of situations in which a 50:50 party split leads to failure; this corresponds to $p_A > \frac{1}{2} > p_B$.[3] Our stylized description of mid-1990s British Labour Party reform illustrates: platform A is the "New Labour" agenda, and B represents "old" Labour. Given that B is some kind of status quo, we might specify $p_A > 1-p_B$; contentious aspects of the new agenda require greater unity. If frustration with old failings means that greater consensus is needed to retain existing policies then we might set $p_A < 1-p_B$.

As noted in our introduction, we can envisage a secondary interpretation of the model developed here. For the 1970 New York senatorial election (Table 1) the parameter $n = 3{,}605{,}704$ is the number of liberal voters. To defeat the conservative, a critical super-majority exceeding 2,288,190 needed to coordinate behind either Goodell or Ottinger. Dividing by n then yields the specification $p_A = 1-p_B \approx 63.5\%$.

We turn to payoffs. A failure to coordinate is undesirable, and so generates a zero payoff to everyone, whereas successful coordination yields a strictly positive payoff.[4] For some $u_A > 0$ and $u_B > 0$,

$$\text{Common Payoff} = \begin{cases} u_A & \text{if } \frac{x}{n} > p_A, \\ 0 & \text{if } p_A \geq \frac{x}{n} \geq p_B, \text{ and} \\ u_B & \text{if } p_B > \frac{x}{n}. \end{cases}$$

[2] This definition of a global game (which, taken literally, defines only a game of incomplete information) differs from the original use of the term: Carlsson and van Damme (1993) considered a specification in which each strategy is strictly dominant in a particular region of the state space.

[3] Our formal results continue to hold when either $1 > p_A > p_B > \frac{1}{2}$ or $\frac{1}{2} > p_A > p_B > 0$.

[4] Each activist must choose to back one of the policies; indifferent abstention is disallowed. However, given that mis-coordination is possible and everyone strictly dislikes failure, abstention from participation is a (weakly) dominated strategy. This contrasts with the work of Feddersen and Pesendorfer (1996). They studied a world in which either A or B always wins, and they found a strict incentive for abstention.

FIGURE 1. Coordinating Activism

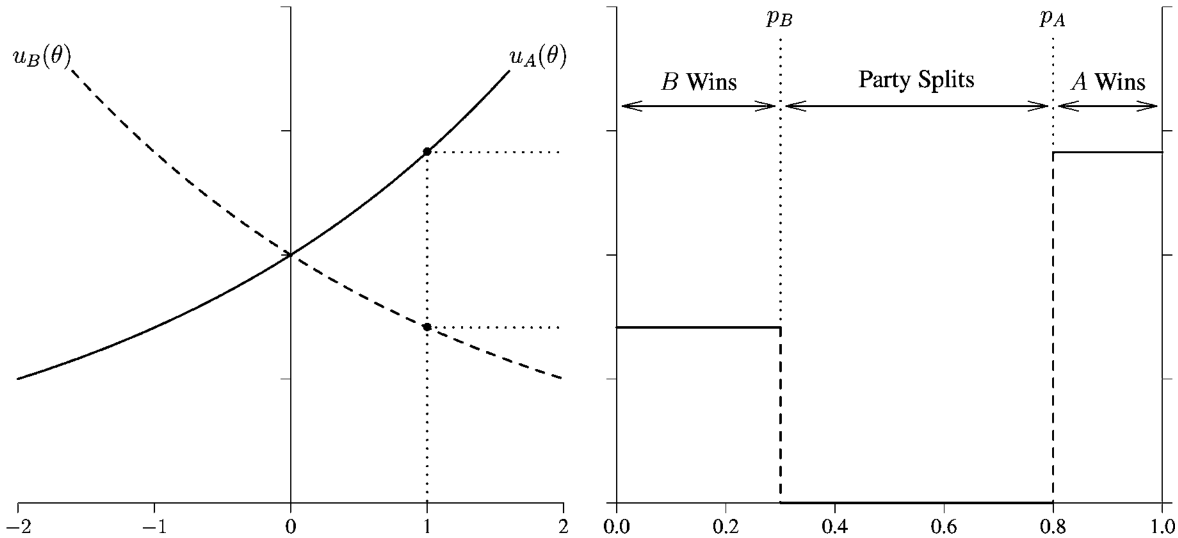

(a) From Fundamentals to Policy Payoffs

(b) The Coordination Problem

Note: Panel (a) illustrates the relationship between the underlying political situation and the payoffs from successful coordination. For the particular value of $\theta = 1$ highlighted, we see that $\frac{u_A(\theta)}{u_B(\theta)} = 2$, and so (if only they knew the true underlying state of the world) activists would wish to coordinate on platform A. Panel (b) illustrates the barriers to coordination faced by activists. In this case, $p_A > 1 - p_B$; hence policy platform A might be described as a "loftier" goal. When the proportion of activists supporting A falls into the central region, the party splits and fails.

All party members would like to coordinate on the best policy. When information is complete there are equilibria with full coordination, so that either $x = 0$ or $x = n$. There are many other pure-strategy equilibria in which the party either partially coordinates or splits and fails.[5] Despite this multiplicity, one equilibrium seems focal: if $u_A > u_B$, then surely everyone should coordinate on policy A? Alas, this obvious solution assumes that everyone shares a common understanding of the policies' merits.

We build a richer game where the desirabilities of the policies are uncertain. Formally, the payoffs $u_A(\theta)$ and $u_B(\theta)$ depend on an underlying (and *ex ante* uncertain) real-valued state variable θ. We assume that $u_A(\theta)$ is increasing, and $u_B(\theta)$ is decreasing; an increase in θ favors A relative to B. The state variable θ represents the underlying political situation. It might depend on socioeconomic variables, the preferences of an electorate, or the ideology of the party membership at large. We use a "need for direction" parameter $\lambda > 0$ to index the the extent to which payoffs respond to these fundamentals, and adopt the functional forms $u_A(\theta) = \exp[\frac{\lambda\theta}{2}]$ and $u_B(\theta) = \exp[-\frac{\lambda\theta}{2}]$. Thus an activist's preference for policy A relative to policy B is a log-linear function of the fundamental underlying state-of-the-world θ:

$$\log\left[\frac{u_A}{u_B}\right] = \lambda \times \theta \quad \text{where} \quad \lambda = \textit{need for direction}.$$

Figure 1(a) illustrates our specification. When θ is zero an activist is indifferent between the two policies. As θ swings to the right then the payoff from platform A relative to platform B grows.

Since θ is unknown, activists must use any informative signals at their disposal to form beliefs about it. Choices relying on such signals depend on the realization of θ. An activist must contemplate the underlying fundamentals for two reasons: firstly, he assesses the merits of the policies; and, secondly, he considers the likelihood of pivotal events in which his advocacy makes a difference. We will specify the signals via which activists form their assessments, but before doing so we pause to discuss optimal behavior.

An activist's choice is relevant only when he is *pivotal*. For instance, he is pivotal for policy A when support for A is one step short of p_A. $\mathcal{P}_A$ is shorthand for this event. Similarly, $\mathcal{P}_B$ is the situation in which the activist is pivotal for policy B. If choices respond to signal realizations, then these two pivotal events have positive probability. If an activist is pivotal for A, his support for A yields a payoff of $u_A(\theta)$. Since θ is uncertain his expected payoff is $\mathrm{E}[u_A(\theta) \mid \mathcal{P}_A]$; the merits

[5] Any pattern of support where no individual is pivotal is an equilibrium. The two pure-strategy profiles which are not equilibria correspond to values of x satisfying either (i) $\frac{x+1}{n} > p_A \geq \frac{x}{n}$ or (ii) $\frac{x}{n} \geq p_B > \frac{x-1}{n}$. For instance, in the first situation a player who chooses B could switch to A, avoiding a failure and hence yielding a payoff gain of u_A. Such pivotal situations are exceptional, because other values of x (where no individual can affect the outcome) yield a pure-strategy Nash equilibrium. We could, of course, characterize mixed-strategy equilibria but this would only expand the embarrassment of riches.

of policy A are conditional on the pivotal event $\mathcal{P}_A$. This event happens with probability $\Pr[\mathcal{P}_A]$, and so the net payoff from backing A is $\Pr[\mathcal{P}_A] \times \mathrm{E}[u_A(\theta) \mid \mathcal{P}_A]$. Similarly, the net payoff from backing B is $\Pr[\mathcal{P}_B] \times \mathrm{E}[u_B(\theta) \mid \mathcal{P}_B]$. It is strictly optimal to support A, if and only if $\Pr[\mathcal{P}_A]\mathrm{E}[u_A(\theta) \mid \mathcal{P}_A] > \Pr[\mathcal{P}_B]\mathrm{E}[u_B(\theta) \mid \mathcal{P}_B]$, or equivalently

$$\text{Choose } A \iff \underbrace{\log \frac{\Pr[\mathcal{P}_A]}{\Pr[\mathcal{P}_B]}}_{\text{(i) relative likelihood}} + \underbrace{\log \frac{\mathrm{E}[u_A(\theta) \mid \mathcal{P}_A]}{\mathrm{E}[u_B(\theta) \mid \mathcal{P}_B]}}_{\text{(ii) relative conditional preference}} > 0. \qquad (\star)$$

Hence the activist optimally balances (i) the relative likelihood of his pivotality for A versus B; and (ii) his relative preference evaluated conditional on his pivotality in each of the two cases.

THE PARTY CONFERENCE

Here we study behavior in a leaderless world. A "party conference" is a metaphor for intraparty discussion of the underlying political situation. We think of a gathering where each activist gains (via meetings, conversations, and eavesdropping) an informative but private signal of the true state of the world. If such a conference is representative of the wider party membership, then signals are correct on average—they correspond to the true value of θ. Although the aggregate conference mood perfectly reflects θ, different activists form different assessments: an activist who attends fringe meetings may get a different sense of the party mood than one who spends his time in the conference bar. This is important because, as an activist wanders around conference, he wonders whether others share his sense of the direction in which the party is going.[6]

We think of conference as providing a learning environment. Of course, the party might instigate a formal mechanism to aggregate opinions; for instance, a simple vote among the party caucus. Alas such a mechanism does not provide a resolution to the coordination problem. For example, a 50:50 split in the vote share would publicly reveal any underlying division; the party would be seen to be split and to have failed in its quest to coordinate. Essentially, then, voting and advocacy are equivalent. To circumvent this problem, we might think of activists submitting their opinions in private to a mediator. We return to this issue in the penultimate section of the paper, but for now focus attention on our policy-advocacy game.

Activists begin with no knowledge of the best policy to pursue; formally they entertain uniform (improper) priors over θ. Each activist $i \in \{1, 2, \ldots, n\}$ then sees a private signal, where conditional on the party mood (henceforth we sometimes refer to θ thus) signals are identically and independently distributed:[7]

$$m_i \mid \theta \sim N\left(\theta, \frac{1}{\psi}\right) \quad \text{where} \quad \psi = \textit{sense of direction}.$$

Here ψ, the inverse of the variance, is the signal's precision; it represents the "sense of direction" of an activist. Conditional on his signal, his updated beliefs satisfy $\theta \mid m_i \sim N(m_i, [1/\psi])$.

We consider a natural class of strategies: an activist operates a *threshold rule* if he backs A rather than B if and only if his assessment of the party mood exceeds some threshold m.

Definition. Activists employ a *threshold rule* if, for some m, each activist chooses to advocate policy A if and only if $m_i > m$. When m takes on a finite real value, a threshold rule is *signal responsive*.

Allowing the threshold to take values $m \in \{-\infty, \infty\}$, the class of threshold rules includes those where an activist ignores his private signal and always advocates A (corresponding to $m = -\infty$) or always advocates B (for $m = \infty$). However, for any finite threshold, actions respond to signals in a nontrivial way. Conditional on the state of the world θ, an activist backs A with probability $p \equiv \Pr[m_i > m \mid \theta]$. Writing $\Phi(\cdot)$ for the cumulative distribution function of the standard normal, $p = \Phi[\sqrt{\psi}(\theta - m)]$. Clearly, support for A grows as the party's mood swings to the right (θ increases). It is also decreasing in the threshold used by activists: a reduction in m yields a move toward A and away from platform B.

Whereas we call p the "support" for platform A, the actual fraction of activists who advocate A is $\frac{x}{n}$. However, if n is large then the Law of Large Numbers ensures that this fraction is almost always close to p. Given that this is so, policy A succeeds if and only if $p > p_A$. Equivalently, it succeeds if and only if the true party mood exceeds θ_A, where θ_A satisfies $p_A = \Phi[\sqrt{\psi}(\theta_A - m)]$; similarly, B succeeds if and only if $\theta < \theta_B$ where $p_B = \Phi[\sqrt{\psi}(\theta_B - m)]$. Summarizing,

$$\text{Outcome} = \begin{cases} \text{Platform } A & \text{if } \theta > \theta_A, \\ \text{Failure} & \text{if } \theta_A \geq \theta \geq \theta_B, \\ \text{Platform } B & \text{if } \theta_B > \theta, \end{cases}$$

$$\text{where} \quad \begin{cases} \theta_A = m + \frac{\pi_A}{\sqrt{\psi}}, & \text{and} \\ \theta_B = m - \frac{\pi_B}{\sqrt{\psi}}, \end{cases}$$

where we have used the substitution $\pi_A \equiv \Phi^{-1}(p_A)$ and $\pi_B \equiv \Phi^{-1}(1 - p_B)$. The parameters π_A and π_B measure the height of the barriers to coordination faced by the party. The two critical values, θ_A and θ_B, partition the range of party moods into three segments according to the outcome.

We make two observations. First, when activists use a threshold m they react to their assessments and so

[6] Our results extend to a setting in which the aggregate mood of the party conference imperfectly reflects the wider political situation and so fails to capture fully the merits of the policy alternatives.

[7] Our results extend easily to environments in which the activists' signals are conditionally correlated.

policy responds to the fundamentals. Second, such strategies bring the risk of mis-coordination: when $\theta_A > \theta > \theta_B$ the party's mood (relative to the threshold m) is one of indifference and the party fails. The size of this "zone of mis-coordination" is determined by $(\pi_A + \pi_B)/\sqrt{\psi}$ and so is increasing with the combined height of the barriers to coordination but is decreasing in activists' sense of direction.

We now return to consider an individual's decision. He uses his signal to assess the underlying party mood and hence (1) the relative likelihood of his pivotality for A versus B, and (2) his relative preference for A versus B. We study each factor in turn, before characterizing behavior in equilibrium.

When the party conference is large, and others adopt a threshold m, the support p is determined by θ. An activist realizes he is pivotal for A only when $\theta \approx \theta_A$, and for B only when $\theta \approx \theta_B$. Thus, conditional on his private signal m_i, activist i computes the relative likelihood of θ_A versus θ_B. His posterior beliefs are normally distributed around m_i with precision ψ and so, as the following lemma confirms, the log relative likelihood takes a simple form. (All proofs are contained in the technical appendix.)

Lemma 1. *Fixing a threshold m used by other members of the party, and conditional on the private signal m_i of activist i, the log relative likelihood of being pivotal for A versus B satisfies*

$$\log \frac{\Pr[\mathcal{P}_A]}{\Pr[\mathcal{P}_B]} \longrightarrow (\pi_A + \pi_B)\sqrt{\psi} \times (m_i - m) \quad \text{as } n \to \infty.$$

This is increasing in the activist's private assessment of the true relative merits of the policies, and decreasing in the threshold used by others. The relative-likelihood effect increases with the barriers to coordination $\pi_A + \pi_B$ and an activist's sense of direction ψ, but is unaffected by the need for direction λ.

When thinking of pivotal events, an activist judges the party's mood relative to the threshold used by others. As he perceives a swing to the right (a rise in θ) it is more likely that, if he is pivotal, he will be pivotal for A. The effect depends not only on his sense of direction, but also on the height of the barriers to coordination; as π_A and π_B grow the zone of mis-coordination widens (Figure 2). The two critical values θ_A and θ_B move further apart and so are easier to distinguish. This being so, an activist can more readily ascertain which of the two critical events (whether the is pivotal for A or for B) is more likely.[8]

Lemma 1 also reveals a positive-feedback effect. As others bias toward A (a fall in m) the log-likelihood ratio grows, pushing activist i toward the advocacy of A. So, if an activist believes that others lean toward a policy, then he is tempted to follow them. This is not the whole story, however, because we must also consider the activist's relative preference for A versus B. When contemplating the policies' merits, an activist considers his strategic environment and recognizes that his payoff only matters when his action is critical to a policy's success. When the party is large he is pivotal for A only when $\theta \approx \theta_A$, and hence $\mathrm{E}[u_A(\theta) \mid \mathcal{P}_A] \longrightarrow u(\theta_A)$ as $n \to \infty$. Similar considerations for the payoff from policy B lead to the following lemma.

Lemma 2. *Fixing a threshold m, the log relative conditional preference for A over B satisfies*

$$\log \frac{\mathrm{E}[u_A(\theta) \mid \mathcal{P}_A]}{\mathrm{E}[u_B(\theta) \mid \mathcal{P}_B]} \longrightarrow \lambda \left[m + \frac{\pi_A - \pi_B}{2\sqrt{\psi}} \right] \quad \text{as } n \to \infty.$$

This relative preference for A increases with the bias of others toward B and with the relative height $\pi_A - \pi_B$ of the barriers to coordination. The size of the second effect decreases with an activist's sense of direction ψ, and the size of both effects increases with the overall need for direction λ.

Lemma 2 reveals a negative-feedback effect: as others lean toward A (a fall in the threshold m) then the relative preference for A falls, which (cf. Lemma 1) pushes an activist toward B. Similarly, when A is easy to achieve, in the sense of a lower barrier to coordination, an activist swings against it and toward B.

Why does an activist swing toward B in a situation where the odds are stacked against it? He asks himself what must be true of the world given the situation he finds himself in (namely, that he is pivotal) and so the realization of his own signal does not matter. When he is pivotal for B so that B is on the verge of success then, because other factors work against B, the party mood must heavily favor it and so policy B must be extremely desirable. The implications of such considerations depend on the strength of the connection between the policies' payoffs and the underlying state of the world; this is simply the need for direction λ.

Lemmas 1 and 2 characterize the effect of a threshold strategy deployed by the party membership on the incentives of an individual. We now seek equilibria within the class of threshold strategies.

In general, (Bayesian Nash) equilibria depend on the precise size n of the party's membership. Our aim, however, is to characterize equilibria in large parties. Put somewhat informally, we seek an equilibrium threshold advocacy strategy that "works for large n." One obvious approach would be to find equilibria (if they exist) for each n and examine the properties of the associated sequence of equilibria (assuming that the sequence converges) as $n \to \infty$. Here, however, we follow previous work (Myatt 2007) by defining a simpler and arguably natural solution concept directly over a sequence of policy-advocacy games.

Definition. A threshold equilibrium *is a threshold rule with threshold m^* such that when all other party members use it: (1) an activist $m_i > m^*$ optimally backs policy A when the party size is sufficiently large; and (2)*

[8] The need for direction does not enter Lemma 1. This is because the activist considers only the likelihood of $\mathcal{P}_A$ and $\mathcal{P}_B$; the need for direction comes into play only when an activist considers the policies' merits.

FIGURE 2. Critical Values of the State of the World

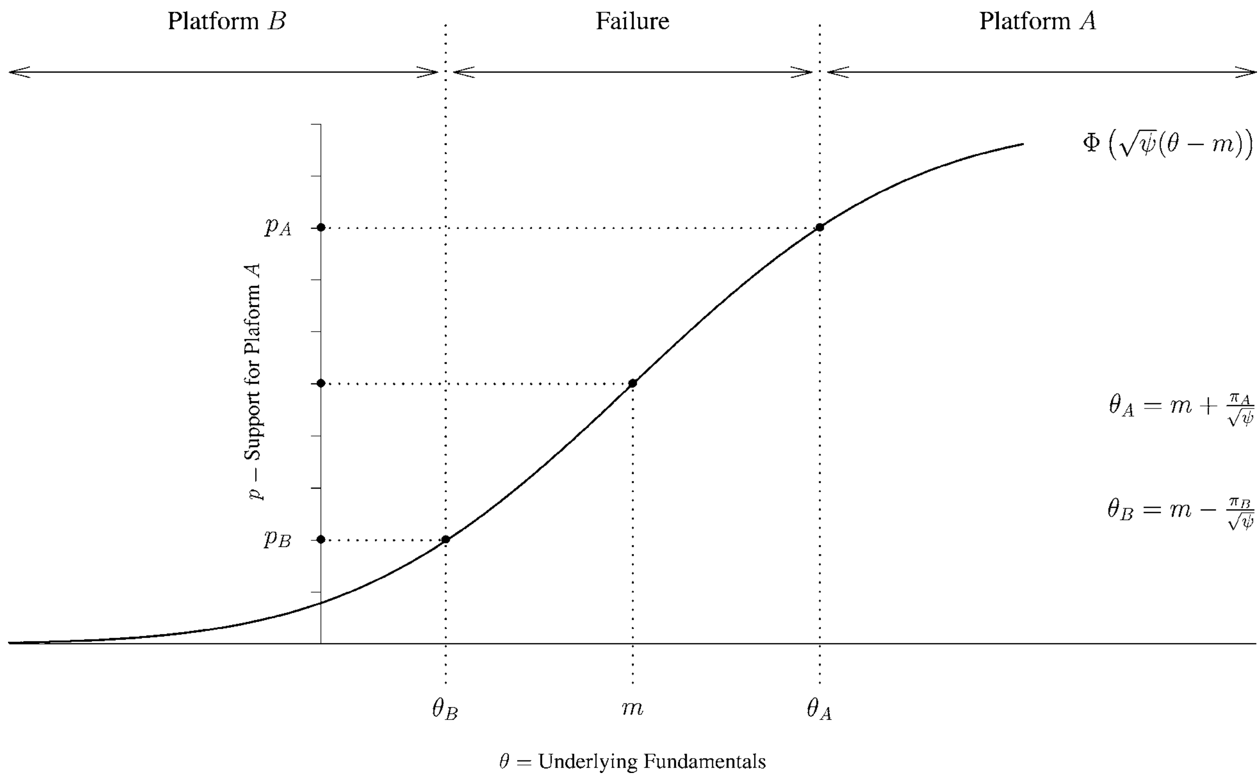

Note: This figure illustrates the relationship between the underlying fundamentals θ and the support p for policy A (the support for B is $1 - p$) when activists use the threshold rule:

$$\text{Choose } A \iff m_i > m.$$

When $\theta = m$, the party is 50:50 split between the platforms. For A to succeed, the party mood needs to swing right to $\theta > \theta_A$, and for B to succeed the mood must swing left to $\theta < \theta_B$. The mood is indecisive when $\theta_A < \theta < \theta_B$. Observe that

$$\theta_A - \theta_B = \frac{\pi_A + \pi_B}{\sqrt{\psi}},$$

and so the "zone of mis-coordination" for an indecisive party mood increases with the barriers to coordination but decreases with activists' sense of direction.

an activist $m_i < m^*$ optimally backs policy B when the party size is sufficiently large.

Heuristically, this defines an "ε-equilibrium" for large parties and arbitrarily small ε. Hence, an equilibrium threshold rule leads to the play of a best reply by (almost) all activists in parties that are sufficiently large. In one sense, this is less stringent than the Bayesian-Nash solution concept: there may be some activists (with signals very close to m^*) who do not play a best reply for a particular finite n. In a second sense, our concept is more stringent: whereas a Bayesian-Nash equilibrium would involve the play of a best reply by almost all activists for a particular finite n, it would not necessarily be robust to increases in n.[9]

Adopting our solution concept, two equilibria, $m^* = \infty$ and $m^* = -\infty$, involve full coordination: the pivotal events $\mathcal{P}_A$ and $\mathcal{P}_B$ do not occur, and actions are (trivially) optimal. In contrast, when an equilibrium threshold m^* takes a finite value, decisions respond to signal realizations. For a signal-responsive equilibrium, the outcome is uncertain and the events $\mathcal{P}_A$ and $\mathcal{P}_B$ occur with positive probability. A threshold m^* forms a signal-responsive equilibrium when an activist with a signal $m_i = m^*$ is essentially indifferent between the alternatives for large n. Assembling the relative-likelihood and relative-preference effects,

$$m_i = m^* \quad \Longrightarrow \quad \lim_{n\to\infty}\left[\log\frac{\Pr[\mathcal{P}_A]}{\Pr[\mathcal{P}_B]} + \log\frac{\mathrm{E}[u_A(\theta)\mid\mathcal{P}_A]}{\mathrm{E}[u_B(\theta)\mid\mathcal{P}_B]}\right] = 0.$$

[9] Myatt (2007) offered further discussion of our solution concept. Roughly speaking, if a sequence of threshold-based Bayesian-Nash equilibria converges, then the limiting threshold yield a threshold equilibrium in the sense used in this paper. If it does not converge, then it must mean that the Bayesian-Nash equilibria depend sensitively on the precise size n of the party's membership. In this latter situation, a desire for robustness suggests the use of the solution concept used here.

Inspecting Lemma 1, when $m_i = m^*$ the relative-likelihood effect disappears. Given that a private signal is equal to the threshold used by others, $\mathcal{P}_A$ and $\mathcal{P}_B$ are equally likely in an arbitrarily large party: formally, $\log[\Pr[\mathcal{P}_A]/\Pr[\mathcal{P}_B]] \to 0$ as $n \to \infty$. This means that the equilibrium threshold m^* is tied down by the relative preference term from Lemma 2, which must equal zero. This term is increasing in the difference between the barriers to coordination π_A and π_B, so that there is a preference bias toward policy A whenever A is more difficult to achieve (i.e., $\pi_A > \pi_B$). The equilibrium threshold m^* must, therefore, offset this effect. Setting the relative preference to zero, and solving for m^* we obtain the next result.

Proposition 1. *There is a unique signal-responsive threshold equilibrium, with threshold*

$$m^* = \frac{\pi_B - \pi_A}{2\sqrt{\psi}}.$$

There is a bias toward policy A if and only if its barrier to coordination is higher than that of policy B. The size of the bias decreases with the sense of direction ψ, and is invariant to the need for direction λ.

The signal-responsive equilibrium (where activists pay attention to their signals) seems, at face-value, more plausible than its contenders. Nevertheless it has surprising properties. Whilst one might think an exogenous increase in π_A makes a success for policy A less likely, in equilibrium activists are biased in their advocacy toward a policy that is more difficult to achieve: there is an endogenous reaction to any increase in π_A. The effect we are pointing to has been seen elsewhere: the negative feedback that drives this result is present in work by Feddersen and Pesendorfer (1996, 1997, 1998). Furthermore, the effect arguably accords with a stylized fact related to our motivating example of the British Labour Party: during periods when its conference has asserted its sovereignty (during the 1980s in particular), the party swung to the left and toward policies less favored by the majority of voters. The full effect of the reaction to changes in the barriers to coordination can be seen by substituting for m^* in our expressions for θ_A and θ_B:

$$\theta_A = \frac{\pi_A + \pi_B}{2\sqrt{\psi}} \quad \text{and} \quad \theta_B = -\frac{\pi_A + \pi_B}{2\sqrt{\psi}}.$$

The zone of mis-coordination is a symmetric interval around zero. If $\pi_A > \pi_B$, for instance, then the endogenous equilibrium bias toward A precisely offsets the exogenous bias toward B. Turning attention back to Figure 2, the political outcome (A, B, or failure) depends only on the location of the zone of mis-coordination which, in turn, depends only on the aggregate height of the barriers to coordination. If we were to raise the barrier faced by one policy, while lowering the other, then this zone would not move. The important point to be made is that a key feature of the institutional backdrop, namely, that some policies require greater consensus, does not influence the policy adopted by conference.

We note two further features. First, a signal-responsive threshold ensures that policy choices react to the party mood (which in turn reflects the underlying fundamentals) and so favors the use of a conference as an institutional mechanism for aggregating opinions. Second, there is a risk of mis-coordination: the party fails when $\theta_A > \theta > \theta_B$. Note, however, that the relative importance of these effects depends on the need for direction and yet, despite this, λ plays no role in the equilibrium threshold strategy.

Before assessing the performance of the conference we highlight the feedback effects that are central to the incentives of activists and justify our focus on the signal-responsive equilibrium described in Proposition 1. An activist's advocacy will depend on the threshold rule used by others. As m increases, both his relative likelihood of influencing the party's chosen policy and his relative preference for the policies are affected. Inspecting Lemma 1, notice that as m rises (so that others bias toward B), the relative likelihood falls at rate $(\pi_A + \pi_B) \times \sqrt{\psi}$, so that he too is led toward policy B; this is *positive feedback*. On the other hand, Lemma 2 reveals that, as others bias toward A, his relative preference for A versus B rises at a rate λ; this is *negative feedback*. Combining these two effects, we see that positive feedback exceeds negative feedback if and only if $(\pi_A + \pi_B) \times \sqrt{\psi} > \lambda$. When this holds, the net effect of a bias among other members is to push an activist in the same direction. We write this criterion in a modified form.

Definition. *We define R to be a single index combining the aggregate height $\pi_A + \pi_B$ of the barriers to coordination, the need for direction λ, and an activist's sense of direction ψ. It satisfies*

$$R \equiv \frac{(\pi_A + \pi_B) \times \sqrt{\psi}}{\lambda}.$$

Positive feedback exceeds negative feedback if and only if $R > 1$. We use "Michels' Ratio" to refer to R.

As we shall see, the index R is fundamental to both the feasibility and the desirability of leadership. When $R > 1$, it seems that a "bandwagon" might form behind a policy. To explore this hypothesis, consider a situation where $\pi_A = \pi_B$, so that $m^* = 0$. Now suppose that party members switch to use a lower threshold strategy $m < 0$, so that they are biased toward A. When $R > 1$, this increases the incentive for an activist to back A; in particular, an activist who observes a message $m_i = 0$ will certainly advocate policy A.

Although this may suggest that a bandwagon begins to roll in favor of A, this is not quite so. To see why, consider an activist who observes a message $m_i = m$. For such a party member, the relative likelihood term disappears and thus the only factor he will consider is his relative preference term. Of course, negative feedback is integral to this term and ensures that he now faces a strict incentive to back B. The attempt to induce a bandwagon effect in favor of policy A fails, because an activist with a signal equal to this new threshold will strictly wish to back B. This discussion leads to a formal definition of stability.

Definition. *A threshold equilibrium m^* is stable if: (1) for larger thresholds $m > m^*$ an activist with a private signal $m_i = m$ would strictly prefer to back policy A; and (2) for smaller thresholds $m < m^*$ an activist with a private signal $m_i = m$ would strictly prefer to back policy B.*

This definition allows us to assess formally the stability of the different equilibria.

Proposition 2. *In a conference environment, the signal-responsive equilibrium (Proposition 1) is stable. The fully coordinated equilibria in which activists ignore their signals are unstable.*

This provides one justification (there are others) for a focus on signal-responsive equilibria in a conference environment.[10] The important observation is that the relative-likelihood term depends on the difference $m_i - m$ between that threshold and the message received by an activist, and so the positive-feedback effect is eliminated when we consider an activist whose signal is equal to the threshold used by others. Hence, in equilibrium, negative feedback dominates. Interestingly, the positive-feedback term returns to play a significant role once we introduce the possibility (later in the paper) of a publicly observed leader's speech.

THE NEED FOR LEADERSHIP

Michels (1915, 35–41) suggested that the transition from democracy to oligarchy is driven by a "need for leadership felt by the mass." We analyze this need by assessing the desire of an individual to abandon conference and dictate the party's policy. We do this via a thought experiment that eventually yields Proposition 3 below. We ask activist l (we call her a leader) to choose the threshold used by others. She has three options: (1) $m = -\infty$, so that everyone backs A; (2) $m = \infty$, so that everyone backs B; or (3) some interior threshold. Options (1) or (2) reflect a desire to dictate. Option (3) exploits the information-aggregation properties of a conference but risks a coordination failure; it reflects a desire to shape policy, but not to dictate it. Conducting our experiment, we write $U(m, m_l)$ for the expected payoff enjoyed by leading activist l in a large party ($n \to \infty$) when she sees a signal m_l and others use a threshold m. Hence,

$$U(m, m_l) = \Pr[\theta > \theta_A] \times \mathrm{E}[u_A(\theta) \mid \theta > \theta_A] + \Pr[\theta < \theta_B] \times \mathrm{E}[u_B(\theta) \mid \theta < \theta_B],$$

where the probabilities and expectations are conditional on m_l. Consider a small change in the threshold. A fall in m shifts the zone of mis-coordination to the left, moving the party toward A. Formally,

$$\frac{\partial U}{\partial m} < 0 \Leftrightarrow \underbrace{\log \frac{\Pr[\mathcal{P}_A]}{\Pr[\mathcal{P}_B]}}_{\text{(i) relative likelihood}} + \underbrace{\log \frac{\mathrm{E}[u_A(\theta) \mid \mathcal{P}_A]}{\mathrm{E}[u_B(\theta) \mid \mathcal{P}_B]}}_{\text{(ii) relative conditional preference}} + \underbrace{\frac{\pi_B^2 - \pi_A^2}{2}}_{\text{(iii) feasibility effect}} > 0, \qquad (\dagger)$$

where (i) and (ii) are evaluated at their limiting values as $n \to \infty$. These first two terms are identical to those that influence advocacy decisions. This is because the change in m has an effect only when activists are pivotal ($\theta \approx \theta_A$ or $\theta \approx \theta_B$). An increase in the relative likelihood of $\mathcal{P}_A$ versus $\mathcal{P}_B$ generates a desire to reduce m, thus increasing the likelihood of success for A. The relative preference for A is greater when there is a bias toward B: in selecting a threshold to be used by others negative feedback remains.

The third term on the right-hand side of (†) is absent from the advocacy criterion of ($\star$). It measures the difference in the heights of the barriers to coordination. Other activists (if left to their own devices) favor the ambitious policy objective. Other things equal, activist l wishes to offset this bias; if $\pi_B > \pi_A$ then the leader would prefer to lower the threshold m. To see why, note that the number of activists who switch sides following a change in m depends on the party's underlying mood. Those who switch have signals close to m. If $|m - \theta|$ is large then there are relatively few "marginal" activists since they are far from the median signal of θ; Figure 3 illustrates this. From the definitions of θ_A and θ_B we obtain $|m - \theta_A| = \frac{\pi_A}{\sqrt{\psi}}$ and $|m - \theta_B| = \frac{\pi_B}{\sqrt{\psi}}$. Suppose that $\pi_A < \pi_B$, so that the barrier to coordination on policy A is lower. When $\theta = \theta_A$, there are relatively many marginal activists, whereas when $\theta = \theta_B$ there are relatively few. Because more activists switch in the former case the pivotal event $\mathcal{P}_A$ becomes more important than the event $\mathcal{P}_B$. This provides an enhanced incentive to push the threshold m down and bias party members toward A.

In sum, delegating the choice of threshold to a leading activist, she will (other things equal) bias its level toward the policy that is "more feasible" in that it faces a lower barrier to coordination; her intervention acts as a counterweight to the endogenous bias toward an ambitious goal. However, given the presence of the term (i) in the criterion (†), she also moves the threshold in the direction of her own private signal. For instance, if $m_l > 0$, then she wishes to depress the threshold. If so, then how low will she go?

To answer, suppose that the criterion (†) holds. A reduction in m feeds back into terms (i) and (ii): the relative likelihood of $\mathcal{P}_A$ versus $\mathcal{P}_B$ increases due to positive feedback, whereas the relative conditional preferences falls due to the negative feedback. (The feasibility effect (iii) is unaffected.) When $R > 1$, the positive effect on (i) is greater than the negative effect on (ii); because positive feedback exceeds negative feedback

[10] Another justification is that the responsive equilibrium is optimal *ex ante*. Prior to the revelation of their private signals activists would all wish to commit to the play of the signal-responsive equilibrium.

FIGURE 3. Explaining the Feasibility Effect

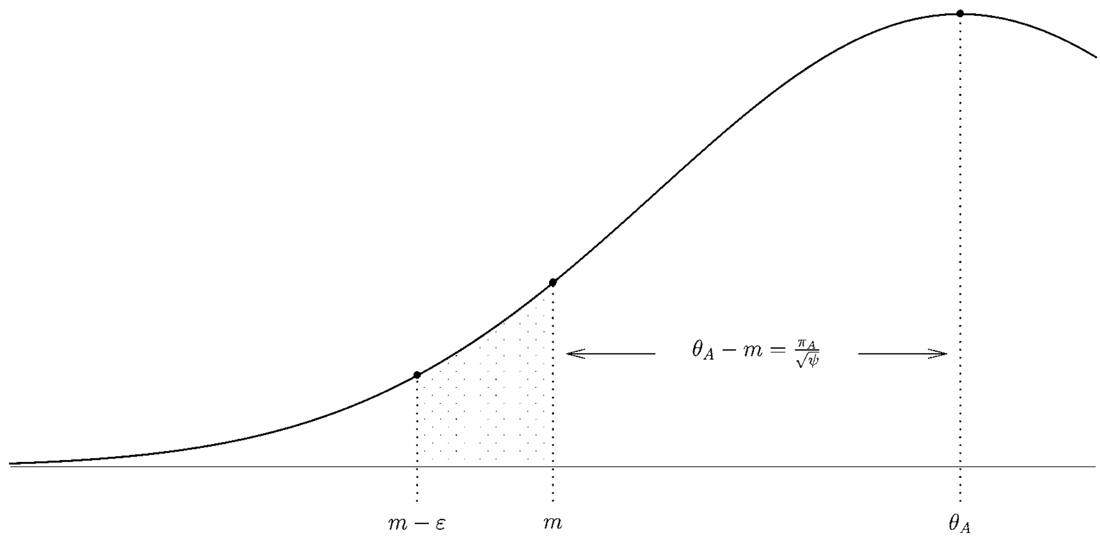

Note: This figure illustrates the effect of a reduction in the threshold m when the underlying state of the world θ hits the critical value θ_A. The solid line represents the distribution of private signals across the party membership. A reduction in the threshold from m to $m-\varepsilon$ increases the mass of activists who back A. This increase is represented by the shaded area. As θ_A and m move further apart, the effect of the ε fall in the threshold is weakened. The distance between m and θ_A is $\pi_A/\sqrt{\psi}$, and so a reduction in the threshold is less effective when the barrier to coordination for policy A is high.

the criterion remains positive and thus the leader faces a further incentive to lower m. Iterating this logic, she wishes to lower m all the way to $m = -\infty$. Similarly, if she wishes to raise m then she wishes to raise it all the way to $m = +\infty$. When $R > 1$, $U(m, m_l)$ is a quasi-convex or "U-shaped" function of m, and so the leader would prefer to choose one of the two extreme values for the threshold. She prefers to dictate the policy and sees no role for a democratic party conference.

If $R < 1$, so that negative feedback exceeds positive feedback, then a reduction in m will lower the criterion of (†). In fact, $U(m, m_l)$ is a quasi-concave or "hill-shaped" function of m. Thus, when $R < 1$, there is a unique interior threshold $m^\dagger$ that maximizes the leader's payoff. In essence, she would prefer a situation where, although she chooses the threshold, conference remains sovereign in that it determines the overall policy outcome. We summarize our thought experiment in the next proposition.

Proposition 3. *If $R > 1$ then a leader prefers to dictate the party's policy rather than allow a conference threshold strategy to operate. If $R < 1$, then she prefers others to use a threshold*

$$m^\dagger = m^* - \frac{R}{1-R} m_l$$

where m^* is the equilibrium threshold from Proposition 1. This new threshold $m^\dagger$ is decreasing in her own private signal, and so, relative to the equilibrium threshold, the leader prefers others to operate a threshold that is biased toward her own private opinion of the underlying fundamentals.

The "Michels' Ratio" index R brings together three factors driving a desire to lead. This desire arises when there are high barriers to coordination, since the zone of mis-coordination (Figure 2) is wide: ensuring that activists work together is more important than choosing the correct policy. In addition, a willingness to take up the reins of leadership is affected by an activist's sense of direction: when ψ is large, she has confidence in her signal and in her ability to lead the party in the right direction; if ψ is low, it is more prudent to exploit the information-aggregation properties of the advocacy game, although the desire to guide conference remains. Finally, when ultimate power resides with conference, policy is more likely to be in tune with the true state of the world. This is desirable when λ is large: when there is a greater need for policy direction, leaving control with the conference reduces the risk that a leader chooses to implement an inferior policy; when λ is small such risks are less salient and she retains a desire to dictate.

In our thought experiment, we asked what threshold a leading activist would choose if she could ensure others use it. Of course, other activists use the equilibrium threshold m^* rather than her preferred threshold $m^\dagger$; these thresholds coincide if and only if the leader's signal is neutral. When not, the leader may prefer to dictate even though $R < 1$. So, $R > 1$ is sufficient but not necessary to generate a desire to lead.

When $R < 1$, some are content with conference, but others, given the opportunity, would like to lead. A leader with a signal $m_l = 0$ lacks a clear signal of the best policy direction, and so is content with the information-aggregation properties of a conference; formally $m^\dagger = m^*$ in this case. She is a *moderate* with an equivocal assessment of the policies. In contrast, a leading activist with a very high or very low private signal is an *extremist* who is confident in her assessment and so feels that she is well suited to dictating party policy. This suggests a taxonomy where only the moderates favor the party conference.

Proposition 4. *If $R < 1$ then a leading activist is content with the conference equilibrium if and only if $|m_l| < \tilde{m}$, for some $\tilde{m} > 0$. Otherwise, if $|m_l| > \tilde{m}$, she would prefer to dictate by choosing the policy herself. $\tilde{m}$ rises with the need for direction λ, but falls with the aggregate height $\pi_A + \pi_B$ of the barriers to coordination and the sense of direction ψ. Changing these parameters, $\tilde{m} \to 0$ as $R \to 1$.*

Although it might seem obvious that an extremist prefers to dictate, the issue is more subtle. An extremist also believes that the party mood leans heavily in one direction and so is confident that a conference selects the best policy. This suggests that an extremist is (almost) indifferent between the two options. Yet, as the proof of Proposition 4 shows, she has relatively more confidence in herself than in the conference.

Our taxonomy of extremists and moderates, although differing somewhat from common usage of these terms, suggests a conflict within a party organization concerning the best mechanism for achieving policy coordination. This conflict coincides with differing private interpretations of the collective interest, and emerges despite the lack of inherent divergence of preferences in our model specification.

FOLLOWING THE LEADER

Having analyzed an individual's desire to take a leading role, and following Michels' (1915, 36) claim that "the renouncement of the exercise of democratic rights is voluntary," we now assess the willingness of others to follow her lead. A party member l, chosen at random, makes a stand as leader. She gives a perfectly communicated speech $s = m_l$ describing her views. Hence, in addition to their private signals, party members observe a public signal $s \mid \theta \sim N(\theta, [1/\psi])$. Activist i's updated beliefs satisfy

$$\theta \mid (s, m_i) \sim N\left(\frac{s + m_i}{2}, \frac{1}{2\psi}\right).$$

The expectation is the average of the private and public signals; the precision has doubled. Importantly, the same public signal is used by all activists, and so party members begin to share a similar perspective.

Once again, activist i supports A if and only if $m_i > m$ for some threshold m. When he considers the likelihood of the pivotal events $\mathcal{P}_A$ and $\mathcal{P}_B$, his beliefs will now be influenced by the public signal.

Lemma 3. *Fixing a threshold m used by others, and conditional on the private signal m_i of activist i and the speech s of the leader, the log relative likelihood of being pivotal for A versus B satisfies*

$$\log \frac{\Pr[\mathcal{P}_A]}{\Pr[\mathcal{P}_B]} \longrightarrow \frac{\pi_B^2 - \pi_A^2}{2} + 2(\pi_A + \pi_B)\sqrt{\psi} \times \left(\frac{m_i + s}{2} - m\right) \quad \text{as } n \to \infty.$$

This is increasing in the activist's private assessment of the merits of policy A and the leader's public speech expressing her views about the fundamentals. It is decreasing in the threshold used by others.

In the absence of the leader's intervention a marginal activist (with a signal $m_i = m$) is unconcerned by the relative likelihood of pivotal events (Lemma 1). This is no longer so when a public signal is available. Inspecting Lemma 3, observe that for a marginal activist the relative likelihood of the pivotal events depends upon $s - m$. This opens the possibility of a bandwagon in favor of one of the policies. For instance, a speech in favor of policy A ($s > 0$) can push activists toward A ($m < 0$). This generates positive feedback via the relative-likelihood effect. To explore whether the negative feedback from the relative-preference term (Lemma 2) can slow the bandwagon, we consider the stability of signal-responsive equilibria.

Proposition 5. *Following the observation of the leader's speech s, there is a unique signal-responsive threshold equilibrium in which party members use the threshold*

$$m^\diamond = m^* - \left[\frac{R}{1-R}\right] s.$$

If $R < 1$, then this signal-responsive equilibrium is stable. Hence for $R < 1$, if a leading activist l is able to make a public speech $s = m_l$ then the threshold used by others in a stable equilibrium is precisely her preferred threshold $m^\dagger$ emerging from Proposition 3.

When $R < 1$, a perfectly communicated leadership speech results in the use of the threshold preferred by the leader. When she announces her signal, other activists know what she knows. As they have all information at the leader's disposal and share her preferences, they act in the way that she would like them to. The leader's original frustration with the conference-based threshold equilibrium stemmed from a difference of opinion caused by different signal realizations. Her speech bridges that difference.

We recall (Proposition 3) that R provides an index of the desirability of leadership. It also (Proposition 5) indexes the willingness of a conference to adapt its behavior to a leader's speech. As R increases conference responds, and more so when the leader's message is strong (so that $|s|$ is large). Because conference moves its threshold in response to her speech, there is a direct effect of the leader's intervention on policy. Suggestively put, Michel's Ratio R indexes the feasibility of a leader's influence.

We note that, when $R < 1$, a leader always wishes to guide conference to her preferred threshold $m^\dagger$. If she communicates perfectly, her speech is persuasive; activists share her assessment and adopt her desired threshold. If perfect communication is elusive, however, then the situation is more complex. Suppose that the leader is unable to explain her views, but is allowed (if she wishes) to intervene and choose the policy. It is in this situation (at least for $R < 1$) that only an extremist with a strong signal of the correct thing to do would choose to dictate (Proposition 4). No such difference arises when the leader can explain herself: she would always do so and hence ensure the use of her preferred threshold rule. Thus the extremist–moderate divide arises only when a clear public speech is impossible. This suggests the importance of rhetoric in the classical sense: a democratic assembly is constrained by any limits to communication.

Of course, our discussion here has restricted to the case $R < 1$. When $R > 1$, the signal-responsive equilibrium threshold $m^\diamond$ has strange properties. If activists were to use this threshold, then they would shift away from the policy suggested by the leader's speech. For $R > 1$, however, it makes more sense to look toward the fully coordinated equilibria in which activists ignore their signals. To see why, consider a fall in the threshold to $m < m^\diamond$. This increases the relative likelihood term: a bias toward A is self-reinforcing, putting further downward pressure on m. However, it also lowers the relative conditional preference for A versus B: because B is now relatively harder to achieve, activists bias toward it. When $R > 1$ positive feedback dominates negative feedback. This suggests that a bandwagon toward full coordination can begin to accelerate. (The comparison of feedback effects did not work in a conference environment, because the relative-likelihood term disappeared when we considered a marginal activist with a signal $m_i = m$.)

Proposition 6. *If $R > 1$ then the signal-responsive equilibrium with threshold $m^\diamond$ is unstable. However, the fully coordinated equilibria in which activists ignore their signals (either eveyone advocates A or everyone advocates B) are stable in the following sense: for $m < m^\diamond$, an activist with a signal $m_i = m$ strictly prefers to back A, and for $m > m^\diamond$ an activist $m_i = m$ strictly prefers to back B.*

Using stability as an equilibrium-selection criterion and when $R > 1$, Proposition 6 tells us that we must look toward the two fully coordinated equilibria in which activists ignore their private assessment of the party mood. But which of these two equilibria, if any, will the mass adopt?

Following Proposition 3, a leader is delighted to provide a focal resolution to this coordination problem because she wishes to dictate. For instance, a clear and unambiguous announcement that "everyone should back A" is an obvious focal point. Of course, we are saying nothing new here; many examples of focal points were suggested in a range of complete-information coordination games described in the classic work of Schelling (1960). However, we can offer a further justification. Consider the following rhetoric:

> My assessment is $s = m_l$. I would like you to advocate policy A if and only if your own assessment is higher than $m^\dagger$. Prior to assessing the policies, you would unanimously wish to commit to following my recommendation. Now that you have formed your own private opinions, you have no reason to deviate from my recommendation. Hence you should follow my advice. Given that you do, I have no reason to mislead you.

The leader conveys valid reasons for her recommendation, and it is this feature of her appeal that makes it focal. In the absence of her informed assessment, her speech would lack focal properties, despite the fact that it is commonly understood. The leader is simply asking others to do what they would wish to commit themselves to *ex ante*, and noting that they have no reason to deviate *ex post*. Finally, the leader notes that she has no reason to misrepresent her views. In all, her request for others to follow her is compelling.

Note that, because activists can calculate the threshold they mutually prefer *ex ante* based on their common observation of the leader's public signal, the leader need not make an explicit policy recommendation (although she can). Our argument is independent of R and leads to the following proposition.

Proposition 7. *Suppose that an activist stands as leader and makes a perfectly communicated speech describing her views and describing a policy choice. If she does so, then party members ignore their private information and precisely adopt her recommended policy if and only if $R > 1$.*

Moving beyond the leader's rhetoric, when $R > 1$, there are further justifications for unification behind the leader's preferred policy. Consider a world in which, prior to the leader's speech, activists employ the (stable) threshold m^* from Proposition 1. The leader then speaks, with a speech $s > 0$ in favor of policy A. This speech will cause an individual party activist to reappraise positively the relative likelihood of $\mathcal{P}_A$ versus $\mathcal{P}_B$. Given that others use the threshold m^*, he now finds it optimal to use a lower threshold $m < m^*$. Of course, an activist might then anticipate that other party members will follow the same thought process. If he does, then he now expects them to use a lower threshold than before. Since $R > 1$, positive feedback exceeds negative feedback; hence he will find it optimal to push down his own threshold still further. This heuristic "ficticious play" exercise continues until full coordination on policy A results.

OLIGARCHY

We now study the emergence of an organized group (a clique) within the party's ranks. If activists follow the lead provided by this clique, then the democratic rule by conference is replaced by an oligarchy.

We suppose that a clique of k activists join together and share their views clearly among themselves. They reach a consensus and develop common beliefs about the party mood. Thus k is small enough to allow a mutual understanding to form. Whereas the beliefs of an individual have precision ψ, those of the clique (an average of their signals) have precision $k\psi$, because they are based on k conditionally independent signals of θ.[11] Hence, k indexes the sense of direction of the clique relative to that of an individual. This scenario is equivalent to one in which a leader is better able to assess policies than other party members; equivalently we can think of a leader who canvasses private opinions of party members. By studying the formation of a clique, we merely provide a microfoundation for such a sharpened sense of direction.

Following our earlier thought experiment, we ask whether the k-strong clique wishes to dictate or to defer back to conference. Because the precision of their shared beliefs is $k\psi$ rather than ψ, they have greater confidence in their ability to do the right thing. The same logic presented in our analysis of the leader's thought experiment applies here. However, the criterion determining the desire for leadership becomes $kR > 1$.

Proposition 8. *If $kR > 1$ then a k-strong clique prefers to dictate policy rather than allow a conference threshold strategy to operate. If $kR < 1$, then they prefer others to use a threshold*

$$m^{\ddagger} = m^{*} - \frac{kR}{1 - kR}\,\bar{m}$$

where m^* is the equilibrium threshold from Proposition 1 and where $\bar{m}$ is the average private signal among the clique. $m^{\ddagger}$ is decreasing in the clique's shared assessment of the fundamentals; they prefer others to operate a threshold that is biased toward their shared opinion of the policies' relative merits.

Michels' Ratio R continues to drive the desire to lead and the comparative-static properties of the clique's preferred threshold precisely match those for the case of a single leader. (In fact, Proposition 3 is a special case of Proposition 8 for $k = 1$.) We can also extend the extremist-moderate taxonomy of Proposition 4 to a k-strong clique. Recall that this classification arose when we asked a leader to choose between dictating policy herself and retaining the equilibrium threshold m^* from the party conference.

Proposition 9. *If $kR < 1$ then a k-strong clique is content with conference if and only if $|\bar{m}| < \tilde{m}$, for some $\tilde{m} > 0$. Otherwise, if $|\bar{m}| > \tilde{m}$, they would prefer to dictate by choosing the policy themselves. $\tilde{m}$ rises with the need for direction λ but falls with the aggregate height $\pi_A + \pi_B$ of the barriers to coordination and the sense of direction ψ. Changing these parameters, $\tilde{m} \to 0$ as $kR \to 1$.*

We now turn to consider the feasibility of group-based leadership. Suppose that the k-strong clique becomes an elite: a group of activists who are able to communicate clearly their views to the party membership. Equivalently, the elite is able to put forward a single representative who can perfectly express their views $s = \bar{m}$ by a speech to conference. Following this speech, activist i's updated beliefs will satisfy

$$\theta \mid (s, m_i) \sim N\left(\frac{ks + m_i}{k+1}, \frac{1}{(k+1)\psi}\right).$$

He pays more attention to the views of the elite than to his own signal. His perspective is shared by others, enhancing the positive-feedback effect. Unsurprisingly, positive feedback exceeds negative feedback if and only $kR > 1$, and this criterion is central to our final proposition which extends Propositions 5–7.

Proposition 10. *Following a speech by a k-strong elite, there is a unique signal-responsive equilibrium in which party members use a threshold $m^{\ddagger}$. If $kR < 1$, then this equilibrium is stable. If $kR > 1$, then it is unstable, but the fully coordinated equilibria are stable in the sense used in Proposition 6. Activists following the advice of the elite will always play a stable equilibrium. They ignore their private information (they defer to a de facto oligarchy) if and only if $kR > 1$. Hence $1/R$ is the minimum size of a successful Michelsian oligarchy. This size increases with the need for direction λ, but decreases with the height $\pi_A + \pi_B$ of the barriers to coordination and the sense of direction ψ.*

The first element concerns a stable signal-responsive equilibrium when $kR < 1$. The equilibrium threshold $m^{\ddagger}$ is that preferred by the elite. The party follows the elite's advice, because this is how they would play if they could commit *ex ante*. The elite shapes policy, but conference remains sovereign.[12] The second element concerns the case $kR > 1$ when a signal-responsive equilibrium is no longer stable and the elite would ideally like to see full coordination behind their chosen policy. The third element describes the emergence of an oligarchy. As it was under the leadership of an individual, the advice of the elite is compelling. Prior to the realization of their signals, but after listening to the elite, activists would unanimously wish to follow the elite's advice. They perfectly coordinate and the elite becomes a *de facto* oligarchy.

Finally, Proposition 10 reveals that the desirability and feasibility of leadership, by either an individual or an elite, is intrinsically linked by R. The inverse of Michels' Ratio provides a lower bound to the size of an oligarchy. The precision of the elite's aggregate signal of the party mood is increasing in k. An activist mass gives way to the elite only when this precision is sufficiently high. Adopting a more general interpretation

[11] Alternatively, we might specify conditional correlation of the clique's signals. For instance, if the clique forms from an impromptu meeting in the conference bar, then their signals might be based on similar information sources, and so might be correlated conditional on the party mood θ. Whereas the average of their signals continues to provide a sufficient statistic for beliefs about θ, its effective precision is lower. Thus, when k activists share conditionally correlated signals, the effective size of the clique is lower than k.

[12] When $k > 1$ the elite has a better sense of direction than an individual. However, our results also apply when $k < 1$. Letting $k \to 0$ the criterion $kR < 1$ is satisfied and the elite's preferred threshold converges to m^*. Hence the equilibrium-selection argument that justifies the focus on an elite's preferred threshold equilibrium can also be applied to the threshold in a conference environment.

of the parameter k, an elite becomes an oligarchy so long as its sense direction is sufficiently sharp.

DISCUSSION

Although $k>1/R$ turns an elite into an oligarchy, further increases in k enhance the quality of the oligarchy's leadership. Given that this is so, will the oligarchy grow without bound and incorporate the entire party membership? As it does so, it will begin to reflect the aggregation properties of the general conference. The problem that arises is one of communication: when the elite's membership is large it will find it more difficult to aggregate successfully its views and communicate a coherent and easily understood message to a wider audience. This places an upper bound on the feasible size of a Michelsian oligarchy.

This discussion relates to a further institutional possibility. Suppose that a leader mediates by privately canvassing opinion among the membership. How many opinions should she solicit? Ideally, she would incorporate all views. Alas, the constraints on communication will prevent her from aggregating and articulating such a large and diverse range of views. Nevertheless, our results provide a lower bound on the number of activists she would need to caucus if she is to successfully coordinate the activist mass: the lower bound is $1/R$, just as it is for the Michelsian oligarchy.

We have already noted that the communication could be enhanced if a mediator were to call for a party vote. This, however, brings back the original coordination problem: if the vote is close, then the public see a divided party. Certainly, it is a stylized fact that divided parties tend not to achieve electoral success. Why might this be? We tentantively suggest two reasons. First, close calls are more likely to arise when each activist's sense of direction is limited, and so the quality of the party's assessments is low. Second, outsiders erroneously attribute the split to fundamental differences in the values of the party's membership.

In our introductory remarks we mentioned an alternative interpretation of our game: the coordination problem faced by instrumental voters in plurality-rule elections. Beyond the New York case (Table 1), there are many examples: in the United Kingdom's General Election of 1997 many left-leaning voters wished to coordinate in order to oust the incumbent Conservative administration. In such scenarios a strategic voter switches away from his preferred candidate and toward a better-placed alternative (Fisher 2004). This definition does not fit neatly within our model, since our activists share common preferences but are unsure of them; they must use their signals and envisage pivotal events in order to work out what they like.

This feature (common but unknown preferences) distinguishes our model from the theory of strategic voting developed by Myatt (2007). In his theory, voters share a dislike of a third candidate but differ in their preferences over the two challengers. The voters know what they like, and so face a genuine trade-off between their preferences over candidates and the relative likelihood of influencing the different possible pivotal events. We believe that his specification provides a better insight into the coordination problem faced by voters, whereas our common-value game is better suited to the analysis of the dynamics of leadership. Furthermore, whereas our model could be used to depict a strategic-voting scenario, it is less clear that leadership can provide a resolution to the voters' coordination problem. In the voting scenario, we can think of a leader's speech as a publicly announced signal of the candidates' strengths. When the public signal is good enough, there is perfect coordination with only two candidates receiving votes (one conservative and one liberal in the New York case) as predicted by the psychological effect of Duverger's Law. Of course, for this to work the leader must share the common objectives of the coordinating voters and must not be motivated by other concerns. Such commonality of interest seems rather less plausible in a world with heterogeneous preferences; hence our preferred focus is the coordination of party activists.

Nevertheless, there are common elements to these papers. In Myatt's (2007) paper a public coordination device is not a leader but rather a commonly observed opinion poll. He shows that voters fully coordinate (so that only two candidates receive votes) if and only if this opinion poll is sufficiently precise relative to the privately observed information of voters. We find, therefore, that both Michels' Iron Law of Oligarchy, on which we have focused here, and Duverger's Law, can be linked via the same kind of criterion; it is the precision of public versus private information that leads to coordination in both cases.

TOWARD A THEORY OF LEADERSHIP

Leadership was central to Levi's (2006, 5) "desire to understand what makes for good governments and how to build them." In our view a leader's assessment of policies' relative merits provides a coordinating focal point. Our analysis thus adds to the literature on the role of institutions in helping coordination (Calvert 1995; Myerson 2004; Weingast 1997). We note, however, a subtle feature of our approach: leadership is important not only because it provides a common understanding of play but also because it provides payoff-relevant information. Of course, there are other means of providing such information: we have compared leadership to a more democratic information source, namely a stylized party conference.

In assessing these institutional forms we were inspired by Michels (1915). His conceptualization of two mutually incompatible types of internal governance motivated our formal analysis. We have provided micro-foundations for his claim that a "need for leadership" exists in mass psychology:

> Though it grumbles occasionally, the majority is really delighted to find persons who will take the trouble to look after its affairs. In the mass, and even in the organized mass of the labor parties, there is an immense need for direction and guidance. (Michels, 1915, 38)

In our world, the need for direction, barriers to coordination, and an activist's sense of direction combine to give a single measure (Michels' Ratio) of leadership. It indexes not only the willingness of activists to modify their behavior in the light of a leader's speech, but also their willingness to abandon the conference forum altogether and follow a leader's prescription; it is a complete index of the feasibility of leadership. Moreover, the desirability of leadership is determined by the same combination of variables.

Echoing Michels' claim, the need for leadership is felt when barriers to coordination are high (so that the coordination problem is severe) and when a leader's sense of direction is sharp (so that she knows what to do). In contrast, our "need for direction" works in favor of a party conference: although leadership enhances the clarity of intra-party communication and so avoids the penalty of mis-coordination, it lessens the response of policy to the underlying fundamentals. Our analysis thus contributes to an understanding of the trade-off between the responsiveness of policy outcomes and concentration of power in the form of dictatorship or oligarchy, central to the formal analysis of social-choice mechanisms.

Our conceptualization of leaders differs from previous formal studies which cast leaders in the role of agents under the control of a legislative body (Fiorina and Shepsle 1989). In those studies, a leader possesses skills necessary to the achievement of collective goals; the gap in expertise between leaders and followers underlies a common-agency problem. An interesting feature of our analysis is that neither the desirability nor the feasibility of leadership depend critically upon the skill set of a leader. Nevertheless, our framework can further illuminate this issue. For example, the establishment of an oligarchy allows members of an elite to pool their information; this, in turn, allows an oligarchy to convey more precise information. Of course, individuals differ in their ability to evaluate information and convey messages. An extension of our model would allow for an exploration of this and other individual traits.

Our analysis of the coordination problem faced by activists is devoid of factional conflict: activists share common values but differ in their informed opinions of the path the party should take. We have captured a key element of intraparty division; that which pertains not to core values, but how best to achieve goals related to those values. Uncertainty over how to achieve goals underpins any division. Even in our common value game, a degree of factionalism may, nevertheless, emerge: those with neutral signals are more likely to place their trust in the sovereignty of conference, whereas those with extreme signals (pointing strongly in favor of a policy option) are willing to abandon conference as a central democratic institution.

The absence of any conflict of interest in our model helps a leader to communicate meaningful information. Our results suggest, however, that the clarity of a leader's communication is also important. Conference, acting as a central democratic institution, can (in aggregate) correctly assess the merits of policy, but as a mechanism for communication its performance is poor. A leader, by contrast, can convey only her private assessment of the party mood, but is able to do so with clarity. Moreover, the ability of a leader to convey clearly her message is relevant to institutional choice. For example, when a leader has a moderate signal then, faced with a choice, she would always wish to guide conference toward the use of her preferred threshold rather than to dictate; her willingness and ability to do so depends on her ability to communicate perfectly. Our focus on communication thus contributes to a broader understanding of different forms of governance, such as democracy and oligarchy, which until now have been studied formally only under the guise of commitment problems with regard to economic redistribution (Acemoglu and Robinson 2000, 2001). Finally, and perhaps most importantly, our results suggest a formal analysis of the role of rhetoric in effective leadership. Our next step (Dewan and Myatt 2007) pursues this line of enquiry.

TECHNICAL APPENDIX

Here we develop a formal model which encompasses the three scenarios (conference, leadership, and oligarchy) considered in the text, and provide proofs of Lemmas 1–3 and Propositions 1–10.

Beliefs. Activist i updates a diffuse prior over θ following his observation of signals $m_i \mid \theta \sim N(\theta, 1/\psi)$ and $s \mid \theta \sim N(\theta, 1/[k\psi])$. ($k = 0$ is a conference, $k = 1$ is a leader, and $k > 1$ is a k-strong elite.) Conditional on θ, signals are independent. Updating to form a posterior $G(\theta \mid s, m_i)$ with density $g(\theta \mid s, m_i)$,

$$\theta \mid (m_i, s) \sim N\left(\frac{ks + m_i}{k+1}, \frac{1}{(k+1)\psi}\right) \Rightarrow G(\theta \mid s, m_i) = \Phi\left(\sqrt{(k+1)\psi}\left[\theta - \frac{ks + m_i}{k+1}\right]\right), \qquad \textbf{(1)}$$

where $\Phi(\cdot)$ is the distribution function of the standard normal. Now suppose that other activists employ a threshold strategy. Conditional on θ, an activist backs A with probability p where $p = \Phi(\sqrt{\psi}(\theta - m))$. Writing $F(p \mid s, m_i)$ and $f(p \mid s, m_i)$ for the distribution and density of beliefs about p,

$$f(p \mid s, m_i) = \frac{1}{dp/d\theta} \times g(\theta \mid s, m_i) = \frac{g(\theta \mid s, m_i)}{\sqrt{\psi} \times \phi(\Phi^{-1}(p))}, \quad \text{where} \quad \theta = m + \frac{\Phi^{-1}(p)}{\sqrt{\psi}}, \qquad \textbf{(2)}$$

and where $\phi(\cdot)$ is the density of the standard normal.

Pivotal Probabilities. Fixing activist i and abusing notation slightly, write $x \in \{0, 1, \ldots, n-1\}$ for the number of others who advocate A. Conditional on θ, party members back A with probability p; hence x is a draw from the binomial with parameters p and $n-1$. However, p is uncertain and activist i must take expectations to form $\Pr[x \mid s, m_i] = \int_0^1 \binom{n-1}{x} p^x (1-p)^{n-1-x} f(p \mid s, m_i)\, dp$. Activist i is pivotal for the success of A if and only if $x \leq p_A n < x + 1$. We write x_n^A

for the unique value of x that satisfies these inequalities, and note that $[x_n^A/n] \to p_A$ as $n \to 0$. Clearly,

$$\Pr[\mathcal{P}_A \mid s, m_i] = \int_0^1 \binom{n-1}{x_n^A} p^{x_n^A}(1-p)^{n-1-x_n^A} f(p \mid s, m_i)\, dp.$$

This probability vanishes as $n \to \infty$ but, applying Proposition 1 of Chamberlain and Rothschild (1981), $n \times \Pr[\mathcal{P}_A \mid s, m_i] \to f(p_A \mid s, m_i)$ and $n \times \Pr[\mathcal{P}_B \mid s, m_i] \longrightarrow f(p_B \mid s, m_i)$ as $n \to \infty$. This (in essence) is the Law of Large Numbers: when n is large, the proportion supporting policy converges in probability to p. Pivotal probabilities are determined by beliefs about p via the density $f(p \mid s, m_i)$. The probability of being pivotal vanishes with $1/n$. However, an activist cares about the *relative* likelihood of $\mathcal{P}_A$ and $\mathcal{P}_B$:

$$\begin{aligned}\lim_{n\to\infty}\left[\log\frac{\Pr[\mathcal{P}_A \mid s, m_i]}{\Pr[\mathcal{P}_B \mid s, m_i]}\right] &= \log\frac{f(p_A \mid s, m_i)}{f(p_B \mid s, m_i)} = \log\frac{\phi(\Phi^{-1}(p_B))}{\phi(\Phi^{-1}(p_A))} + \log\frac{g(\theta_A \mid s, m_i)}{g(\theta_B \mid s, m_i)}\\ &= \log\frac{\phi(\pi_B)}{\phi(\pi_A)} + \log\frac{g(\theta_A \mid s, m_i)}{g(\theta_B \mid s, m_i)}\\ &= \frac{\pi_A^2 - \pi_B^2}{2} + \log\frac{g(\theta_A \mid s, m_i)}{g(\theta_B \mid s, m_i)}. \end{aligned} \tag{3}$$

The first equality follows from Chamberlain and Rothschild (1981); the second from Equation (2); the third from substitution of π_A and π_B and the symmetry of the normal which ensures that $\phi(\pi_B) = \phi(-\pi_B)$ and $-\Phi^{-1}(1 - p_B) = \Phi^{-1}(p_B)$; and the fourth from substitution into the standard normal density $\phi(z) \propto \exp(-z^2/2)$. The notation θ_A and θ_B is from the text: $\theta_A \equiv m + [\Phi^{-1}(p_A)/\sqrt{\psi}] = m + [\pi_A/\sqrt{\psi}]$ and $\theta_B \equiv m + [\Phi^{-1}(p_B)/\sqrt{\psi}] = m - [\pi_B/\sqrt{\psi}]$; the second equality again exploits from the symmetry of the normal. Taking the posterior beliefs of activist i from Equation (1) and evaluating at θ_A, we obtain

$$\begin{aligned} g(\theta_A \mid s, m_i) &= \sqrt{(k+1)\psi}\,\phi\left(\sqrt{(k+1)\psi}\left[\theta_A - \frac{ks+m_i}{k+1}\right]\right)\\ &\propto \exp\left(-\frac{(k+1)\psi}{2}\left[\theta_A - \frac{ks+m_i}{k+1}\right]^2\right)\\ &= \exp\left(-\frac{(k+1)\psi}{2}\left[m + \frac{\pi_A}{\sqrt{\psi}} - \frac{ks+m_i}{k+1}\right]^2\right)\\ &= \exp\left(-\frac{k+1}{2}\left[\psi\left(m - \frac{ks+m_i}{k+1}\right)^2 + 2\sqrt{\psi}\pi_A\left(m - \frac{ks+m_i}{k+1}\right) + \pi_A^2\right]\right), \end{aligned} \tag{4}$$

where for the second step we have applied the formula for the density of the normal, and we have omitted the multiplicative constant that will be shared with the density $g(\theta_B \mid s, m_i)$. The final two equalities follow from substitution of θ_A and algebraic manipulation. Similarly,

$$g(\theta_B \mid s, m_i) \propto \exp\left(-\frac{k+1}{2}\left[\psi\left(m - \frac{ks+m_i}{k+1}\right)^2 - 2\sqrt{\psi}\pi_B\left(m - \frac{ks+m_i}{k+1}\right) + \pi_B^2\right]\right). \tag{5}$$

Combining the expressions from Equations (4) and (5), we obtain

$$\log\frac{g(\theta_A \mid s, m_i)}{g(\theta_B \mid s, m_i)} = (k+1)\left[\frac{\pi_B^2 - \pi_A^2}{2} + \sqrt{\psi}(\pi_A + \pi_B)\left(\frac{ks+m_i}{k+1} - m\right)\right]. \tag{6}$$

Substituting Equation (6) into Equation (3) we obtain

$$\lim_{n\to\infty}\left[\log\frac{\Pr[\mathcal{P}_A \mid s, m_i]}{\Pr[\mathcal{P}_B \mid s, m_i]}\right] = \frac{k[\pi_B^2 - \pi_A^2]}{2} + (k+1)\sqrt{\psi}(\pi_A + \pi_B)\left(\frac{ks+m_i}{k+1} - m\right). \tag{7}$$

Conditional Preference. Here we study payoffs conditional on the events $\mathcal{P}_A$ and $\mathcal{P}_B$. In a large party, the proportion of activists supporting policy A converges in probability to p. Hence $\mathcal{P}_A$ occurs if and only if $p \approx p_A$, or equivalently $\theta \approx \theta_A$. So, as $n \to \infty$, $\mathrm{E}[u_A(\theta) \mid \mathcal{P}_A] \to u_A(\theta_A) = \exp(\lambda\theta_A/2)$. Hence,

$$\begin{aligned}\lim_{n\to\infty}\log\frac{\mathrm{E}[u_A(\theta) \mid \mathcal{P}_A]}{\mathrm{E}[u_B(\theta) \mid \mathcal{P}_B]} &= \log\frac{u_A(\theta_A)}{u_B(\theta_B)} = \frac{\lambda(\theta_A + \theta_B)}{2}\\ &= \lambda\left[m + \frac{\pi_A - \pi_B}{2\sqrt{\psi}}\right]. \end{aligned} \tag{8}$$

Proof of Lemmas 1–3. Apply Equations (7) and (8). □

Optimal Advocacy. We now consider the decision of an activist given that the party is large.

$$\begin{aligned}&\lim_{n\to\infty}\left[\log\frac{\Pr[\mathcal{P}_A]}{\Pr[\mathcal{P}_B]} + \log\frac{\mathrm{E}[u_A(\theta) \mid \mathcal{P}_A]}{\mathrm{E}[u_B(\theta) \mid \mathcal{P}_B]}\right]\\ &= \underbrace{\frac{k[\pi_B^2 - \pi_A^2]}{2} + (k+1)\sqrt{\psi}(\pi_A + \pi_B)\left(\frac{ks+m_i}{k+1} - m\right)}_{\text{from Equation (7)}}\\ &\quad + \underbrace{\lambda\left[m + \frac{\pi_A - \pi_B}{2\sqrt{\psi}}\right]}_{\text{from Equation (8)}}\\ &= \lambda\left[\frac{k(\pi_B - \pi_A)(\pi_A + \pi_B)\sqrt{\psi}}{2\lambda\sqrt{\psi}} + (k+1)\right.\\ &\quad \left.\times\frac{\sqrt{\psi}(\pi_A + \pi_B)}{\lambda}\left(\frac{ks+m_i}{k+1} - m\right) + m + \frac{\pi_A - \pi_B}{2\sqrt{\psi}}\right]\\ &= \lambda R\left[m_i + ks + \left(k - \frac{1}{R}\right)m^* - \left(k + 1 - \frac{1}{R}\right)m\right]. \end{aligned} \tag{9}$$

The second equality follows from rearrangement. The third equality follows from the substitution of Michels' Ratio R and m^* from Proposition 1 and further manipulation. Observe that the final expression is increasing in m_i; hence any optimal best reply (in a large party) is a threshold rule. Furthermore, when an activist's signal is equal to the threshold used by others:

$$m_i = m \quad \Rightarrow \quad \lim_{n\to\infty}\left[\log\frac{\Pr[\mathcal{P}_A]}{\Pr[\mathcal{P}_B]} + \log\frac{\mathrm{E}[u_A(\theta) \mid \mathcal{P}_A]}{\mathrm{E}[u_B(\theta) \mid \mathcal{P}_B]}\right] = \lambda R\left[ks + \left(k - \frac{1}{R}\right)(m^* - m)\right], \tag{10}$$

which is increasing in m if and only if $kR < 1$. For a threshold equilibrium, this needs to be zero:

$$\lambda R\left[ks + \left(k - \frac{1}{R}\right)(m^* - m)\right] = 0$$
$$\Leftrightarrow \quad m = m^* - \left[\frac{Rk}{1 - Rk}\right]s. \qquad (\mathbf{11})$$

When a public signal is absent ($k = 0$) then the equilibrium threshold is m^* from the statement of Proposition 1. If a public signal is in favor of policy A (that is, when $s > 0$) then the equilibrium threshold is pushed down so long as $Rk < 1$. The extent of this effect is increasing in R and k.

Proof of Propositions 1 and 2. Apply Equations (10) and (11) with $k = 0$. □

Choosing a Threshold. $U(m, m_l) = \int_{-\infty}^{\theta_B} u_B(\theta)g(\theta \mid m_l)\, d\theta + \int_{\theta_A}^{\infty} u_A(\theta)g(\theta \mid m_l)\, d\theta$. Evaluating at $m_l = m$ and recalling that θ_A and θ_B are linearly increasing in m, differentiate to obtain

$$\begin{aligned}
\frac{\partial U}{\partial m} &= u_B(\theta_B)g(\theta_B) - u_A(\theta_A)g(\theta_A) < 0 \\
&\Longleftrightarrow \log\frac{u_A(\theta_A)}{u_B(\theta_B)} + \log\frac{g(\theta_A \mid m_l)}{g(\theta_B \mid m_l)} > 0 \\
&\Longleftrightarrow \lim_{n\to\infty}\left[\log\frac{\Pr[\mathcal{P}_A \mid m_l]}{\Pr[\mathcal{P}_B \mid m_l]} + \log\frac{\mathrm{E}[u_A(\theta) \mid \mathcal{P}_A]}{\mathrm{E}[u_B(\theta) \mid \mathcal{P}_B]}\right] \\
&\quad + \frac{\pi_B^2 - \pi_A^2}{2} > 0,
\end{aligned}$$

where the final equivalence is from (3). The criterion is (†) from the text. The sum of the first two terms (it is the same as (9) setting $m_i = m_l$ and $k = 0$) is strictly decreasing in m if $R > 1$ and strictly increasing if $R < 1$. Thus (for the generic case $R \neq 1$) a unique $m^\dagger$ satisfies $\partial U/\partial m = 0$. This must be a local minimum (and global minimum, since there is only one stationary point) if $R > 1$, and so $U(m, m_l)$ is maximized by choosing either $m \to \infty$ or $m \to -\infty$. If $R < 1$, then $m^\dagger$ yields a global maximum. Explicitly,

$$\begin{aligned}
\frac{\partial U}{\partial m} = 0 \quad &\Leftrightarrow \quad \underbrace{(\pi_A + \pi_B)\sqrt{\psi}(m_l - m^\dagger)}_{\text{from Equation (7) with } k=0} \\
&\quad + \underbrace{\lambda\left[m^\dagger + \frac{\pi_A - \pi_B}{2\sqrt{\psi}}\right]}_{\text{from Equation (8)}} + \frac{\pi_B^2 - \pi_A^2}{2} = 0 \\
\Leftrightarrow \quad m^\dagger &= \frac{\pi_B - \pi_A}{2\sqrt{\psi}} - \frac{(\pi_A + \pi_B)\sqrt{\psi}m_l}{\lambda - \sqrt{\psi}(\pi_A + \pi_B)} \\
&= m^* - \left[\frac{R}{1 - R}\right]m_l, \qquad (\mathbf{12})
\end{aligned}$$

where the solution for $m^\dagger$ follows from simple algebraic manipulation.

Proof of Proposition 3. Apply Equation (12). □

Extremists and Moderates. If $z \sim N(\mu, \sigma^2)$ then for real-valued constants b and $H > L$,

$$\begin{aligned}
&\int_L^H \exp(bz)\, d\Phi\left(\frac{z - \mu}{\sigma}\right) \\
&= \exp\left(b\mu + \frac{b^2\sigma^2}{2}\right) \\
&\times \left[\Phi\left(\frac{H - \mu - b\sigma^2}{\sigma}\right) - \Phi\left(\frac{L - \mu - b\sigma^2}{\sigma}\right)\right]. \qquad (\mathbf{13})
\end{aligned}$$

We can use (13) to calculate a leader's expected payoff. Suppose that she believes $\theta \sim N(\mu, \sigma^2)$. Write $U_A = \mathrm{E}[u_A(\theta)]$ and $U_B[u_B(\theta)]$ for her payoffs when she dictates the adoption of policies A and B, respectively. For U_A we set $b = \lambda/2$, and for U_B we set $b = -\lambda/2$. Hence

$$U_A = \exp\left(\frac{\lambda\mu}{2} + \frac{\lambda^2\sigma^2}{8}\right) \quad \text{and}$$
$$U_B = \exp\left(-\frac{\lambda\mu}{2} + \frac{\lambda^2\sigma^2}{8}\right) = \exp(-\lambda\mu) \times U_A. \qquad (\mathbf{14})$$

$U_A > U_B$ if and only if $\mu > 0$: a leader implements A if and only if she expects the underlying state of the world to favor it. Next, consider her payoff when others use a threshold m. Policy A wins if $\theta > \theta_A$, policy B wins if $\theta < \theta_B$. So, writing $\mathcal{I}[\cdot]$ for the indicator function and applying (13),

$$\begin{aligned}
U_C &= \mathrm{E}\left[\exp\left(-\frac{\lambda\theta}{2}\right) \times \mathcal{I}[\theta < \theta_B]\right] \\
&\quad + \mathrm{E}\left[\exp\left(\frac{\lambda\theta}{2}\right) \times \mathcal{I}[\theta > \theta_A]\right] \\
&= U_B \times \Phi\left(\frac{\theta_B - \mu}{\sigma} + \frac{\lambda\sigma}{2}\right) \\
&\quad + U_A \times \left[1 - \Phi\left(\frac{\theta_A - \mu}{\sigma} - \frac{\lambda\sigma}{2}\right)\right] \\
&= U_A \times \left[\exp(-\lambda\mu) \times \Phi\left(\frac{\theta_B - \mu}{\sigma} + \frac{\lambda\sigma}{2}\right)\right. \\
&\quad \left. + \Phi\left(\frac{\mu - \theta_A}{\sigma} + \frac{\lambda\sigma}{2}\right)\right],
\end{aligned}$$

where the final equality stems from $U_B = \exp(-\lambda\mu) \times U_A$ and from the symmetry of the normal. Now, suppose that the leader observes a signal m_l with precision $k\psi$ so that $\theta \mid m_l \sim N(m_l, [1/k\psi])$. Without loss of generality, we set $m_l > 0$, so that $U_A > U_B$. Now, setting $\mu = m_l$ and $\sigma^2 = 1/k\psi$,

$$\begin{aligned}
\frac{U_C}{U_A} &= \exp(-\lambda m_l) \times \Phi\left(\sqrt{k\psi}(\theta_B - m_l) + \frac{\lambda}{2\sqrt{k\psi}}\right) \\
&\quad + \Phi\left(\sqrt{k\psi}(m_l - \theta_A) + \frac{\lambda}{2\sqrt{k\psi}}\right). \qquad (\mathbf{15})
\end{aligned}$$

If others adopt a threshold m^* then $\theta_A = (\pi_A + \pi_B)/2\sqrt{\psi}$ and $\theta_B = -(\pi_A + \pi_B)/2\sqrt{\psi}$. Hence

$$\frac{U_C}{U_A} = \exp(-\lambda m_l) \times \Phi(X - \sqrt{k\psi}m_l) + \Phi(X + \sqrt{k\psi}m_l)$$
$$\text{where } X \equiv \frac{\lambda(1-kR)}{2\sqrt{k\psi}} = \frac{\lambda}{2\sqrt{k\psi}} - \frac{\sqrt{k}(\pi_A + \pi_B)}{2}. \qquad (\mathbf{16})$$

Proof of Proposition 4. Consider a leader with a signal $m_l > 0$. (The case $m_l < 0$ is symmetric.) If she were to dictate then she would implement policy A and enjoy a payoff of U_A. By deferring to the equilibrium threshold of conference she enjoys a payoff U_C. She strictly prefers conference if

$$\begin{aligned}\frac{U_C}{U_A} &> 1 \Leftrightarrow \exp(-\lambda m_l) \times \Phi(X - \sqrt{k\psi}m_l) \\ &> 1 - \Phi(X + \sqrt{k\psi}m_l) \\ &\Leftrightarrow Y(m_l) \equiv \log\left[\frac{\Phi(X - \sqrt{k\psi}m_l)}{1 - \Phi(X + \sqrt{k\psi}m_l)}\right] - \lambda m_l > 0,\end{aligned}$$

where $k = 1$ for a single leader. At $m_l = 0$ (a neutral signal) this criterion becomes

$$\begin{aligned}\frac{U_C}{U_A} > 1 &\Leftrightarrow \log\left[\frac{\Phi(X)}{1 - \Phi(X)}\right] > 0 \\ &\Leftrightarrow 2\Phi(X) > 1 \quad \Leftrightarrow \quad X > 0.\end{aligned}$$

This last inequality holds if and only if $kR < 1$. Since $Y(m_l)$ is continuous in m_l there is some region of signals close to zero for which the leader strictly prefers to follow conference. In fact, there is a unique $\tilde{m}$ such that (for positive signals) $Y(m_l) > 0$ if and only if $m_l < \tilde{m}$. The proof follows from the claim that $Y(m_l)$ is strictly decreasing in m_l with a derivative that is bounded away from zero. To prove this, write

$$Y(m_l) = \log\left[\frac{1 - \Phi(\sqrt{k\psi}m_l - X)}{1 - \Phi(\sqrt{k\psi}m_l + X)}\right] - \lambda m_l.$$

Next differentiate to obtain

$$\begin{aligned}&Y'(m_l) \\ &= \sqrt{k\psi}\left[\frac{\phi(\sqrt{k\psi}m_l + X)}{1 - \Phi(\sqrt{k\psi}m_l + X)} - \frac{\phi(\sqrt{k\psi}m_l - X)}{1 - \Phi(\sqrt{k\psi}m_l - X)}\right] - \lambda \\ &= \sqrt{l\psi}[h(\sqrt{k\psi}m_l + X) - h(\sqrt{k\psi}m_l - X)] - \lambda,\end{aligned}$$

where $h(z) \equiv \phi(z)/[1 - \Phi(z)]$ is the hazard rate of the standard normal. Applying the mean-value theorem, there is some z satisfying $\sqrt{k\psi}m_l + X > z > \sqrt{k\psi}m_l - X$ such that

$$Y'(m_l) = 2\sqrt{k\psi}Xh'(z) - \lambda = \lambda(1 - kR)h'(z) - \lambda,$$

where the second equality follows from the definition of X in (16). Now, the hazard $h(z)$ of the standard normal is an increasing and convex function of its argument, and is asymptotically linear, so that $h(z) - z \to 0$ as $z \to \infty$; hence $h'(z) \leq 1$ for all z. Hence, for $kR < 1$,

$$Y'(m_l) \leq \lambda(1 - kR) - \lambda = -\lambda kR < 0.$$

Thus $Y(m_l)$ is strictly decreasing in m_l, and the derivative is bounded away from zero. Hence $Y(m_l) < 0$ for m_l sufficiently large, and there is a unique $\tilde{m}$ such that $Y(\tilde{m}) = 0$. □

Proof of Propositions 5 and 6. Apply Equations (10) and (11) with $k = 1$. □

Proof of Propostion 7. From the argument given in the main text. □

Proof of Proposition 8. With an average signal of $\bar{m}$, a k-strong clique's beliefs about the θ are, modifying (1) appropriately, captured by the density $g(\theta \mid \bar{m}) = \sqrt{k\psi}\phi(\sqrt{k\psi}(\theta - \bar{m}))$ where $\phi(\cdot)$ is the density of the standard normal. Following derivations analogous to those leading up to (6),

$$\log\frac{g(\theta_A \mid \bar{m})}{g(\theta_B \mid \bar{m})} = k\left[\frac{\pi_B^2 - \pi_A^2}{2} + \sqrt{\psi}(\pi_A + \pi_B)(\bar{m} - m)\right]. \quad (\mathbf{17})$$

The clique's expected payoff $U(m, \bar{m})$ is locally decreasing in m (so that they favor a shift toward policy A) if and only if $g(\theta_A \mid \bar{m})u_A(\theta_A) > g(\theta_B \mid \bar{m})u_B(\theta_B)$, or, upon substitution,

$$\begin{aligned}&\log\frac{g(\theta_A \mid \bar{m})}{g(\theta_B \mid \bar{m})} + \log\frac{u_A(\theta_A)}{u_B(\theta_B)} \\ &= \underbrace{k\left[\frac{\pi_B^2 - \pi_A^2}{2} + \sqrt{\psi}(\pi_A + \pi_B)(\bar{m} - m)\right]}_{\text{from Equation (17)}} \\ &+ \underbrace{\lambda\left[m + \frac{\pi_A - \pi_B}{2\sqrt{\psi}}\right]}_{\text{from Equation (8)}} > 0. \qquad (\mathbf{18})\end{aligned}$$

The left-hand side of this inequality is strictly decreasing in m (implying that $U(m, \bar{m})$ is quasi-convex in m) if and only if $kR > 1$. So, if $kR > 1$, the clique prefers to dictate policy. For $kR < 1$, setting the expression in (18) to zero and solving for m yields $m^{\ddagger}$. □

Proof of Proposition 9. For general k, the proof of Proposition 4 applies. □

*Proof of Propostion 10.*The first and second claims follow Equations (10) and (11). The remaining claims follow by inspection or from the arguments given in the main text. □

REFERENCES

Acemoglu, Daron, and James A. Robinson. 2000. "Why Did the West Extend the Franchise? Democracy, Inequality, and Growth in Historical Perspective." *Quarterly Journal of Economics* 115 (November): 1167–99.

Acemoglu, Daron, and James A. Robinson. 2001. "A Theory of Political Transitions." *American Economic Review* 91 (September): 938–63.

Angeletos, George-Marios, and Alessandro Pavan. 2004. "Transparency of Information and Coordination in Economies with Investment Complementarities." *American Economic Review: AEA Papers and Proceedings* 94 (May): 91–98.

Calvert, Randall L. 1995. The Rational Choice of Social Institutions: Cooperation, Coordination, and Communication. In *Modern Political Economy*, ed. Jeffrey S. Banks and Eric A. Hanushek. Cambridge, UK: Cambridge University Press.

Carlsson, Hans, and Eric van Damme. 1993. "Global Games and Equilibrium Selection." *Econometrica* 61 (September): 989–1018.

Chamberlain, Gary, and Michael Rothschild. 1981. "A Note on the Probability of Casting a Decisive Vote." *Journal of Economic Theory* 25 (August): 152–62.

Cox, Gary W. 1994. "Strategic Voting Equilibria under the Single Nontransferable Vote." *American Political Science Review* 88 (September): 608–21.

Cox, Gary W. 1997. *Making Votes Count*. Cambridge: Cambridge University Press.

Dewan, Torun, and David P. Myatt. 2007. "The Qualities of Leadership: Direction, Communication, and Obfuscation." *Oxford University Economics Discussion Paper* no. 311.

Feddersen, Timothy J. 1992. "A Voting Model Implying Duverger's Law and Positive Turnout." *American Journal of Political Science* 36 (November): 938–62.

Feddersen, Timothy J., and Wolfgang Pesendorfer. 1996. "The Swing Voter's Curse." *American Economic Review* 86 (June): 408–24.

Feddersen, Timothy J., and Wolfgang Pesendorfer. 1997. "Voting Behavior and Information Aggregation in Elections with Private Information." *Econometrica* 65 (September): 1029–58.

Feddersen, Timothy J., and Wolfgang Pesendorfer. 1998. "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting." *American Political Science Review* 92 (March): 23–35.

Fiorina, Morris P., and Kenneth A. Shepsle. 1989. Formal Theories of Leadership: Agents, Agenda-Setters, and Entrepeneurs. In *Leadership and Politics*, ed. Bryan D. Jones. Lawrence: University Press Kansas, pp. 17–41.

Fisher, Stephen D. 2004. "Definition and Measurement of Tactical Voting: The Role of Rational Choice." *British Journal of Political Science* 34 (January): 152–66.

Goldstein, Itay, and Ady Pauzner. 2005. "Demand-Deposit Contracts and the Probability of Bank Runs." *Journal of Finance* 60 (June): 1293–327.

Hamilton, Alexander, James Madison, and John Jay. 1788. *The Federalist: A Collection of Essays, Written in Favour of the New Constitution, as Agreed by the Federal Convention, September 17, 1787*, New York: J. and A. McLean.

Hands, Gordon. 1971. "Robert Michels and the Study of Political Parties." *British Journal of Political Science* 1 (April): 155–72.

Koelble, Thomas A. 1996. "Economic Theories of Organization and the Politics of Institutional Design in Political Parties." *Party Politics* 2 (March): 251–63.

Levi, Margaret. 2006. "Why We Need a New Theory of Government." *Perspectives on Politics* 4 (March): 5–19.

Lupia, Arthur, and Mathew D. McCubbins. 2000. "Representation or Abdication? How Citizens Use Information to Help Delegation Succeed." *European Journal of Political Research* 37 (May): 291–307.

McKenzie, Robert T. 1963. *British Political Parties: The Distribution of Power within the Conservative and Labour Parties (1955)*. London: Heinemann.

Michels, Robert. 1915. *Political Parties: A Sociological Study of the Oligarchical Tendencies of Modern Democracy*. Kitchener, Ontario: Batoche Books. 2001 edition translated by Eden and Cedar Paul and available from http://onlinebooks.library.upenn.edu/.

Morris, Stephen, and Hyun S. Shin. 1998. "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks." *American Economic Review* 88 (June): 587–97.

Morris, Stephen, and Hyun S. Shin. 2002. "Social Value of Public Information." *American Economic Review* 92 (December): 1521–34.

Morris, Stephen, and Hyun S. Shin. 2003. Global Games: Theory and Applications. In *Advances in Economics and Econometrics, the Eighth World Congress*, ed. M. Dewatripont, L. Hansen and S. Turnovsky. Cambridge, UK: Cambridge University Press.

Morris, Stephen, and Hyun S. Shin. 2004. "Coordination Risk and the Price of Debt." *European Economic Review* 48 (February): 133–53.

Morris, Stephen, Hyun S. Shin, and Hui Tong. 2006. "Social Value of Public Information: Morris and Shin (2002) is Actually Pro-Transparency, not Con: Reply." *American Economic Review* 96 (March): 453–55.

Myatt, David P. 2007. "On the Theory of Strategic Voting." *Review of Economic Studies* 74 (January): 255–81.

Myerson, Roger B. 2004. "Justice, Institutions, and Multiple Equilibria." *Chicago Journal of International Law* 5 (Summer): 91–107.

Myerson, Roger B., and Robert J. Weber. 1993. "A Theory of Voting Equilibria." *American Political Science Review* 87 (March): 102–114.

Olson, Mancur. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge, MA: Harvard University Press.

Palfrey, Thomas R. 1989. A Mathematical Proof of Duverger's Law. In *Models of Strategic Choice in Politics*, ed. Peter C. Ordeshook. Ann Arbor: University of Michigan Press, pp. 69–92.

Schelling, Thomas C. 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.

Svensson, Lars E. O. 2006. "Social Value of Public Information: Comment: Morris and Shin (2002) is Actually Pro-Transparency, not Con." *American Economic Review* 96 (March): 448–52.

Weingast, Barry R. 1997. "The Political Foundations of Democracy and the Rule of Law." *American Political Science Review* 91 (June): 245–63.