Cloturing Deliberation

Vincent Anesi* Mikhail Safronov†

October 21, 2020

Abstract

We study how the institutional arrangements for ending deliberation — the "cloture rules" — interact with collective learning to affect the outcomes of decision making in committees. In contrast to much of the previous literature on deliberative committees, this paper makes a distinction between the final votes over policy proposals and the cloture votes that bring them about. Using this approach, we explore how cloture rules influence the course of deliberation, the likelihood of inefficient deliberative outcomes, the circumstances surrounding failures to bring proposals to a final vote, and the distribution of power among committee members in the deliberative process. We also use our simple model to examine the issue of the stability of cloture rules, characterizing the rules that no coalition of committee members is able or willing to overturn. We show in particular that all cloture rules are dynamically stable.

Keywords: Cloture, deliberation, obstruction, pivots, political failure, stability, voting.

*University of Luxembourg. vincent.anesi@uni.lu.
†University of Cambridge. ms2329@cam.ac.uk.

1 Introduction

In a deliberative committee, a certain amount of agreement is required, implicitly or explicitly, from the members of the committee to end deliberation and take a vote on the issue at hand. In particular, a proposal may fail not because it is voted down, but because it is never brought to a vote. The institutional arrangements for "cloturing" deliberation vary across committees but they appear to be very persistent: committees seldom try and typically fail to change their cloture rules. In the US Senate, for example, the threshold for invoking cloture on a motion is three-fifths of the whole Senate. A well-known implication of this cloture rule is that Senate minorities can (and regularly do) prevent votes on bills that have majority support by extending deliberation — a practice referred to as a "filibuster." In recent years, observers of Senate's daily life have noted that it is "replete with filibusters, threats of extended debates, and cloture votes" (Oleszek, 2016). Although the subject of much discussion and criticism, the Senate's cloture rule has been extremely resilient: with the exception of an amendment in 1975 that changed the number of Senators required to invoke cloture from two-thirds of the Senators present and voting to the current threshold, all attempts to reduce the cloture threshold over the past century have been unsuccessful (e.g., Eidelson 2013). In addition, evidence suggests that Senators are well aware of how changes in the rules of deliberation affect policy outcomes and, historically, their decisions about these rules have been motivated by strategic considerations (e.g., Binder and Smith 1998, and Wawro and Schickler 2010).

The US Senate is probably the most notorious illustration of the strategic implications of cloture rules and obstructionism on collective decision making, but examples can be found in a wide variety of other contexts. Those range from international organizations, where cloture rules are typically explicitly formulated;[1] to academic committees, where a chair often has the power to decide whether and when a proposal should be brought to a vote; to households where an implicit rule often requires that a collective decision (buying a new house, moving to a different country, etc.) can only be made once both parents

[1]See, e.g., Rule 75 of the Rules of Procedures on the closure of deliberation in the General Assembly of the United Nations

consider they have the appropriate information to do so.

These observations raise important questions about the role of cloture rules and the strategic incentives they create in explaining the various outcomes that can occur from deliberation. How do they affect the course of deliberation? How do they determine, for example, whether deliberation will end in a final vote over the issue at hand or in an impasse? Do they generate efficient deliberative outcomes? Do they give some committee members an upper hand over others in the process of ending deliberation? And what explains their apparent stability?

There is an extensive political-economy literature on deliberative committees — which we elaborate on below — but there has been relatively little work done to develop a formal theory of cloture. In particular, the broader political economic implications of cloture rules and their apparent stability are, at present, not well understood. In this paper, we address the questions above in the context of the collective-learning role of deliberation, following the approach initiated by Chan et al. (2018).[2] This approach captures to role of deliberation as a collective-thinking and information-acquisition process that generates knowledge.

We consider a committee that has to decide between two alternatives, a risky reform and a safe status quo policy. Members of the committee are uncertain about the merits of the reform, which is either good or bad. They all share the same ex ante belief that the reform is good but differ in the values they place on a good reform. Prior to voting over the alternatives, they can deliberate, i.e., try to collect more information about the type of the reform. The deliberative session that cannot last more than a fixed number of rounds. Each round begins with a *cloture vote.* If cloture passes then the committee takes a *final vote* either to adopt the reform or to maintain the status quo, thus ending the game; otherwise, deliberation continues. If a final vote has not been held by the end of the last round, then the status quo is maintained. Passage of cloture requires the assent of one from a collection of decisive coalitions, which we refer to as the committee's *cloture rule.* Similarly, passage of the reform requires the assent of one from a collection of decisive coalitions, which may differ from the cloture rule.

[2]Bolton and Faure-Grimaud (2009) also use a learning framework to model individual deliberation.

Every new round of deliberation involves a positive cost to committee members. However, it also generates a public signal about the reform, which provides the committee either with "good news" or "no news." Good news reveals that the reform is good, while no news makes committee members more pessimistic about the type of the reform. Each round then brings a new collective-choice problem, characterized by the common belief that the reform is good, along with a new opportunity either to generate more information about the reform or to make a final decision.

We characterize the unique equilibrium outcome generated by the game described above, for any cloture rule satisfying the minimal conditions of properness and monotonicity. The equilibrium has a number of interesting features:

(i) Obstruction: If the deliberation cost is not too large then after the arrival of good news, an anti-reform coalition — i.e., those committee members who prefer the status quo to a good reform — may block cloture until the end of the session, so as to prevent the adoption of the good reform. We refer to such a tactic as a "filibuster."[3] As the end of the session draws near, the cost of a filibuster decreases. It follows that after sufficiently many rounds of deliberation without good news, when the belief about the type of the reform has become sufficiently pessimistic that it would be voted down in a final vote, committee members realize that the status quo will be maintained regardless. Therefore, they unanimously agree to end deliberation, thus avoiding a filibuster: if deliberation continued (off the path), upon receiving good news the anti-reform coalition would filibuster.

A filibuster can however occur *on the equilibrium path* if the deliberation cost is sufficiently low. This happens if the committee receives good news in a round where, initially: (a) the anti-reform coalition is prepared to filibuster, and (b) the other committee members are still sufficiently optimistic that the reform would be voted up in a final vote. In contrast to the case above, the committee does not cloture deliberation to preempt a filibuster. Anticipating that the reform would be adopted if a final vote were held in this round, the members of the anti-reform coalition block cloture in the hope that continuing

[3]While this term is often interpreted as prolonged speech, it refers (and we think of it as) any action that obstructs decision making by extending deliberation.

deliberation will eventually generate sufficiently pessimistic beliefs about the reform.

(ii) Pareto inefficiency: Equilibria in which filibusters occur with positive probability are Pareto inefficient. When the committee learns that the reform is good and a blocking coalition starts filibustering, all committee members would be better off ending deliberation and voting immediately to maintain the status quo. But this is impossible in equilibrium, as the members who prefer a good reform to the status quo cannot commit to reject the reform if a final vote were held.

(iii) The impact of cloture rules: The political-economy literature typically interprets (and models) changes to cloture rules simply as adjustments to the voting thresholds for passing proposals (e.g., Krehbiel 2006). Our analysis demonstrates how cloture rules interact with collective learning in committees, and how the failure to account for this interaction can lead to mistaken inferences about their impacts on deliberative outcomes. In particular, changes to cloture rules aimed at reducing the length of deliberation may in fact have the opposite effect; and enlarging the set of decisive coalitions that can terminate deliberation reduces the likelihood of filibusters, *even if an anti-reform coalition remains blocking.*

(iv) Pivots: As in standard models of spatial voting, the left and right pivots of the cloture rule — i.e., the pivotal members of the committee with the highest and lowest valuations for a good reform, respectively (e.g., Dziuda and Loeper 2017) — both play a critical role in determining equilibrium outcomes. In some cases, however, the right pivot has the same influence on the course of deliberation as a committee chair who could unilaterally terminate deliberation. The cloture rule puts her in a pivotal position to trigger filibusters, entirely canceling the influence of the left pivot.

Furthermore, it follows from the unique equilibrium prediction that for any cloture rule, committee members can deduce the outcome of deliberation before the session begins. This allows us to derive their individual preferences over the family of all cloture rules, which constitute the basis for our analysis of the stability of cloture rules. Fixing an arbitrary voting rule for reforming cloture rules, we provide a full characterization of the set of (core-)stable cloture rules, i.e., those rules that induce equilibrium outcomes that no decisive coalition is able or willing to overturn. Our analysis suggests that most rules are stable. In particular, all cloture rules satisfy a self-stability criterion, in which the reform rule coincides with the existing cloture rule. Drawing on this result, we also show that cloture rules are "dynamically stable": in a variant of our game where committee members can propose to amend the existing rule at the start of *every* round, we find that in equilibrium proposers are unable or unwilling to amend any rule throughout the session. In accord with our observations above, these results suggest that cloture rules are highly persistent, even when they generate Pareto inefficient outcomes.

Related literature. Deliberation is also modeled as a collective-learning game in Chan et al. (2018), where information arrives continuously according to a Weiner process without an (exogenous) time limit. This contrasts with our model where the central assumption is that the final vote must be taken before a fixed deadline, allowing us to study the occurrence and impact of filibusters. Their study and this one are thus complementary in their objectives: while Chan et al. (2018) mostly focus on the length of deliberation and the accuracy of the decision, we examine the incentives created by the deadline and, in particular, the risk of obstruction.

As our model is one of collective learning, what we do is closely related to the growing literature on collective experimentation, pioneered by Strulovici (2010). In those models, committee members also collective learn about the types of feasible policies, but learning about the merits of a given policy requires that it be implemented. In particular, the informational structure of our model is the same as in Anesi and Bowen (forthcoming), who study the impact of redistributive constraints on policy experimentation. In their benchmark case without redistribution, the left and right pivots of the voting rule *both* play a critical role in the characterization of equilibrium outcomes, which can be Pareto inefficient. In our model, collective deliberation gives rise to situations where the left pivot is in effect the only decisive voter and to different forms of Pareto inefficiency. Other more distantly related papers on collective experimentation include Callander (2011), Messner and Polborn (2012), Acemoglu et al. (2015), Callander and Harstad (2015), and Freer et al. (forthcoming).

Another related but more distant literature studies the stability of voting rules and political institutions in contexts different from collective learning and cloture — e.g., Barberà and Jackson (2004), Lagunoff (2009), and Acemoglu et al. (2012). We use stability concepts similar to those developed by these authors when we turn to the issue of persistence of cloture rules in Section 4.

There is an abundance of theoretical research on collective deliberation. Perhaps the predominant perspective comes from an extensive family of cheap-talk models of information transmission, in which the main question of interest is whether deliberation can serve as a vehicle for sharing and eliciting the committee members' private information — see Landa and Meirowitz (2009) for a review. Like us, other work emphasizes different aspects of collective deliberation, focusing instead on non-Bayesian models of argumentation (Hafer and Landa 2007, and Patty 2008), legitimacy (Patty and Penn 2011, 2014), coalitions and agenda formation (Anesi and Seidmann 2014), or the comparative analysis of various modes of deliberation (Chung and Duggan 2020).

Finally, this paper is also related to the political science literature on legislative obstruction, recently reviewed by Wawro and Schickler (2010). These papers do not consider collective learning.

The paper proceeds as follows. After presenting our collective-deliberation framework in the next section, we provide a characterization of the unique equilibrium outcome and discuss its implications in Section 3. In Section 4, we analyze the stability of cloture rules. We conclude in Section 5, and collect in the appendix the proofs of the results not included in the main text.

2 Deliberation Framework

There is a finite set $N \equiv \{1, \ldots, n\}$ of committee members, with $n \geq 2$, who deliberate about whether a safe status quo policy, S, should be changed to a risky reform, R. The reform is either good or bad. Committee member i places value $v_i > 0$ on a good reform, and 0 on a bad reform; the safe status quo S gives each member a common benefit of $s > 0$. For expositional ease, we assume that $v_1 > \cdots > v_n$.

As do Chan et al. (2018), we model deliberation as a collective-learning process. The initial probability that the risky reform R is good is given by $p_0 \in (0,1)$. Committee members update their (common) belief about R's type participate over the course of a deliberative session that may last at most $T \in \mathbb{N}$ rounds, indexed $t = 0, \ldots, T-1$. We denote by b_t the common belief at the start of round t. There are two possible outcomes at the end of a deliberation round, good news or no news. The probability of good news is $\lambda \in (0,1)$ if the reform is good, and the probability of good news is 0 if it is bad. It follows that the first arrival of good news reveals to all committee members that R is good, and hence the common belief updates to one (i.e., $b_t = 1$). In the event that no news has been received after t rounds of deliberation, the belief b_t is equal to

$$p_t \equiv \frac{p_0(1-\lambda)^t}{p_0(1-\lambda)^t + 1 - p_0} \ . \tag{1}$$

Two types of votes are held over the course of the session: *cloture votes* over whether to terminate deliberation and make a decision, and a *final vote* over whether to change the status quo to the reform. More specifically, at the start of round $t \in \{0, 1, \ldots, T-1\}$, all committee members vote simultaneously to cloture or continue deliberation. Let $\mathcal{D} \subseteq 2^N \setminus \{\emptyset\}$ be the collection of decisive coalitions for cloture votes. If a coalition $C \in \mathcal{D}$ of committee members vote in favor of cloture, then deliberation is terminated and a final vote must be held; otherwise, the committee continues to deliberate until the next cloture vote at the start of round $t+1$. In a final vote, all committee members vote simultaneously to pass the reform R or reject it in favor of the status quo S, thus ending the game. The committee's choice is to implement the reform if a coalition $C \in \overline{\mathcal{D}}$ of committee members vote in favor of R, and to maintain the status quo otherwise, where $\overline{\mathcal{D}} \subseteq 2^N \setminus \{\emptyset\}$ is the collection of decisive coalitions for the final vote.

We only impose minimal and common conditions on the cloture rule $\mathcal{D}$ (e.g., Austen-Smith and Banks, 2000). We assume that it is monotonic, e.g., $C \in \mathcal{D}$ and $C \subseteq C'$ imply $C' \in \mathcal{D}$. We further assume that it is proper, e.g., $C, C' \in \mathcal{D}$ implies $C \cap C' \neq \emptyset$. Common examples include *quota rules* ($\mathcal{D} = \{C \subseteq N \colon |C| \geq q\}$, where $n/2 < q \leq n$) and *dictator-*

ships ($\mathcal{D} = \{C \subseteq N : C \ni i\}$ for some $i \in N$). In particular, dictatorships correspond to situations where there is a committee chair, who has the authority to call a vote on the issue under consideration. The same conditions are imposed on the voting rule $\overline{\mathcal{D}}$.

If all the cloture votes fail until the end of round $T-1$, then session ends without a final vote and the status quo is maintained. Every round of deliberation imposes a cost of $c > 0$ to each committee member. Thus, if the reform is voted up in round t and the belief is $b_t \in \{p_t, 1\}$, then committee member i's expected payoff is $b_t v_i - tc$; if the reform is voted down in round t, then her payoff is $s - tc$; and if no final vote is taken by the end of the session, then each member receives $s - Tc$.

We study sequential equilibria in pure strategies. As is common in voting games, we require the committee members' strategies to be stage-undominated (Baron and Kalai, 1993), thus eliminating implausible equilibria in which, for example, under majority rule everyone votes in favor of the reform independently of payoffs. In order to limit the number of possible cases (without affecting the paper's conclusions), we further assume that in case of a tie, committee members prefer to end rather than continue deliberation in cloture votes, and to maintain rather than change the status quo in the final vote. Henceforth, we will refer to any sequential equilibrium that satisfies these conditions as simply an equilibrium. (We provide a formal definition of an equilibrium in the first section of the appendix.)

3 Collective Deliberation and Cloture

Our first series of results concerns the impact of fixed cloture rules on the outcomes of collective deliberation.

3.1 Preliminary Analysis

In this subsection, we conduct some preliminary analysis to prepare the foundations for the presentation of our main results. Equilibrium characterization for the final-vote stage is provided along with some analysis of individual preferences over deliberation policies.

3.1.1 The Final Vote

Before we can turn to the main subject of this paper — i.e., the collective decision to terminate deliberation — we must consider how the committee responds in the final vote to the cloture of deliberation. Suppose the final vote is taken in round $t = 0, \ldots, T-1$, so that the committee members' belief that the reform is good is $p_t < 1$ if no news has been revealed during the deliberation rounds, and is equal to one otherwise. It follows from our tie-breaking rule that in equilibrium, each member i votes to change the status quo S to the reform R if and only if $p_t v_i > s$. By the monotonicity of the voting rule $\overline{\mathcal{D}}$ (and the assumption that $v_1 > \cdots > v_n$), the committee's decision is then entirely determined by the choice of its member $\bar{r} \equiv \min\big\{j \in N \colon \{1, \ldots, j\} \in \overline{\mathcal{D}}\big\}$: the reform is adopted if and only if $p_t v_{\bar{r}} > s$.

For future reference, let $\bar{t} \equiv \min\{t \in \mathbb{N} \colon p_t \leq s/v_{\bar{r}}\}$, so that we have the following:

Observation 1. *In the round-t cloture vote, unless good news has been revealed, the outcome of the final vote is R if $t < \bar{t}$, and S if $t \geq \bar{t}$.*

To avoid trivialities, we make the following assumption throughout:

A1. $p_0 v_{\bar{r}} > s > p_{T-1} v_{\bar{r}}$.

This assumption implies that, conditional on receiving no news, the committee adopts the reform (resp. maintains the status quo) if the final vote is taken in the first (resp. final) stage. This in turn implies that a decisive coalition of committee members in $\overline{\mathcal{D}}$ prefer the good reform to the status quo.

Finally, observe that there is some non-monotonicity in the preferences over cloture: the committee members with the highest valuations wish to cease deliberating first before $\bar{t}$, but last after $\bar{t}$.

3.1.2 Pro-reform Committee Chairs

To help gain some intuition about the incentives underlying committee members' choices in cloture votes, it is useful to consider first the particular case in which the cloture rule

$\mathcal{D}$ is dictatorial, i.e., there is a committee member i such that $\mathcal{D} = \{C \subseteq N : i \in C\}$. This case can be interpreted as a situation in which member i is the committee chair, who has the power to unilaterally call a final vote.

Suppose i supports the reform if it is good, i.e., $v_i > s$. To characterize her ideal deliberation policy, observe first that if she ever chooses to cloture deliberation before round $\bar{t}$ while still uncertain about the reform's type, then she does so in the first round. Specifically, waiting until round $t \in (0, \bar{t})$ does not have any implication for the final voting outcome, as R will be the chosen policy regardless of whether good news has arrived, but this creates a deliberation cost of $tc > 0$ to i. Hence, if she does not call a vote in the first round then, unless the committee receives good news, she will not call a vote before round $\bar{t}$.

Every new round of deliberation involves an additional cost of c but also generates some information that i uses to update her beliefs. It is routine to show that, from round $\bar{t}$ onward, the solution to this tradeoff is a *stopping rule.* That is, there is a *cutoff* $\hat{t}_i$ such that it is optimal for i to continue deliberation until round $\hat{t}_i$ unless she receives good news before that round, in which case she immediately calls a vote; and to end deliberation for S to be selected in round $\hat{t}_i$.

To formally define the optimal cutoff $\hat{t}_i$, let $V_i(b_t, t)$ denote i's continuation value at the start of any round $t \geq \bar{t}$ that begins with a belief $b_t \in \{p, 1\}$. As $v_i > s$, we must have $V_i(1, t) = v_i$ for all t. If i chooses to continue deliberation when the belief is $b_t = p_t < 1$, then her continuation payoff is equal to $\lambda p_t v_i + (1 - \lambda p_t) V_i(p_{t+1}, t+1) - c$; if she calls a vote, her continuation payoff is s. Therefore, we have

$$V_i(p_t, t) = \max\left\{\lambda p_t v_i + (1 - \lambda p_t) V_i(p_{t+1}, t+1) - c, s\right\} ,$$

for all $\bar{t} \leq t \leq T - 2$. Evidently, $V_i(p_{T-1}, T-1) = s$: it is never optimal to deliberate in round $T - 1 \geq \bar{t}$, as the status quo will be maintained regardless. The optimal cutoff $\hat{t}_i$ is the first round t in which the belief p_t is smaller than or equal to the belief that would make i indifferent between continuing and ending deliberation. Setting $V_i(p_{\hat{t}_i}, \hat{t}_i) = s$, we obtain

$$\hat{t}_i = \min\left\{\min\left\{t \geq \bar{t} : p_t \leq \frac{c}{\lambda(v_i - s)}\right\}, T - 1\right\} .$$

Let $\Delta_i(t)$ denote the difference between i's payoff when she calls a vote in round 0 and her payoff when she adopts the stopping rule with cutoff $t \in \{1, \ldots, T-1\}$, i.e.,

$$\begin{aligned}\Delta_i(t) &\equiv p_0 v_i - \left[p_0\left[\left(1-(1-\lambda)^{t-1}\right)v_i + (1-\lambda)^{t-1}s\right] + (1-p_0)(s-ct) - p_0 c\frac{1-(1-\lambda)^t}{\lambda}\right] \\ &= p_0(1-\lambda)^t\left(v_i - s - \frac{c}{\lambda}\right) + (1-p_0)ct - (1-p_0)s + p_0\frac{c}{\lambda}\ , \end{aligned} \tag{2}$$

where the second equality is obtained by using equation (1) and rearranging terms. It follows from the arguments above that in equilibrium, the chair i calls a vote in round 0 if $\Delta_i(\hat{t}_i) \geq 0$, and adopts the stopping rule with cutoff $\hat{t}_i$ otherwise.

3.1.3 Anti-reform Committee Chair

When characterizing the equilibrium outcome for the dictatorial cloture rule $\mathcal{D} = \{C \subseteq N : i \in C\}$, we assumed that $v_i > s$. Now suppose instead that $v_i < s$, that is, the committee chair i prefers the status quo to a good reform. In this case, the chair's decision in rounds $t \geq \bar{t}$ is trivial: she ends deliberation, thus preserving the status quo. In round $t < \bar{t}$, however, ending deliberation leads to the adoption of the reform. Hence, she must either let the committee adopt the reform without further deliberation costs, or wait until round $\bar{t}$ to call a final vote. The latter alternative is risky: with positive probability, the committee may learn that the reform is good between rounds t and $\bar{t}$, in which case R becomes the only possible outcome of a final vote.

At the arrival of good news before round $\bar{t}$, the chair is left with only two options. One is to cloture deliberation immediately, accepting R and avoiding the cost of deliberation; the other is to let the session end without a final vote, thereby preventing to committee from adopting the reform. We refer to the latter action as a *filibuster,* in reference to the well-known procedure used to prevent proposals from being brought to a vote in legislative bodies. This alternative is optimal for i if deliberation is not too costly and/or v_i is sufficiently small relative to s. Specifically, she is better off filibustering rather than ending deliberation in round $t < \bar{t}$ whenever $s - (T-t)c > v_i$.

Let t_f^i be the last round before it becomes profitable for i to filibuster, i.e.,

$$t_f^i \equiv T - \left\lceil \frac{s - v_i}{c} \right\rceil\ . \tag{3}$$

It follows from the arguments above that there is a unique equilibrium outcome when i is an anti-reform committee chair. If $t_f^i > \bar{t} - 1$, then this equilibrium resembles those described in the previous subsection: it is optimal for the chair to let the committee adopt the reform in the first round if $\Delta_i(\bar{t}) \geq 0$ (where $\Delta_i(\cdot)$ is defined as in (2)), and to adopt the stopping rule with cutoff $\hat{t}_i \equiv \bar{t}$ otherwise. However, the equilibrium outcome for the case where $t_f^i \leq \bar{t} - 1$ may not be a stopping rule, as it may be optimal for the chair to adopt the following deliberation policy: she waits until round $\bar{t}$ to cloture deliberation; if good news arrives in some round $t < t_f^i$, then she ends deliberation in the next round; and if good news arrives in round $t \in \{t_f^i, \dots, \bar{t} - 1\}$, she lets the session end without a final vote. Calculations reveal that she is better off adopting this strategy than cloturing deliberation in the first round if and only if $\Delta_i^f < 0$, where

$$\Delta_i^f \equiv p_0 c \left[(1-\lambda)^{t_f^i}(T - t_f^i) - (1-\lambda)^{\bar{t}}(T - \bar{t}) + \frac{1 - (1-\lambda)^{t_f^i}}{\lambda} \right] + p_0(1-\lambda)^{t_f^i}(v_i - s) + (1 - p_0)(c\bar{t} - s) \ . \quad (4)$$

We record the results in this and the previous subsection in the following observation.

Observation 2. *Suppose $\mathcal{D} = \{C \subseteq N \colon i \in C\}$ for some $i \in N$. There is a unique equilibrium outcome:*

(i) Suppose $t_f^i \geq \bar{t} - 1$. Then the committee chair, i, clotures deliberation at $t = 0$ if $\Delta_i(\hat{t}_i) \geq 0$, and adopts the stopping rule with cutoff $\hat{t}_i$ otherwise, where

$$\hat{t}_i \equiv \begin{cases} \min\left\{\min\left\{t \geq \bar{t} \colon p_t \leq \frac{c}{\lambda(v_i - s)}\right\}, T-1\right\} & \text{if } v_i > s \ , \\ \bar{t} & \text{otherwise.} \end{cases}$$

(ii) Suppose $t_f^i < \bar{t} - 1$. Then she clotures deliberation at $t = 0$ if $\Delta_i^f \geq 0$, and adopts the following strategy otherwise: she clotures deliberation in round $\bar{t}$; unless good news arrives in some round $t < t_f^i$, i which case she ends deliberation in the next round; or unless good news arrives in round $t \in \{t_f^i, \dots, \bar{t} - 1\}$, in which case she lets the session end without a final vote.

3.2 Equilibrium Characterization

In this subsection, we draw on the intuitions developed above to characterize the equilibrium outcome for arbitrary cloture rules. Except where explicitly stated, we take the collection of decisive coalitions $\mathcal{D}$ to be any nonempty family of coalitions that satisfies the monotonicity and properness conditions imposed in Section 2. As we will see, while some equilibrium patterns of deliberation are consistent with Observation 2, there are possibilities that arise with general cloture rules that do not have counterparts in the dictatorial-rule case.

Let $\ell(\mathcal{D})$ and $r(\mathcal{D})$ denote the *left* and *right pivots* of $\mathcal{D}$, respectively (e.g., Dziuda and Loeper 2017); i.e,

$$\ell(\mathcal{D}) \equiv \min\left\{i \in N\colon \{i+1,\dots,n\} \notin \mathcal{D}\right\} \text{ , and } r(\mathcal{D}) \equiv \max\left\{j \in N\colon \{1,\dots,j-1\} \notin \mathcal{D}\right\} .$$

Throughout this section, the cloture rule is treated parametrically and to ease the notation, we will often suppress the dependence of ℓ and r on $\mathcal{D}$, with the understanding that they depend on a specific rule. Note that, as $\mathcal{D}$ is proper, $\ell \leq r$. For example, if n is odd and $\mathcal{D}$ is simple majority rule, then the left and right pivots are the same: $\ell = r = (n+1)/2$. More generally, for every quota rule with quota $q > n/2$, we have $\ell = n - q + 1$ and $r = q$. In the dictatorial case of Observation 2, $\ell = r = i$.

Let $\hat{t}_\ell$ and t_f^r be defined as in Observation 2 and in equation (3), respectively; i.e.,

$$\hat{t}_\ell \equiv \begin{cases} \min\left\{\min\left\{t \geq \bar{t}\colon p_t \leq \frac{c}{\lambda(v_\ell - s)}\right\}, T-1\right\} & \text{if } v_\ell > s \text{ ,} \\ \bar{t} & \text{otherwise.} \end{cases}$$

and

$$t_f^r \equiv T - \left\lceil \frac{s - v_r}{c} \right\rceil \text{ .}$$

It is useful to divide the analysis into three cases:

(i) Case 1: Suppose first that $\mathcal{D}$ satisfies $\hat{t}_\ell \leq t_f^r$. If the reform is already known to be good in round $T-1$, committee member i votes in favor of ending deliberation either if she prefers a good reform to the status quo or if it is still too costly for her to filibuster (so that $t_f^i \geq T-1$). As t_f^i is decreasing in i (recall that $v_1 > \cdots > v_n$), committee member r's

decision is pivotal: if $t_f^r \geq T-1$, then the decisive coalition $\{1, \ldots, r\}$ clotures deliberation and the reform is adopted in the ensuing final vote; otherwise, the blocking coalition $\{r, \ldots, n\}$ prevents the committee from adopting the reform by letting the session end without a final vote. Applying the same argument recursively to every round $t > 0$ in which the belief is equal to one, we obtain the same outcome for any such round: deliberation is terminated and the reform adopted if $t \leq t_f^r$, and a filibuster occurs otherwise.

Consider now cloture votes in which the committee has not yet received good news. We ignore for now the possibility of a filibuster, assuming that $t_f^r \geq T-1$. In the last round $T-1 \geq \bar{t}$, cloture is unanimously agreed in equilibrium. As the status quo will be maintained regardless, there is no point to an additional (costly) round of deliberation. It follows that in the cloture vote of round $T-2$, each committee member i faces the same decision problem as the committee chair in the previous subsection. In particular, if $\hat{t}_i \leq T-2$, then she votes in favor of taking the final vote. Committee members who place lower values on a good reform prefer to end deliberation earlier than the left pivot, i.e., $\hat{t}_i \leq \hat{t}_\ell$ for all $i \geq \ell$. Therefore, if $\hat{t}_\ell \leq T-2$, then the decisive coalition $\{\ell, \ldots, n\}$ ends deliberation. Applying the same logic recursively to all rounds $t \in \{\hat{t}_\ell, \ldots, T-2\}$, we obtain that in equilibrium, the committee never deliberates beyond round $\hat{t}_\ell$. Combined with the cloture rule's property that $t_f^r \geq \hat{t}_\ell$, this implies that the possibility of a filibuster cannot have any impact on the equilibrium outcome (on the path).

Conditional on receiving no news, in round $\hat{t}_\ell - 1$, the left pivot is better off continuing deliberation for an additional round, as are all the members of the blocking coalition $\{1, \ldots, \ell\}$. Unless it learns that the reform is good, the latter coalition will then prevent the committee from taking the final vote in this round and, by the same logic, in all rounds $t \in \{\bar{t}, \ldots, \hat{t}_\ell - 1\}$. Therefore, conditional on receiving no news, deliberation will end in round $\hat{t}_\ell$.

As explained in Subsection 3.1.2, if it is optimal for any committee member i to end deliberation and take a vote in any round $t \in \{2, \ldots \bar{t} - 1\}$, then it is also optimal for her to do so in the first round. Therefore, either the committee takes the final vote in the first round or it adopts the stopping rule with cutoff $\hat{t}_\ell$. Since $\Delta_i(t)$ is decreasing in i for all t

(recall (2)), the right pivot is the pivotal committee member in this stage: if $\Delta_r(\hat{t}_\ell) \geq 0$, then r and consequently all the members of the decisive coalition $\{1, \ldots, r\}$ vote to cloture deliberation in round 0; otherwise, r and consequently all the members of the blocking coalition $\{r, \ldots, n\}$ opt for the left pivot's optimal stopping rule. Note in passing that this stopping rule is the ideal stopping rule of the left pivot, ℓ.

(ii) Case 2: Now suppose that $\mathcal{D}$ satisfies $\bar{t} \leq t_f^r < \hat{t}_\ell$. In every round $t \geq t_f^r+1$, every member of the coalition $\{r, \ldots, n\}$ is prepared to filibuster to ensure that the reform will never be adopted. As the status quo will be maintained irrespective of the belief about the reform's type, all committee members agree to take a final vote so as to avoid pointless deliberation.

The characterization of equilibrium behavior in rounds $t \leq t_f^r$ parallels case 1, with t_f^r playing the role of the cutoff $\hat{t}_\ell$: the blocking coalition $\{1, \ldots, \ell\}$ opposes cloture from round $\bar{t}$ to round t_f^r unless the committee learns that the reform is good, in which case the final vote is held and the reform is adopted; and in round 0, the committee effectively chooses between ending deliberation immediately and adopting the stopping rule with cutoff t_f^r. As above, the outcome of the round-0 cloture vote is determined by the right pivot's preferences; i.e., deliberation is immediately ended if and only if $\Delta_r(t_f^r) \geq 0$. Note that in contrast to the previous case, the possibility of a filibuster off the equilibrium path has a significant impact on collective deliberation on the path. Note also that, even though the right pivot's valuation, v_r, determines the cutoff value of t_f^r, it may not be r's ideal cutoff. Indeed, with $v_r < s$, r would prefer to end deliberation as soon as possible — that is, at $\bar{t}$, rather than at $t_f^r \geq \bar{t}$.

(iii) Case 3: Finally, suppose that $t_f^r < \bar{t}$. By the same logic as before, the committee unanimously approves cloture in every round $t \geq \bar{t}$, as every member realizes that maintaining the reform is the only possible outcome.

Now suppose $\bar{t} - 1 \geq t_f^r$. Suppose further that the committee is still uncertain about the type of the reform at the start of round $\bar{t} - 1$. If it chooses to end deliberation in this round, then each member i's expected payoff is equal to $p_{\bar{t}-1} v_i$. If it chooses to continue deliberating, then it learns that the reform is good with probability $\lambda p_{\bar{t}-1}$ and receives no news with probability $1 - \lambda p_{\bar{t}-1}$. In the former case, the blocking coalition $\{r, \ldots, n\}$

filibusters and i's payoff is $s - (T - \bar{t} + 1)c$; in the latter case, the committee votes to maintain the reform in the next round, so that i's continuation payoff is $s - c$. Therefore, she votes in favor of cloture in round $\bar{t} - 1$ if and only if

$$p_{\bar{t}-1} v_i \geq \lambda p_{\bar{t}-1} \big[s - (T - \bar{t} + 1)c \big] + (1 - \lambda p_{\bar{t}-1})(s - c) \ .$$

As $t_f^i \leq \bar{t} - 1$ for all $i \geq r$, we have $p_{\bar{t}-1} v_i < s - (T - \bar{t} + 1)c < s - c$ for all the members of the coalition $\{r, \ldots, n\}$, who consequently vote against cloture. It follows that deliberation continues in round $\bar{t} - 1$, giving rise to a filibuster *on the equilibrium path* with probability $\lambda p_{\bar{t}-1} > 0$. The same argument obtains in all rounds $t \in \{t_f^r, \ldots, \bar{t} - 1\}$: conditional on no news, deliberation is not clotured; and if good news arrive, the blocking coalition $\{r, \ldots, n\}$ filibusters until the end of deliberation session.

As in the previous cases, it is readily checked that if committee member i prefers to end deliberation in any round $\in \{1, \ldots, t_f^r - 1\}$, then she also does in round 0. Calculations reveal that this is the case if and only if

$$p_0 c \left[(1 - \lambda)^{t_f^r} (T - t_f^r) - (1 - \lambda)^{\bar{t}} (T - \bar{t}) + \frac{1 - (1 - \lambda)^{t_f^r}}{\lambda} \right] + p_0 (1 - \lambda)^{t_f^r} (v_i - s) + (1 - p_0)(c\bar{t} - s) \geq 0 \ .$$

As the left-hand side of the above inequality decreases with i, it follows that committee member r is again pivotal, and the equilibrium is analogous to the case in which she is a committee chair — recall Observation 2(ii). Recall from (4) that

$$\Delta_r^f \equiv p_0 c \left[(1 - \lambda)^{t_f^r} (T - t_f^r) - (1 - \lambda)^{\bar{t}} (T - \bar{t}) + \frac{1 - (1 - \lambda)^{t_f^r}}{\lambda} \right] + p_0 (1 - \lambda)^{t_f^r} (v_r - s) + (1 - p_0)(c\bar{t} - s) \ .$$

If $\Delta_r^f \geq 0$, then a final vote is held in the first round, otherwise, the committee deliberates until round $\bar{t}$ unless one of the following events occurs: (i) deliberation generates good news in round $t \leq t_f^r$, in which case it takes a vote to adopt the reform; or (ii) deliberation generates good news in round $t \in \{t_f^r + 1, \ldots, \bar{t} - 1\}$, in which case the coalition $\{r, \ldots, n\}$ filibusters.

The following proposition summarizes the results of this subsection.

Proposition 1. *For every cloture rule $\mathcal{D}$, there is a unique equilibrium outcome:*

(i) Suppose $\hat{t}_\ell \leq t_f^r$. If $\Delta_r(\hat{t}_\ell) > 0$, then the committee votes to implement R in round 0; otherwise it adopts the stopping rule with cutoff $\hat{t}_\ell$.

(ii) Suppose $\bar{t} \leq t_f^r < \hat{t}_\ell$. If $\Delta_r(t_f^r) \geq 0$, then the committee votes to implement R in the first round; otherwise it adopts the stopping rule with cutoff t_f^r.

(iii) Suppose $t_f^r < \bar{t}$. If $\Delta_r^f \geq 0$, then the committee votes to implement R in round 0; otherwise it votes to maintain S in round $\bar{t}$, unless

(iii.a) deliberation generates good news in round $t < t_f^r$, in which case it votes to implement R in round $t+1$, or

(iii.b) deliberation generates good news in round $t \in \{t_f^r, \ldots, \bar{t}-1\}$, in which case it deliberates until the session ends without holding a vote.

3.3 Implications

Institutional versus real power. An immediate implication of Proposition 1 is that the relevant actors on whom to focus when discussing the impacts of cloture rules are the left and right pivots, ℓ and r. Inspection of the proposition indeed reveals that the equilibrium outcome induced by any cloture rule $\mathcal{D}$ coincides with the equilibrium outcome induced by the oligarchic rule $\mathcal{D}^o \equiv \{C \subseteq N : \ell, r \in C\}$. Thus, *the left and right pivots are, in effect, oligarchs in the cloture process.* If follows that in cases where there is a single pivot (i.e., $\ell = r$), such as simple majority rule with an odd number of committee members, the cloture rule induces the same equilibrium outcome as the dictatorial rule under which the pivot is the dictator. In such a case, the pivot is effectively as powerful as if she were the committee chair.

The observations above are reminiscent of standard spatial models of voting in which a committee chooses a real number from the one-dimensional line. As the comparison of the equilibrium outcome in Proposition 1(iii) to the committee-chair benchmark in Observation 2(ii) illustrates, however, the logic is different. *In the case of "filibuster-prone cloture rules" (case 3 above), the right pivot has always the same influence on the deliberative outcome* *as a committee chair, even if there are multiple pivots.* In equilibrium, her pivotal position to trigger a filibuster suppresses the influence of the left pivot on collective deliberation.

Filibusters and Pareto inefficiency. In effect, the right to continue deliberation carries with it a right to filibuster, and conflicting policy preferences mean that filibustering is sometimes individually optimal even though the implied deliberative outcomes are Pareto inefficient. At the arrival of good news in round $t \in (t_f^r, \bar{t})$, all committee members would be better off taking a vote to immediately maintain the status quo rather than pointlessly deliberating until the end of the session. However, the members of the coalition $\{1, \ldots, \bar{r}\} \in \overline{\mathcal{D}}$ cannot commit to vote down the reform if a final vote were held in round t, as they prefer the good reform to the status quo. The only way to maintain the status quo is therefore to filibuster; and, since $t > t_f^r$, this is the optimal alternative for the members of coalition $\{r, \ldots, n\}$ (which is blocking in the cloture vote).

This form of Pareto inefficiency, created by filibustering incentives, is absent from existing models of collective learning. A different form of inefficiency arises in some models of collective experimentation where, because of the endogeneity of the status quo, members of the coalition $\{r, \ldots, n\}$ fear that experimentation would go on far too long and consequently prefer not to experiment at all, although some experimentation would be Pareto improving — e.g., Anesi and Bowen (forthcoming). In our model, if member r is better off bringing the reform to a vote in round 0 than adopting a stopping rule with cutoff $\hat{t}_\ell$ or t_f^r, so is member $\ell \leq r$ — recall that $\Delta_i(t)$ is decreasing in i, for all t. We thus have $\Delta_\ell(\hat{t}_\ell) \geq \Delta_\ell(t_f^r) \geq 0$, so that taking a vote in round 0 is the left pivot's ideal deliberative outcome and, therefore, cannot be Pareto improved. It follows that the circumstances in which cloture rules cause Pareto inefficiencies are confined to those in which filibusters occur on the equilibrium path.

Deliberation patterns. As conventional wisdom would predict, changes to the cloture rule that reduce the amount of agreement required to terminate deliberation may prevent inefficient filibusters because they make anti-reform minorities less powerful. Indeed, let $\mathcal{D}$ be a rule under which filibusters occur with positive probability in equilibrium; and let $\{j,\ldots,n\}\neq\emptyset$ be the anti-reform coalition — i.e., those committee members whose valuations are smaller than s — that may block cloture, so that $r(\mathcal{D})\geq j$. Consider the impact of an (exogenous) reform changing $\mathcal{D}$ to some $\mathcal{D}'\supset\mathcal{D}$, thus enabling more coalitions to end deliberation. (For example, if $\mathcal{D}$ is a quota rule, then this can be done simply by reducing the cloture quota.) Such a reform decreases the right pivot; and, in particular, if $r(\mathcal{D}')<j$ then it makes it impossible for $\{j,\ldots,n\}$ to filibuster.

But Proposition 1 shows that more is true: *such a reform reduces the likelihood of filibusters, even if the anti-reform coalition remains blocking.* Indeed, even if $r(\mathcal{D}')\geq j$, the decrease in $r(\cdot)$ means that the right pivot's valuation for the reform is higher and, consequently, her incentives to filibuster are lowered: $t_f^{r(\mathcal{D}')}\geq t_f^{r(\mathcal{D})}$. This in turn shrinks the "filibuster interval" $\big[t_f^{r(\cdot)},\bar{t}-1\big]$ in which the commitment failure described above is likely to occur. It follows that the equilibrium probability of a filibuster decreases.

More generally, the web of functional relationships in the equilibrium makes general comparative static results on the cloture rule difficult to obtain. Common wisdom would suggest that institutional changes that enlarge $\mathcal{D}$ (by set inclusion) can only reduce the expected length of deliberation. After all, enabling more coalitions to end deliberation should ease cloture. To see whether this intuition is correct, suppose for example that $\bar{t}\leq t_f^{r(\mathcal{D})}<\hat{t}_{\ell(\mathcal{D})}$ and $\Delta_{r(\mathcal{D})}\big(t_f^{r(\mathcal{D})}\big)<0$, so that the committee adopts the stopping rule with cutoff $t_f^{r(\mathcal{D})}$ in equilibrium (Proposition 1(ii)). How is the expected length of deliberation affected if $\mathcal{D}$ is changed to $\mathcal{D}'\supset\mathcal{D}$, with $r(\mathcal{D}')<r(\mathcal{D})$? In cases where the increase in the right pivot's valuation $v_{r(\cdot)}$ is not sufficient for $\Delta_{r(\cdot)}\big(t_f^{r(\cdot)}\big)$ to become positive and for $t_f^{r(\cdot)}$ to exceed $\hat{t}_{\ell(\cdot)}$, the expected length of deliberation *increases*. To avoid a filibuster, the committee unanimously agrees to cloture deliberation at the cutoff $t_f^{r(\mathcal{D}')}$, which is larger than $t_f^{r(\mathcal{D})}$.[4]

[4]Increases in the expected length of deliberation can also be obtained in cases where $\hat{t}_\ell(\mathcal{D})\leq t_f^{r(\mathcal{D})}$ for some parametric configurations.

4 Stability of Cloture Rules

In this section, we take up the question of the persistence of cloture rules, i.e., how vulnerable they are to reforms by committee members. Committees typically choose a substantial portion of their own rules, and may change them at any time (e.g., Diermeier et al. 2015). We have established in the previous section that the equilibrium outcome for the deliberation game is unique for any cloture rule $\mathcal{D}$ (and any profile of values $(v_1, \ldots, v_n)$). Therefore, for any $\mathcal{D}$, committee members can deduce the final outcome of deliberation, which is a known function of $\mathcal{D}$. This will allow us to define the committee members' preferences over cloture rules, which will constitute the basis for analyzing their incentives to amend $\mathcal{D}$.

Preferences over cloture rules. Let $\mathfrak{R}$ denote the family of feasible cloture rules, i.e., the family of nonempty collections $\mathcal{D}$ of coalitions that are both monotonic and proper; let $\mathfrak{R}_0$ denote the subfamily of cloture rules in $\mathfrak{R}$ under which the committee adopts the reform in round 0 in equilibrium; and let $\mathfrak{R}_1 \equiv \mathfrak{R} \setminus \mathfrak{R}_0$ be the subfamily of rules under which the committee deliberates before making a decision. It follows from Proposition 1 that we can define the equivalence relation $\sim$ on $\mathfrak{R}$ as follows: for all $\mathcal{D}, \mathcal{D}' \in \mathfrak{R}$, $\mathcal{D} \sim \mathcal{D}'$ if and only if the equilibrium outcomes under $\mathcal{D}$ and $\mathcal{D}'$ are the same. Let $[\mathcal{D}]$ denote the equivalence class of rule $\mathcal{D}$ relative to $\sim$; and, for every $\mathfrak{R}' \subseteq \mathfrak{R}$, let $\mathfrak{R}'/_\sim$ the quotient set of $\mathfrak{R}'$ relative to $\sim$. We can thus define for each committee member $i \in N$, a preference relation $\succeq_i$ over the equivalence classes of cloture rules: for all $[\mathcal{D}], [\mathcal{D}'] \in \mathfrak{R}/_\sim$, we have $[\mathcal{D}] \succeq_i [\mathcal{D}']$ if and only if i's equilibrium payoff under $\mathcal{D}$ is greater than or equal to her equilibrium payoff under $\mathcal{D}'$. Let $\succ_i$ denote the asymmetric part of $\succeq_i$.

Core stability and self-stability. Our first step toward a stability result is to establish that it is possible to order the equivalence classes of rules that induce deliberation beyond round 0 in such way that committee members' preferences are single-peaked over these classes. Single-peakedness over the classes that induce stopping rules in equilibrium is not surprising, as using the equilibrium cutoffs is a natural way to order such rules. Extending the result to the other classes of cloture rules in $\mathfrak{R}_1/_\sim$ is more subtle.

Lemma 1. *There exists a linear order $\trianglerighteq$ on $\mathfrak{R}_1/_\sim$ such that all committee members' preferences on $\mathfrak{R}_1/_\sim$ are single-peaked with respect to $\trianglerighteq$.*

When studying the stability of cloture rules in $\mathfrak{R}$, the relevant question is: Given the current cloture rule $\mathcal{D}$, does there exist a consensus for change, that is, does there exist a decisive coalition that prefers some other rule? Before answering this question, we face a decision as to what we want the term "decisive coalition" to mean for the reform of cloture rules. We discuss the issue concerning which collection of decisive coalitions is appropriate below when we introduce the notion of self-stability. But, for now, in order to obtain a general characterization of stability, we do not impose any restriction on the collection of decisive coalitions.

Formally, let $\mathbf{D}$ be any subset of $2^N \setminus \{\emptyset\}$ that is monotonic and proper; and let $\boldsymbol{\ell} \equiv \min\big\{i \in N\colon \{i+1,\dots,n\} \notin \mathbf{D}\big\}$ and $\mathbf{r} \equiv \max\big\{j \in N\colon \{1,\dots,j-1\} \notin \mathbf{D}\big\}$ denote the left and right pivots of $\mathbf{D}$. We will refer to $\mathbf{D}$ as the *reform rule.* Say that a cloture rule $\mathcal{D} \in \mathfrak{R}$ is $\mathbf{D}$-*stable* if for all $\mathcal{D}' \in \mathfrak{R}$, $\big\{i \in N\colon [\mathcal{D}'] \succ_i [\mathcal{D}]\big\} \notin \mathbf{D}$. In other words, the $\mathbf{D}$-stable cloture rules are those that cannot be overturned by decisive coalitions in $\mathbf{D}$. This is the standard notion of core stability, typically used in positive political theory (e.g., Austen-Smith and Banks 2000).

In the next proposition, we use the property of individual preferences over cloture rules described in Lemma 1 to obtain a general characterization of $\mathbf{D}$-stable rules for any $\mathbf{D}$.

Proposition 2. *Let $[\mathcal{D}_{\boldsymbol{\ell}}]$ and $[\mathcal{D}_{\mathbf{r}}]$ be the $\succeq_{\boldsymbol{\ell}}$- and $\succeq_{\mathbf{r}}$-maxima of $\mathfrak{R}_1/_\sim$, respectively; and let*

$$\mathfrak{R}^* \equiv \big\{\mathcal{D} \in \mathfrak{R}_1\colon [\mathcal{D}_{\boldsymbol{\ell}}] \trianglerighteq [\mathcal{D}] \trianglerighteq [\mathcal{D}_{\mathbf{r}}] \text{ and } [\mathcal{D}] \succeq_{\mathbf{r}} \mathfrak{R}_0\big\} .$$

Then, every cloture rule $\mathcal{D} \in \mathfrak{R}_1$ is $\mathbf{D}$-stable if and only if $\mathcal{D} \in \mathfrak{R}^*$. Furthermore, every cloture rule in $\mathfrak{R}_0$ is $\mathbf{D}$-stable if and only if $\mathfrak{R}_0 \succeq_{\boldsymbol{\ell}} [\mathcal{D}_{\boldsymbol{\ell}}]$.

The set $\mathfrak{R}^*$ in Proposition 2 is reminiscent of the gridlock interval in the spatial theory of voting, where voters have single-peaked preferences over the one-dimensional real line: the linear order obtained in Lemma 1 allows us to treat the cloture rules in $\mathfrak{R}_1$ (or equivalently, their equivalence classes) as points on a line. It then follows from the single-peakedness property established above that stable cloture rules must belong to a "gridlock interval" whose endpoints are the left and right pivots' ideal rules in $\mathfrak{R}_1$. One key difference, however, is that stability of any rule $\mathcal{D}$ in this interval also requires that no decisive coalition in $\mathbf{D}$ can profitably change it to a no-deliberation cloture rule in $\mathfrak{R}_0$. It is readily checked that the latter condition is satisfied if and only if the right pivot $\mathbf{r}$ prefers $\mathcal{D}$ to no deliberation. In this case, the other members of the blocking coalition $\{\mathbf{r}, \ldots, n\}$, whose valuations for the reform are smaller than $\mathbf{r}$'s, are also better off deliberating under cloture $\mathcal{D}$ than adopting the risky reform without further information. Conversely, the no-deliberation rules are stable if and only if the left pivot of $\mathbf{D}$, $\boldsymbol{\ell}$, prefers to adopt the reform without deliberating to any other outcome — as do all the other members of the blocking coalition $\{1, \ldots, \boldsymbol{\ell}\}$, whose valuations for the reform are higher than $\boldsymbol{\ell}$'s.

The criterion of $\mathbf{D}$-stability can be weakened by choosing a smaller family of decisive coalitions, $\mathbf{D}' \subset \mathbf{D}$, and requiring only that cloture rules be $\mathbf{D}'$-stable. Alternatively, the stability criterion can be strengthened by requiring that it hold for a larger family of decisive coalitions. What is the appropriate criterion? This question does not have a single answer, as a variety of reform rules can be justified according to different examples of how actual procedural reforms take place in committees. One minimal restriction one should impose on the reform rule, however, is that it be as coarse as the exiting cloture rule itself, i.e., $\mathbf{D} \subseteq \mathcal{D}$. Put differently, the coalitions $\mathbf{D}$ should also be decisive coalitions in $\mathcal{D}$. Since in essence, the coalitions in $\mathbf{D}$ have the power to instigate institutional changes, one should expect these coalitions to also have the power to bring deliberation to an end. This leads us to use Barberà and Jackson's (2004) concept of self-stability:[5] a cloture rule $\mathcal{D}$ is *self-stable* if it is $\mathcal{D}$-stable. Evidently, if $\mathcal{D}$ is self-stable, then it is also $\mathbf{D}$-stable for all $\mathbf{D} \subseteq \mathcal{D}$.

Proposition 3. *Every cloture rule in $\mathfrak{R}$ is self-stable.*

[5]Barberà and Jackson (2004) study societies' decisions to amend their constitutions — not committees that can change their own rules *at any time* — and, therefore, they apply self-stability before agents learn their valuations for the reform.

The intuition for this result is as follows. An implication of the structure of the equilibrium characterized in Proposition 1 is that every cloture rule $\mathcal{D}$ in $\mathfrak{R}_1$ must lie "between" the ideal rules of its own left and right pivots in $\mathfrak{R}_1$, i.e., $[\mathcal{D}_\ell] \trianglerighteq [\mathcal{D}] \trianglerighteq [\mathcal{D}_r]$. Moreover, since $\mathcal{D} \in \mathfrak{R}_1$, the members of $\{r, \ldots, n\}$ must prevent the committee members with higher valuations from cloturing deliberation in round 0 in equilibrium; hence, $[\mathcal{D}] \succeq_r \mathfrak{R}_0$. As $\ell = \boldsymbol{\ell}$ and $r = \mathbf{r}$ in the self-stability criterion, it follows that $\mathcal{D}$ belongs to $\mathfrak{R}^*$ and, therefore, is self-stable (Proposition 2). In the case where $\mathcal{D} \in \mathfrak{R}_0$, the members of the decisive coalition with the highest valuations for the reform prefer to adopt it without deliberation in equilibrium. In particular, this implies that the right pivot of $\mathcal{D}$ always prefers $\mathcal{D}$ to any cloture rule in $\mathfrak{R}_1$ and, therefore, so does the left pivot $\boldsymbol{\ell} = \ell < r$. It then follows from Proposition 2 that $\mathcal{D}$ is self-stable. These intuitions are made precise in the proof of Proposition 3.

Dynamic stability. In the above we assumed that the committee could only amend its cloture rule at the start of the session. This is a reasonable assumption in cases where the deliberation time T available for the committee is short. In many cases of interest, however, time for deliberation is plentiful, sometimes months or even years. In such cases, it makes more sense to also allow committee members to amend procedural rules during the course of a "session."

To study the persistence of cloture rules in such contexts, we take the bargaining approach to dynamic stability developed by Acemoglu et al. (2012), i.e., we augment each round of the model introduced in Section 2 with a rule-making phase. Specifically, each round t begins with a status-quo rule $\mathcal{D}^{t-1}$, inherited from the previous round. Nature first selects a finite list of "proposers" $(\pi_1, \ldots, \pi_m)$ (possibly with repetition) according to some fixed (possibly history-dependent) distribution on $\big\{(\pi_1, \ldots, \pi_m) \colon m \in \mathbb{N} \; \& \; \pi_\ell \in N \text{ for each } \ell = 1, \ldots, m\big\}$. Then, proposer π_1 makes the first proposal $\mathcal{D} \in \mathfrak{R}$;[6] once the proposal is made, committee members vote sequentially (in an arbitrary order) over whether to accept it. The proposal is accepted if a coalition $C \in \mathcal{D}^{t-1}$ of committee

[6] Proposing $\mathcal{D} = \mathcal{D}^{t-1}$ is equivalent to passing.

members vote to accept, and it is rejected otherwise. If it is accepted, then events unfold as in the basic model under cloture rule $\mathcal{D}^t = \mathcal{D}$; if it is rejected, proposer π_2 is called upon to make a proposal and the same process is repeated. If the m proposers all make unsuccessful proposals, then the status-quo rule is maintained, i.e., $\mathcal{D}^t = \mathcal{D}^{t-1}$. If the committee does not terminate deliberation in round t, then the game transitions to round $t+1$ with status-quo rule $\mathcal{D}^t$. Rule-making phases take and negligible amount of time. The *initial rule* $\mathcal{D}^0$ is exogenously given.

In this framework, we say that a cloture rule $\mathcal{D} \in \mathfrak{R}$ is *dynamically stable* if there exists a sequential equilibrium of the game with initial rule $\mathcal{D}^0 = \mathcal{D}$ in which $\mathcal{D}$ is never amended on the equilibrium path. In other words, a cloture rule is dynamically stable if proposers are unable or unwilling to amend it throughout the session.

Our next result is an easy implication of the analysis above; although technically obvious, the corollary has substantive implications for rule persistence. Observe that Proposition 3 holds regardless of the initial belief p_0 and the horizon T.[7] Applying the proposition to any subgame starting at $T-1$ (and a backward-induction argument), we therefore obtain that the status-quo rule $\mathcal{D}^{T-2}$ is never amended in an equilibrium where indifferent committee members reject proposals to amend $\mathcal{D}^{T-2}$. It follows that committee members anticipate that the rule they adopt at the start of round $T-2$ will remain in place until the end of the session. We can then apply Proposition 3 to any subgame starting at $T-2$ to obtain that the status-quo rule $\mathcal{D}^{T-3}$ will not be amended on the path. The same argument can be applied recursively to yield the following corollary. (Note that it holds regardless of the protocol used to select proposers and the order in which committee members vote over proposals.)

Corollary 1. *Every cloture rule in* $\mathfrak{R}$ *is dynamically stable.*

We conclude from the results in this section that cloture rules are highly persistent, even when they generate Pareto inefficient outcomes. Strategic anticipation of the use of these rules always leads a blocking coalition of committee members to reject amendment

[7] Proposition 3 was established for the nontrivial cases in which assumption A1 holds. It readily checked that the result applies in the other cases.

proposals. Of course, our model does not provide a complete story of the persistence of cloture rules: real-world committees do occasionally change their cloture rules. After all, the adoption of a certain rule $\mathcal{D}'$ can itself be interpreted as changing $\mathcal{D} = \{N\}$ to some $\mathcal{D}' \supset \mathcal{D}$. There are presumably many other factors that can explain why committees choose to maintain or amend existing rules, but our focus here is on the collective-learning dimension of deliberation. Proposition 3 and Corollary 1 suggest that the reasons for rule changes must lie elsewhere.

5 Discussion

We have developed a model of collective deliberation and learning that explicitly incorporates a cloture vote in the collective decision-making process. Its simplicity notwithstanding, the model generates new insights and predictions about the impacts and stability of cloture rules, emphasizing the strategic incentives created by limited deliberation times. More generally, the analysis demonstrates that important aspects of the deliberative process cannot be fully understood without taking account of its collective-learning dimension.

Our simple and stylized model of collective learning abstracts from all but the bare essentials necessary to illustrate our main ideas. This said, there are of course other important aspects of collective deliberation worth exploring. We close the paper by discussing a few possible extensions of our model, which we regard as particularly important for further work. In our model, each round deliberation only involves a fixed cost, which we think of as an opportunity cost. However, the acts of producing information, and communicating and absorbing its contents require effort, and hence are subject to moral hazard in teams (Dewatripont and Tirole, 2005). This clearly creates new strategic incentives, since the decision to terminate deliberation will depend on the committee members' expectations about their future incentives to generate information, which in turn will depend on the expected course of deliberation.

We say nothing here about argumentation. Reforms are typically proposed by one or several members of the committee, who may possess private information and seek to use it to their own advantage. By strategically transmitting information about their proposals to the other members over the course of deliberation, they can influence the collective learning process. To the best of our knowledge, the interaction between information transmission and collective learning in deliberative committees has so far remained unexplored.

Our analysis of rule stability has mainly focused on dynamic stability, where committees can change their rules at any time. As do Barberà and Jackson (2004) in the case of constitutional rules, one could consider ex ante stability, where the reform to be deliberated on, and therefore the members' valuations for that reform, are still unknown at the time of rule change. Committee members would decide on a cloture rule that would then be applied to several reforms, perhaps, chosen at random. Preliminary analysis suggests that not all cloture rules are ex ante stable. In particular, the members' degree of risk aversion seems to be a key factor: as they become more risk averse, they prefer to have more restrictive cloture rules, i.e., they prefer having a higher chance of filibustering rather than risking that an unfavorable reform be implemented.

Finally, like all the existing papers on collective learning, we have assumed that each committee member's valuation for the reform is known to the other members. In the case of committees that meet infrequently, such as ad hoc committees, it would make sense to assume that policy preferences are private information. We leave these, and other extensions, for future work.

Appendix

A Equilibrium Definitions

We denote a typical "state" as (b_t, t, d_t), where t is the round of deliberation, $b_t \in \{p_t, 1\}$ is the current belief in round t, and $d_t \in \{0, 1\}$ denotes whether deliberation continues ($d_t = 1$) or not ($d_t = 0$). For $t \leq T - 1$, the state $(b_t, t, 1)$ can either transition to $(b_t, t, 0)$ if deliberation is clotured in period t, or transition to $(b_{t+1}, t+1, 1)$ if deliberation is not clotured. If deliberation is not clotured, belief $b_t = 1$ changes to $b_{t+1} = 1$ with probability 1, while $b_t = p_t$ changes to $b_{t+1} = 1$ with probability λp_t and to $b_{t+1} = p_{t+1}$ with probability $1 - \lambda p_t$. We let $cl_i(b_t, t) \in \{0, 1\}$ denote member i's cloture vote on whether to stop deliberation ($cl_i = 1$) or not ($cl_i = 0$) in state $(b_t, t, 1)$. Moreover, we let $f_i(b_t, t) \in \{0, 1\}$ denote i's final vote on whether to pass the reform ($f_i = 1$) or not ($f_i = 0$), in state $(b_t, t, 0)$.

We describe i's voting strategy s_i as a cloture vote $cl_i(b_t, t, 1)$, and a final vote $f_i(b_t, t, 0)$, for all possible states (b_t, t, d_t). Given the (pure) strategies of all committee members, and the collections of decisive coalitions $\mathcal{D}, \overline{\mathcal{D}}$, one can identify the voting outcome, as well as calculate the continuation payoff of any member i, at any state (b_t, t, d_t).

Definition A1. *Given voting strategies of committee members, define*
(i) $V_i(b_t, t)$ *as* i*'s expected continuation payoff in state* $(b_t, t, 1)$*, without counting any costs of deliberation spent before* t*; and*
(ii) $F_i(b_t, t)$ *as* i*'s expected payoff in state* $(b_t, t, 0)$*, without including any costs of deliberation.*

That is, $V_i(b_t, t)$ is calculated at the beginning of period t, before committee members cast their cloture votes. Only the cost of deliberation from t onwards is included in $V_i(b_t, t)$. Moreover, $F_i(b_t, t)$ is i's expected payoff in state $(b_t, t, 0)$, as the final vote is held.

Having defined continuation payoffs of the committee members in the various states, we can now introduce the definition of the sequential equilibrium with iterative elimination of weakly dominated strategies:

Definition A2. *In an* equilibrium, *committee member* i *votes as follows:*
(i) For $t \leq T - 1$, $f_i(b_t, t) = 1$ *if* $b_t v_i > s$, *and* $f_i(b_t, t) = 0$ *otherwise;*
(ii) For $t \leq T - 1$, $cl_i(1, t) = 0$ *if* $V_i(1, t + 1) - c > F_i(1, t)$, *and* $cl_i(1, t) = 1$ *otherwise. Respectively,* $cl_i(p_t, t) = 0$ *if* $\lambda p_t V_i(1, t + 1) + (1 - \lambda p_t) V_i(p_{t+1}, t + 1) - c > F_i(1, t + 1)$, *and* $cl_i(p_t, t) = 1$ *otherwise.*

That is, during the final vote, any committee member i votes for R if and only if i strictly prefers R to S. During the cloture vote, i votes for continuing deliberation if and only if she receives a strictly higher continuation payoff. Note that $V_i(b_T, T) = s$: If no decision is made in rounds $t \in \{0, ..., T - 1\}$, then the status quo is maintained.

B Proof of Lemma 1

We begin by defining the linear order $\trianglerighteq$ on $\mathfrak{R}_1/_\sim$. To this end, we first define the function $\mathtt{T}\colon \mathfrak{R}_1/_\sim \to \{1, \ldots, T\}$ as follows: if the equilibrium outcome induced by equivalence class $[\mathcal{D}]$ is a stopping rule with cutoff τ, then $\mathtt{T}\big([\mathcal{D}]\big) = \tau$; if the equilibrium outcome under $[\mathcal{D}]$ is as in Proposition 1(iii), then $\mathtt{T}\big([\mathcal{D}]\big) = t_f^{r(\mathcal{D})}$. We then define $\trianglerighteq$ as: for all $\mathcal{D}_1, \mathcal{D}_2 \in \mathfrak{R}$, $[\mathcal{D}_1] \trianglerighteq [\mathcal{D}_2]$ if and if $\mathtt{T}\big([\mathcal{D}_1]\big) \geq \mathtt{T}\big([\mathcal{D}_2]\big)$.

It is readily checked that each committee member i has single-peaked preferences over the set of $[\mathcal{D}]$'s such that $\mathtt{T}\big([\mathcal{D}]\big) \geq \bar{t}$ — see, e.g., Anesi and Bowen's (forthcoming) Proposition 1 (they establish single-peakedness over the cutoff belief but, as p_t is strictly decreasing in t, the same applies to stopping times). Now suppose that

$$t_f^i \equiv T - \left\lceil \frac{s - v_i}{c} \right\rceil \geq \bar{t} - 1 \ , \tag{B1}$$

so that committee member i is better off deliberating until $\bar{t}$ than risking a filibuster before $\bar{t}$. It follows that her ideal equivalence class $[\mathcal{D}_i]$ satisfies $[\mathcal{D}_i] \trianglerighteq [\mathcal{D}]$ for all $[\mathcal{D}]$ such that $\mathtt{T}\big([\mathcal{D}]\big) < \bar{t}$. To show that her preferences are single-peaked over $\mathfrak{R}_1/_\sim$, therefore, it suffices to show that they are "increasing" over the $[\mathcal{D}]$'s such that $\mathtt{T}\big([\mathcal{D}]\big) < \bar{t}$, in the sense that $[\mathcal{D}_2] \trianglerighteq [\mathcal{D}_1]$ implies $[\mathcal{D}_2] \succeq_i [\mathcal{D}_1]$. Take any two such classes $[\mathcal{D}_1]$ and $[\mathcal{D}_2]$, and suppose without loss of generality that $[\mathcal{D}_2] \trianglerighteq [\mathcal{D}_1]$. They induce the same equilibrium payoffs to i conditional on the following events happening: the committee receives good news before $\mathtt{T}\big([\mathcal{D}_1]\big)$, or after $\mathtt{T}\big([\mathcal{D}_2]\big)$, or never. However, i receives a higher payoff conditional on good news being received between $\mathtt{T}\big([\mathcal{D}_1]\big)$ and $\mathtt{T}\big([\mathcal{D}_2]\big)$: under $[\mathcal{D}_1]$, deliberation continues till the end of the session; under $[\mathcal{D}_2]$, deliberation stops and the reform is implemented. From (B1), i strictly prefers the latter outcome to the former. This shows that i's preferences are single-peaked over $\mathfrak{R}_1/_\sim$.

We now turn to the case where (B1) does not hold. In this case, the ideal deliberation plan for i would be to adopt the reform if good news is revealed by t_f^i, filibuster if good news is revealed between $t_f^i + 1$ and $\bar{t} - 1$, and stop deliberating if there is no news by $\bar{t}$. But this must be the equilibrium outcome under the rule, say $\mathcal{D}_i$, where i is the committee chair. Hence, she prefers $[\widehat{\mathcal{D}}_i]$ to all the other equivalence classes of cloture rules. Moreover, by the same logic as above, we can check that $\succeq_i$ is "increasing" over the $\mathcal{D}$'s that satisfy $\mathtt{T}([\mathcal{D}]) < \mathtt{T}([\widehat{\mathcal{D}}_i])$, and "decreasing" over the $\mathcal{D}$'s that satisfy $\mathtt{T}([\mathcal{D}]) > \mathtt{T}([\widehat{\mathcal{D}}_i])$; so that i's preferences are single-peaked over $\mathfrak{R}_1/_\sim$.

C Proof of Proposition 2

An immediate implication of Lemma 1 above is that the only possible candidates for **D**-stability are $\mathfrak{R}_0$ and all the equivalent classes $[\mathcal{D}]$ in $\mathfrak{R}_1/_\sim$ such that $[\mathcal{D}_{\boldsymbol{\ell}}] \trianglerighteq [\mathcal{D}] \trianglerighteq [\mathcal{D}_{\mathbf{r}}]$. But it is readily checked that if committee member i prefers $\mathfrak{R}_0$ to some $\mathtt{T}^{-1}(t)$ for $t \geq \bar{t}$, so do all members $j < i$. By definition of the right pivot, therefore, the equivalent classes in $\mathfrak{R}^*$ (if $\mathfrak{R}^* \neq \emptyset$) must be **D**-stable.

Now suppose that the left pivot (under **D**) prefers $\mathfrak{R}_0$ to her ideal equivalence class in $\mathfrak{R}_1/_\sim$, $[\mathcal{D}_\ell]$. This implies that she also prefers $\mathfrak{R}_0$ over all the other $[\mathcal{D}]$'s, and so do all the committee members j with $v_j > v_{\boldsymbol{\ell}}$. As $\{1, \ldots, \boldsymbol{\ell}\}$ is a blocking coalition, $\mathfrak{R}_0$ is **D**-stable. Finally, suppose that $\boldsymbol{\ell}$ strictly prefers $[\mathcal{D}_{\boldsymbol{\ell}}]$ to $\mathfrak{R}_0$. Then, all committee members $j > \boldsymbol{\ell}$ also strictly prefers $[\mathcal{D}_{\boldsymbol{\ell}}]$ to $\mathfrak{R}_0$. As $\{\boldsymbol{\ell}, \ldots, n\} \in \mathbf{D}$, it follows that $\mathfrak{R}_0$ is not **D**-stable.

D Proof of Proposition 3

Let $\mathcal{D}$ be an arbitrary cloture rule. To show that it is $\mathcal{D}$-stable, it suffices from Proposition 2 to show that either $[\mathcal{D}] \in \mathfrak{R}^*$, or that $\mathcal{D} \in \mathfrak{R}_0 \succeq_\ell [\mathcal{D}_\ell]$.

Suppose first that $\hat{t}_\ell \leq t_f^r$, so that the equilibrium outcome is either (i) the optimal stopping rule of the left pivot, or (ii) adoption of the reform in round 0 — recall Proposition 1. In case (i), $\mathcal{D} \in [\mathcal{D}_\ell]$ and the right pivot prefers the stopping rule with cutoff $\mathtt{T}([\mathcal{D}])$ over no deliberation, i.e., $[\mathcal{D}] \succeq_r \mathfrak{R}_0$. Hence, $[\mathcal{D}] \in \mathfrak{R}^*$. In case (ii), it must be the case that the right pivot prefers no deliberation to the left pivot's ideal stopping rule. As $\Delta_i(\hat{t}_\ell)$ is decreasing in i, this must also be true for all members j such that $v_j > v_r$, including the left pivot. Hence, $\mathcal{D} \in \mathfrak{R}_0 \succeq_\ell [\mathcal{D}_\ell]$.

Suppose now that $\bar{t} \leq t_f^r < \hat{t}_\ell$, so that the equilibrium outcome is either (i) the stopping rule with cutoff t_f^r, or (ii) adoption of the reform in round 0. In case (i), we have $\mathtt{T}([\mathcal{D}_r]) = \bar{t}$

(since the right pivot prefers the status quo to a good reform and $\Delta_r(\bar{t}) < 0$) and $\mathtt{T}\big([\mathcal{D}_\ell]\big) = \hat{t}_\ell$, so that $[\mathcal{D}_\ell] \trianglerighteq [\mathcal{D}] \trianglerighteq [\mathcal{D}_r]$. And since $\Delta_r(t_f^r) < 0$, we have $[\mathcal{D}] \in \mathfrak{R}^*$. In case (ii), the same logic as in the previous paragraph applies again.

Finally, suppose that $t_f^r < \bar{t}$. From Proposition 1(iii), either there is deliberation in equilibrium (so that $\mathtt{T}\big([\mathcal{D}_r]\big) = t_f^r = \mathtt{T}\big([\mathcal{D}]\big)$), in which case the right pivot prefers $[\mathcal{D}] = [\mathcal{D}_r]$ to $\mathfrak{R}_0$; or there is no deliberation in equilibrium, in which case she prefers $\mathfrak{R}_0$ to $[\mathcal{D}_r]$. In the former case, we have $[\mathcal{D}] \in \mathfrak{R}^*$. In the latter case, she prefers $\mathfrak{R}_0$ to any other equivalence class, including $[\mathcal{D}_\ell]$. From the definitions of $\Delta_i(\cdot)$ and Δ_i^f, this must also be the case for all committee members $j < r$, including the left pivot ℓ. Hence, $\mathfrak{R}_0 \succeq_\ell [\mathcal{D}]$. This completes the proof of the proposition.

References:

Acemoglu, D., G. Egorov, and K. Sonin (2012) "Dynamics and Stability of Constitutions, Coalitions, and Clubs," *American Economic Review* 102, 1446-1476.

Acemoglu, D., G. Egorov, and K. Sonin (2015) "Political Economy in a Changing World," *Journal of Political Economy* 123, 1038-1086.

Anesi, V., and T.R. Bowen (forthcoming) "Policy Experimentation in Committees: A Case against Veto Rights under Redistributive Constraints," *American Economic Journal: Microeconomics.*

Anesi, V., and D.J. Seidmann (2014) "Bargaining over an Endogenous Agenda," *Theoretical Economics* 9, 445-482.

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

Barberà, S., and M.O. Jackson (2004) "Choosing How to Choose: Self-Stable Majority Rules and Constitutions," *Quarterly Journal of Economics* 119, 1011-1048.

Baron, D., and H. Kalai (1993) “The Simplest Equilibrium of a Majority Rule Division Game,” *Journal of Economic Theory* 61, 290-301.

Binder, S.A., and S.S. Smith (1998) “Political Goals and Procedural Choice in the Senate,” *Journal of Politics* 60, 398-416.

Bolton, P., and A. Faure-Grimaud (2009) “Thinking Ahead: The Decision Problem,” *Review of Economic Studies* 76, 1205-1238.

Callander, S. (2011) “Searching for Good Policies,” *American Political Science Review* 105, 643-662.

Callander, S., and B. Harstad (2015) “Experimentation in Federal Systems,” *Quarterly Journal of Economics* 130, 951-1002.

Chan, J., A. Lizzeri, W. Suen and L. Yariv (2018) “Deliberating Collective Decisions,” *Review of Economic Studies* 85, 929-963.

Chung, H., and J. Duggan (2020) “A Formal Theory of Democratic Deliberation,” *American Political Science Review* 114, 14-35.

Dewatripont, M., and J. Tirole (2005) “Modes of Communication,” *Journal of Political Economy* 133, 217-1238.

Diermeier, D., C. Prato, and R. Vlaicu (2015) “Procedural Choice in Majoritarian Organizations,” *American Journal of Political Science* 59, 866-879.

Dziuda, W., and A. Loeper (2017) “Dynamic Pivotal Politics,” *American Political Science Review* 112, 580-601.

Eidelson, B. (2013) “The Majoritarian Filibuster,” *Yale Law Journal* 122, 980-1023.

Freer, M., C. Martinelli, and S. Wang (forthcoming) “Collective Experimentation: A Laboratory Study,” *Journal of Economic Behavior & Organization.*

Hafer, C., and D. Landa (2007) “Deliberation as Self-Discovery and Institutions for Political Speech,” *Journal of Theoretical Politics* 19, 329-360.

Krehbiel, K. (2006) "Pivots." In *The Oxford Handbook of Political Economy,* ed. B. Weingast and D. Wittman. Oxford: Oxford University Press, 223-241.

Lagunoff, R. (2009) "Dynamic Stability and Reform of Political Institutions," *Games and Economic Behavior* 67, 569-583.

Landa, D., and A. Meirowitz (2009) "Game Theory, Information, and Deliberative Democracy," *American Journal of Political Science* 53, 427-444.

Messner, M., and M.K. Polborn (2012) "The Option to Wait in Collective Decisions and Optimal Majority Rules," *Journal of Public Economics* 96, 524-540.

Oleszek, W.J. (2016) "Changing the Senate Cloture Rule at the Start of a New Congress." In: CRS Report Prepared for Members and Committees of Congress.

Patty, J. (2008) "Arguments-Based Collective Choice," *Journal of Theoretical Politics* 20, 379-414.

Patty, J., and E.M. Penn (2011) "A Social Choice Theory of Legitimacy," *Social Choice and Welfare* 36, 365-382.

Patty, J., and E.M. Penn (2014). *Social Choice and Legitimacy: The Possibilities of Impossibility.* New York: Cambridge University Press.

Strulovici, B. (2010) "Learning While Voting: Determinants of Collective Experimentation," *Econometrica* 78, 933-971.

Wawro, G.J., and E. Schickler (2010) "Legislative Obstructionism," *Annual Review of Political Science* 13, 297-319.