

Voting with Random Proposers: Two Rounds May Suffice^{*}

Hans Gersbach[†]

Kremena Valkanova[‡]

This version: August 2026

Abstract

This paper introduces the Voting with Random Proposers (VRP) procedure to address the challenges of agenda manipulation in voting. In each round of VRP, a randomly selected proposer suggests an alternative that is voted on against the previous round's winner. In a framework with single-peaked preferences, we show that the VRP procedure guarantees that the Condorcet winner is implemented in a few rounds, and in just two rounds under sufficiently symmetric preference distributions, or whenever the initial status quo or the first proposer's type is not extreme. Under the equilibrium proposal strategy, sincere and sophisticated voting induce the same winner in every round. The results have applications for committee decisions, legislative decision-making, and the organization of citizens' assemblies and decentralized autonomous organizations.

Keywords: Condorcet winner, sincere voting, multi-round voting, random proposers.

JEL Classification: C72, D70, D72.

Conflicts of Interest: None.

^{*}We thank two anonymous referees and the advisory editor, as well as, Vincent Anesi, Pio Blieske, Jon X. Eguia, Sean Horan, Elias Kolleck, Markus Pivato, Clemens Puppe, Dimitrios Xefferis, Jan Zapal and workshop and seminar participants at the CEPS PolEcon Workshop in Paris 2024 and at the University of Zurich and ETH Zurich for valuable comments.

[†]KOF Swiss Economic Institute, ETH Zurich and CEPR, hgersbach@ethz.ch.

[‡]Parity Technologies, kremena@parity.io.

1 Introduction

In various collective decision-making scenarios such as committee deliberations, legislative processes, and the governance of proof-of-stake blockchains, members of the corresponding bodies face the challenge of selecting a single alternative from a large set of options. A widely used approach in such contexts is to conduct successive rounds of pairwise majority voting (Rasch, 2000). For instance, legislative bodies often refine bills by sequentially voting on amendments, and negotiating parties in multi-agent resource allocation problems typically converge on a solution through incremental steps. However, with a large number of alternatives, it becomes infeasible to consider all options comprehensively. Moreover, the final outcome is highly sensitive to the subset of alternatives chosen for voting and the sequence in which they are presented. This creates a well-known vulnerability, whereby strategic monopoly agenda setters can manipulate the decision-making process to achieve outcomes aligned with their preferences, as extensively documented in the literature.¹

To address these challenges, we introduce an iterative majority voting procedure where, in each round, a *randomly selected proposer* from the society suggests an alternative to be voted against the winner of the previous round. We refer to this method as “Voting with Random Proposers” (VRP). The random selection of proposers is designed to decentralize and democratize the agenda-setting process, thus preventing any single agent from dominating the proposal sequence. Our analysis focuses on a simplified framework with a fixed number of voting rounds, one-dimensional alternatives, and a society of agents with symmetric single-peaked preferences with publicly known distribution. Additionally, only the final winning alternative is payoff relevant to the members of the society.

Our main goal is to analyze the strategic incentives created by the VRP procedure for both voters and proposers, while assessing its efficiency in selecting the socially optimal

¹See the classic contributions by McKelvey (1979), Rubinstein (1979), Bell (1981), Schofield (1983) and the more recent contributions that have extended and applied the potential and limit of agenda power to various settings, Barberà and Gerber (2017), Nakamura (2020) and Ali, Bernheim, Bloedel, and Console Battilana (2023). We refer to Horan (2021) and Rosenthal and Zame (2022) for a recent account of the literature and to Banks (2002) for a review of the earlier literature.

Condorcet winner. We show that an optimal proposal for a proposer in all rounds except the final one, given sophisticated voting, is the Condorcet winner, unless their peak lies at the tail of the preference distribution, in which case they propose their own peak. Conditional on this proposal strategy, sincere voting, i.e. supporting whichever of the two standing alternatives a voter currently prefers, and sophisticated voting induce the same winner in every round. The procedure’s performance therefore does not depend on voters’ sophistication.

As a result, the probability of implementing the socially optimal alternative converges to one after only a few voting rounds. For balanced distributions of the preferences, where the median and mean are sufficiently close, the Condorcet winner is obtained in only two voting rounds, independently of the status quo and the proposers’ preferences. Even for arbitrary distributions of preferences in society, the Condorcet winner is selected in two rounds whenever the initial status quo or the first proposer’s type is not extreme. The main intuition is that proposal-makers who are not in the final round face the following trade-off: if they propose a policy at or near their own peak, they risk that the final outcome will shift to the opposite side of the median. This risk can be avoided by proposing the median peak instead.

Our results demonstrate that current iterative voting procedures can be optimized by engaging randomly selected proposal-makers, thereby increasing the likelihood of reaching the Condorcet winner or at least come as close as possible. The VRP procedure has a wide range of potential applications, not only within standard collective decision-making bodies, such as committees and legislative bodies, but also in citizens’ assemblies that prepare proposals for city councils or parliaments. In this context, it may be optimal to select two randomly chosen citizen groups to develop proposals, which the assembly would then vote on, advancing the winning alternative to the next decision-making body. Moreover, this procedure could be effectively employed to decide on initiatives in direct democracies and can be easily adapted for decentralized autonomous organizations operating with distributed ledger technologies.

2 Relation to Literature

Since the seminal result of [Moulin \(1980\)](#), it is well known that in the case of single-peaked preferences, the Condorcet winner can be selected by asking agents to announce their peak, adding a number of fixed peaks, and using a suitable generalized median rule.² This result has been considerably extended by [Border and Jordan \(1983\)](#), [Barberà and Jackson \(1994\)](#), and [Klaus and Protopapas \(2020\)](#). Unlike their mechanism design approach, our paper focuses on an alternative implementation problem: There is no central authority with commitment, i.e. a mechanism designer. The procedures can only involve proposal-making by agents and voting with equal proposal-making and agenda rights. Such procedures are sometimes called democratic mechanisms (see [Gersbach \(2009\)](#)) and also mirror common practices in parliamentary decision-making in representative democracies and referenda in direct democracies. In particular, we focus on procedures that present agents with two alternatives at a time.

Our paper also relates to the work of [Miller \(1977\)](#), who demonstrates that a Condorcet winner is selected through sophisticated and potentially strategic voting, regardless of the agenda, as long as it includes the Condorcet winner. Similarly, [Austen-Smith \(1987\)](#) shows that strategic voting and sincere voting may be observationally equivalent since agenda-setters will propose options that can defeat the last winning proposal under sincere voting. We extend these insights by showing that the random selection of agenda-setters in an iterative majority voting process creates strong incentives for them to propose the Condorcet winner, even when they are self-interested and strategic. Under the equilibrium proposal strategy, sincere and sophisticated voting induce the same winner in every round, so implementation does not require voters to coordinate or to know each other's types.

Legislative bargaining models commonly feature random proposer selection and iterative

²[Moulin \(1980\)](#) has shown that any strategy-proof social choice function on the domain of single-peaked preferences is a generalized median-rule that selects the median of the voters' ballots and $n - 1$ fixed ballots, where n is the number of voters. However, not all generalized median rules select the Condorcet winner.

majority voting on proposals, a framework originating from the seminal work of [Baron and Ferejohn \(1989\)](#). While much of the subsequent literature has focused on bargaining with an exogenous status quo³, another line of research, notably [Anesi and Seidmann \(2014\)](#), explores dynamic bargaining models with an evolving status quo, where the status quo is determined by the outcome of the bargaining in the previous round.⁴ Our model features several important deviations from these frameworks, including a finite number of voting rounds, a one-dimensional policy space⁵, and a single payoff-relevant outcome. Within this context, our model can be classified as a spatial legislative bargaining model featuring a finite horizon and an evolving status quo, where the winning alternative of each round becomes the status quo for the next round but is only implemented at the final period, akin to amendment procedures in legislatures ([Eraslan et al., 2022](#)). We obtain moderation due to the threat of future polarization: If I know that my ideological opponent will make a proposal tomorrow, I have an incentive to constrain her through my proposal today. In the presence of a decisive median voter, the way to impose such a constraint is to propose a policy that appeals to the median.

Our paper also relates to the literature on strategic polarization. [Kalai and Kalai \(2001\)](#) demonstrate that strategic behavior within aggregation games tends to foster polarization. In comparison, our work illustrates that well-crafted institutional mechanisms, such as the Voting with Random Proposers procedure, can promote moderation and facilitate convergence toward compromise, more exactly the Condorcet winner. The intuition behind this contrast is the following. Polarization arises when each player directly influences the aggregate outcome. VRP, however, counters this by removing direct aggregation and instead using randomized sequential proposals with majority voting. In another work, [Eraslan, Evdokimov, and Zápal \(2020\)](#) show that equilibrium outcomes evolve gradually through repeated

³See [Eraslan and Evdokimov \(2019\)](#) for a review.

⁴This literature has been recently surveyed in [Eraslan, Evdokimov, and Zápal \(2022\)](#).

⁵In this respect, our paper relates to the literature on spatial bargaining, including [Baron \(1991\)](#), [Banks and Duggan \(2000\)](#), [Kalandrakis \(2016\)](#), and [Zapal \(2016\)](#).

majority voting, but the process may take many rounds. In comparison, our investigations demonstrate that randomizing proposers aligns incentives so that the Condorcet winner emerges almost immediately, often within just two rounds. Thus, while the former relies on gradual stabilization through voting dynamics, the latter achieves rapid convergence through institutional design of the proposal stage.

Our research on voting with random proposers over a large set of alternatives also relates to the literature on algorithmically determined and iterative voting procedures with a random component in various policy spaces, as developed by [Airiau and Endriss \(2009\)](#), [Goel and Lee \(2012\)](#) and [Garg, Kamble, Goel, Marn, and Munagala \(2019\)](#), among others. We add to this literature by studying a simple procedure to implement the Condorcet winner.

Finally, when designing iterative voting procedures with endogenous proposal-making, it is crucial to identify alternatives in each round that are closer to the Condorcet winner.⁶ In this sense, our work contributes to the literature by [Callander \(2011\)](#), [Riboni and Ruge-Murcia \(2010\)](#), and [Barberà and Gerber \(2022\)](#), which examines how to identify desirable policies to put to a vote. We demonstrate that the sequential random selection of proposers plays a pivotal role in guiding proposers toward making socially superior proposals.

3 The Model

Consider a society that consists of a continuum of agents with mass one, each of them privately informed about their type $\theta \in [0, 1]$. Agents' types are drawn from a publicly known cumulative distribution function F . We assume throughout that F is absolutely continuous on $[0, 1]$ with density f , and that its median θ_m is uniquely determined, i.e. $F(\theta_m) = \frac{1}{2}$ and F is strictly increasing in a neighborhood of θ_m . These conditions guarantee that the median peak is well defined and that majorities in a neighborhood of θ_m are decisive.

The utility of an agent of type θ from an alternative $x \in [0, 1]$ is given by $u_\theta(x) =$

⁶We use reaching the Condorcet winner as the objective of designing voting procedures.

$-(x - \theta)^2$, therefore it has a unique maximum at θ and is symmetric around θ . Since preferences are single-peaked, the median peak θ_m is the Condorcet winner (Black, 1948).

We define a function $c : [0, 1] \rightarrow [0, 1]$, $x \mapsto \min\{\max\{2\theta_m - x, 0\}, 1\}$. Thus, $c(x)$ is the farthest positioned alternative in $[0, 1]$ that is at least as preferred as x by a majority of society. Note in particular that whenever $x \in (\max\{0, 2\theta_m - 1\}, \min\{2\theta_m, 1\})$, we have $F\left(\frac{x+c(x)}{2}\right) = F(\theta_m) = \frac{1}{2}$.

Furthermore, we define lower and upper thresholds as the boundaries of the regions in which a proposer strictly prefers the lottery induced by their own peak to the certain outcome θ_m :

$$\begin{aligned}\underline{\theta} &= \sup \left(\{0\} \cup \left\{ \theta \in [0, 1] : \theta \leq 2\theta_m - 1, \int_{\theta}^1 u_{\theta}(v)f(v) dv > u_{\theta}(\theta_m) \right\} \right), \\ \bar{\theta} &= \inf \left(\{1\} \cup \left\{ \theta \in [0, 1] : \theta \geq 2\theta_m, \int_0^{\theta} u_{\theta}(v)f(v) dv > u_{\theta}(\theta_m) \right\} \right).\end{aligned}$$

Intuitively, a type below $2\theta_m - 1$, the lower-tail comparison is between the certain utility $u_{\theta}(\theta_m)$ and $\int_{\theta}^1 u_{\theta}(v)f(v) dv$, the expected utility when the alternative is drawn from F but truncated below at the proposer's peak. The upper-tail comparison is symmetric. Types strictly below $\underline{\theta}$, or strictly above $\bar{\theta}$, strictly prefer the lottery induced by their own peak; at a threshold, and whenever the corresponding strict-preference region is empty, we select θ_m . When $\theta_m < \frac{1}{2}$, no type satisfies $\theta \leq 2\theta_m - 1$; at $\theta_m = \frac{1}{2}$ the only admissible point is $\theta = 0$. Thus $\underline{\theta} = 0$ whenever $\theta_m \leq \frac{1}{2}$; symmetrically, when $\theta_m > \frac{1}{2}$, $\bar{\theta} = 1$.

We observe that the remaining threshold also collapses when the most extreme type on the relevant side weakly prefers the median peak to the lottery induced by its own peak: $\underline{\theta} = 0$ when $\theta_m > \frac{1}{2}$ and

$$\int_0^1 u_0(v)f(v) dv \leq u_0(\theta_m),$$

and, symmetrically, $\bar{\theta} = 1$ when $\theta_m \leq \frac{1}{2}$ and

$$\int_0^1 u_1(v)f(v) dv \leq u_1(\theta_m).$$

We require F to satisfy the following additional assumptions. For $h \in [0, H]$, where $H := \min\{\theta_m, 1 - \theta_m\}$,

$$q(h) := 1 - F(\theta_m + h) - F(\theta_m - h), \quad s(h) := 1 - F(\theta_m + h) + F(\theta_m - h).$$

Both are built from the two tail masses $1 - F(\theta_m + h)$ (above $\theta_m + h$) and $F(\theta_m - h)$ (below $\theta_m - h$): $q(h)$ is their *difference*, a signed measure of tail imbalance that is positive when the upper tail is the heavier of the two at h , and $s(h) \geq 0$ is their *sum*, the combined mass in both tails. Our first restriction signs the tail imbalance in the aggregate.

Assumption 1. The distribution of types satisfies

$$\int_0^{\theta_m} q(h) dh \geq 0 \quad \text{if } \theta_m \leq \frac{1}{2}, \quad \text{or} \quad \int_0^{1-\theta_m} q(h) dh \leq 0 \quad \text{if } \theta_m > \frac{1}{2}.$$

The inequality compares the two tails in the aggregate: for $\theta_m \leq \frac{1}{2}$ it holds when the upper tail is at least as heavy as the lower one, ruling out a disproportionate mass near the lower boundary. As we show in Lemma 1, this is exactly what places the mean on the appropriate side of the median.⁷

Lemma 1. Under Assumption 1, $\mathbb{E}(\theta) - \theta_m \geq \int_{2\theta_m}^1 [1 - F(v)] dv \geq 0$, if $\theta_m \leq \frac{1}{2}$, and $\mathbb{E}(\theta) - \theta_m \leq -\int_0^{2\theta_m-1} F(v) dv \leq 0$ if $\theta_m > \frac{1}{2}$.

Proof. See Appendix A. □

Assumption 1 constrains only the *direction* of the imbalance, that is, the sign of its integral. In addition, we impose a bound on the *magnitude* of the accumulated imbalance relative to the accumulated dispersion.

⁷A more classical sufficient condition is to require $q(h)$ to have constant sign for all h rather than only after integration. This is equation (1.2) of [van Zwet \(1979\)](#). Our Assumption 1 (its integrated form) is equation (1.1). We use the weaker integral condition because it is all the results require.

Assumption 2. For all $h \in [0, H]$,

$$\int_0^h r s(r) dr \geq -(\theta_m - \underline{\theta}) \int_0^h q(r) dr, \quad \text{and} \quad \int_0^h r s(r) dr \geq (\bar{\theta} - \theta_m) \int_0^h q(r) dr.$$

Assumption 1 signs the single quantity $\int_0^h q$ at the outer scale $h = H$ and so fixes the direction of the skewness, Assumption 2 governs the entire profile $\{\int_0^h q\}_{h \leq H}$ and bounds its size relative to the dispersion. The two are complementary, and neither implies the other. Both hold with room to spare when F is close to symmetric about θ_m (if F is exactly symmetric then $\int_0^h q \equiv 0 \leq \int_0^h r s$, so both are slack) and Assumption 2 can fail only for distributions whose tails are strongly imbalanced relative to their spread.

Voting with Random Proposers

The society uses a multi-round voting procedure called Voting with Random Proposers to implement an alternative from the interval $[0, 1]$ in a publicly known fixed number of voting rounds $T \in \{2, 3, \dots\}$. Suppose that initially, there is a status quo alternative $q^1 \in [0, 1]$. In round $t = 1$, a proposer of type θ_{s^1} is randomly drawn from the distribution F and makes a proposal $p^1 \in [0, 1]$. Then, a majority-vote between p^1 and q^1 takes place, where ties are resolved in favor of the proposal. All voters in the society take part in each round of voting. The winner of the first voting round w^1 becomes the status quo in the next round, i.e. $w^1 = q^2$. In round $t = 2$, a new proposer is randomly drawn and the process is repeated. The procedure ends at round T and the winner of the last round w^T is implemented. Proposers and voters are assumed to be strategic, i.e. they would choose a proposal or vote in favour of alternatives in order to maximize the proximity of the final winner to their type.

Equilibrium concept

The VRP procedure defines a finite-horizon dynamic Bayesian game. In principle, the appropriate solution concept is Perfect Bayesian Equilibrium (PBE). However, under the as-

assumptions of the model, the beliefs are simple and remain history independent and the equilibrium analysis reduces to finding Subgame Perfect Nash Equilibria (SPNE) for the following reasons. Although types are privately known and the game is formally one of incomplete information, no payoff-relevant variable is ever uncertain. A player's payoff depends only on the implemented alternative, which depends on the type profile only through the distribution F . Since F is common knowledge and we adopt the standard convention that the empirical distribution of types in the nonatomic population coincides with F , so that the realized median equals θ_m , the entire payoff-relevant environment is common knowledge. A proposal may reveal the proposer's own type, but that type is payoff-irrelevant to every other agent, is uninformative about any other type by independence, and, having measure zero, leaves F unchanged. The proposal-signaling channel of Banks (1990, 1993) is therefore absent. Likewise, no round outcome can revise beliefs about θ_m , so the voter-signaling channel of Piketty (2000) and Razin (2003) is absent as well. Beliefs about payoff-relevant variables thus coincide with the prior at every information set, and PBE reduces to SPNE.⁸

Because an individual voter is non-pivotal in a nonatomic electorate, SPNE alone does not discipline individual ballots. We therefore impose a *pivotal-voting refinement*: at every binary voting information set, each voter chooses the alternative that gives her the higher continuation utility under the assumption that her ballot determines the winner. We call this *sophisticated voting*. This refinement is the nonatomic analogue of the pivotal-voter logic used in finite voting games (Baron & Kalai, 1993).

An SPNE satisfying the pivotal-voting refinement is called an equilibrium in the remainder of the paper. This refinement evaluates each ballot in the decision problem conditional on pivotality; no unconditional dominance claim is made for an individual voter in the nonatomic game. A voter votes *sincerely* in such an equilibrium if she supports whichever of the two standing alternatives she currently prefers, i.e. a voter of type θ votes for p^t whenever

⁸With finitely many voters the realized median is no longer common knowledge, and proposals and votes become informative about it. The analysis would then require Perfect Bayesian Equilibrium with explicit beliefs and would face the multiplicity issues documented in this literature.

$|\theta - p^t| < |\theta - q^t|$ and for q^t whenever $|\theta - p^t| > |\theta - q^t|$. A set of voters is *decisive* in a round if the alternative it supports wins regardless of the remaining votes: it must have mass at least $\frac{1}{2}$ when it supports the proposal and mass strictly greater than $\frac{1}{2}$ when it supports the status quo.

4 Equilibrium Proposals and Implementation

In this section, we study the optimal proposal of the proposers, depending on their type, the distribution of preferences, and the status quo, as well as the outcome of the VRP procedure. As we show in the theorem below, the equilibrium strategy for all proposers in rounds $t \in \{1, 2, \dots, T-1\}$ is identical under sophisticated voting, and they either propose their own type or the Condorcet winner. In the last round of voting, the proposer proposes either their own type or the farthest positioned alternative that wins in a pairwise vote against the status quo.

Theorem 1. *Let Assumptions 1 and 2 hold. Given sophisticated voting, the following rule selects an optimal proposal for every proposer. In rounds $t \in \{1, 2, \dots, T-1\}$*

$$p^t = \begin{cases} \theta_{s^t} & \text{if } \max\{q^t, \theta_{s^t}\} < \underline{\theta} \text{ or } \min\{q^t, \theta_{s^t}\} > \bar{\theta}, \\ \theta_m & \text{else.} \end{cases} \quad \text{and } p^T = \begin{cases} \min\{\theta_{s^T}, c(q^T)\} & \text{if } q^T \leq \theta_m, \\ \max\{\theta_{s^T}, c(q^T)\} & \text{if } q^T > \theta_m. \end{cases}$$

Under this proposal strategy, sincere and sophisticated voting induce the same winner in every round.

Proof. See Appendix B. □

Proposers face a trade-off between proposing their own type to maximize their utility and proposing the Condorcet winner to insure themselves against unfavorable final winners. The closer a proposer's type is to the Condorcet winner, the stronger the incentive to propose it. The share of such voters increases with the symmetry of the distributions of the types.

However, in sufficiently skewed distributions, the Condorcet winner is positioned closer to one extreme, leading proposers with preference peaks at the tail of the distribution to favor their own type. This is the case, because the potential benefit from having all subsequent proposers positioned at the tail outweighs the insurance gained by proposing the Condorcet winner, which is nearly as disliked as the extreme opposite alternatives.⁹ The optimal strategy of the proposers is illustrated in Figure 1, depending on the distribution of the types.




Figure 1: The first row depicts the selected optimal proposals in rounds $t \in \{1, 2, \dots, T-1\}$ for different Beta distributions of types. Types in the green regions propose the Condorcet winner, and those in the red regions propose their own type. The dashed line marks the position of the Condorcet winner. The regions are drawn for a status quo lying in the same extreme region as the proposer; whenever the status quo lies in $[\theta, \bar{\theta}]$, every type proposes the Condorcet winner. The second row shows the probability of implementing the Condorcet winner for $T \in \{2, \dots, 5\}$ for each distribution.

⁹The symmetry of the utility and its dependence on x and θ only through the distance $|x - \theta|$ are what make the median peak the proposers' insurance-optimal proposal and pin down the threshold type $\bar{\theta}$. The curvature of the loss reshapes these thresholds and sets the strength of the distributional condition, but not the qualitative result: a flatter loss tightens the condition and must be offset by a more symmetric F , until in the linear limit only $\mathbb{E}(\theta) = \theta_m$ remains.

This proposal strategy is supported by sophisticated voting. As shown in Lemma 4 in the Appendix, sincere and sophisticated voting induce the same winner in every round under this proposal strategy, so the result does not depend on voters' sophistication.

It follows from the proposers' equilibrium strategy that, for highly skewed distributions of types and a status quo positioned at the tail of the distribution, the final winner may itself lie at the tail. This possibility becomes progressively less likely as the number of rounds grows, because sustaining a tail outcome requires every subsequent proposer to be drawn from the tail. Consequently, the probability of implementing the Condorcet winner approaches one within a few rounds, as illustrated in Figure 1. The following theorem makes this precise: part (i) gives the exact rate at which this probability converges to one, and part (ii) identifies conditions under which VRP always the Condorcet winner with certainty. Part (ii) relies on the following restriction on the distribution of types.

Assumption 3. The distribution of types F satisfies

$$\text{Var}(\theta) \geq \begin{cases} \theta_m^2 - \mathbb{E}(\theta)^2 & \text{if } \theta_m \geq 0.5, \\ (1 - \theta_m)^2 - (1 - \mathbb{E}(\theta))^2 & \text{if } \theta_m < 0.5. \end{cases}$$

The condition holds whenever the median and the mean are sufficiently close, that is, whenever the distribution is not highly asymmetric relative to its dispersion.

Theorem 2. Let Assumptions 1 and 2 hold, and consider the proposal strategy of Theorem 1.

(i) If $q^1 \notin [\underline{\theta}, \bar{\theta}]$,

$$\text{Prob}(w^T = \theta_m) = \begin{cases} 1 - F(\underline{\theta})^{T-1} & \text{if } \theta_m \geq 0.5, \\ 1 - (1 - F(\bar{\theta}))^{T-1} & \text{if } \theta_m < 0.5. \end{cases}$$

For any $q^1 \in [\underline{\theta}, \bar{\theta}]$, we have $\text{Prob}(w^T = \theta_m) = 1$.

(ii) If $q^1 \in [\underline{\theta}, \bar{\theta}]$ or $\theta_{s^1} \in [\underline{\theta}, \bar{\theta}]$, then $w^T = \theta_m$ for every $T \geq 2$. If Assumption 3 holds in addition, then $w^T = \theta_m$ for every q^1 , θ_{s^1} and $T \geq 2$.

Proof. See Appendix C. □

Part (i) shows that the probability of selecting the Condorcet winner rises geometrically in the number of rounds. When $\theta_m > 0.5$, it equals $1 - F(\underline{\theta})^{T-1} \geq 1 - F(\theta_m)^{T-1}$, since $\underline{\theta} \leq \theta_m$, with the symmetric bound applying when $\theta_m < 0.5$. Taking $1 - F(\theta_m)^{T-1}$ as a distribution-free lower bound, the Condorcet winner is implemented with probability exceeding 96.8% within six voting rounds, regardless of the specific distribution of types.

Part (ii) identifies conditions under which the Condorcet winner is proposed in the first round, wins the first vote and is never displaced, so that it is implemented for any horizon, in particular for $T = 2$. This is the case whenever the first-round status quo or the first proposer's type is not extreme, i.e. lies in $[\underline{\theta}, \bar{\theta}]$. Under Assumption 3 the threshold interval is the whole policy space, so the condition holds for every status quo and every proposer type. Intuitively, under Assumption 3 all types prefer to propose the Condorcet winner rather than their own peak, since otherwise the next proposer might implement a policy on the opposite side of the distribution of peaks. Thus, the threat of a polarized outcome tomorrow induces moderation in today's policies.

Assumption 3 is the incentive constraint of the most extreme type. Suppose $\theta_m > 0.5$, so that $\theta_s = 0$ is the type furthest from the Condorcet winner. Proposing θ_m delivers the certain loss θ_m^2 , whereas proposing its own peak leaves the final winner distributed as F on $[0, 1]$, with expected loss $\text{Var}(\theta) + \mathbb{E}(\theta)^2$. The extreme type faces a trade off between location and risk: the lottery is centered at $\mathbb{E}(\theta) \leq \theta_m$ and so lies nearer its peak, but it is dispersed. It proposes the Condorcet winner precisely when the dispersion it would bear outweighs this location advantage, $\text{Var}(\theta) \geq \theta_m^2 - \mathbb{E}(\theta)^2$, which is Assumption 3. Assumption 1 makes the lottery attractive in location, since it implies $\mathbb{E}(\theta) \leq \theta_m$, whereas Assumption 3 requires the dispersion to be large enough that the extreme type nonetheless prefers the certain outcome. Equivalently, Assumption 3 holds if and only if $\underline{\theta} = 0$ (respectively $\bar{\theta} = 1$ when

$\theta_m < 0.5$), i.e. when even the most extreme type weakly prefers the Condorcet winner, so the green region in Figure 1 covers the whole support.¹⁰ Although Assumptions 3 and 2 are not nested, Assumption 3 makes $[\underline{\theta}, \bar{\theta}] = [0, 1]$. It therefore implies that the realisation-level sufficient condition in Theorem 2(ii) holds for every status quo and proposer type. When Assumption 3 fails, that condition may still deliver two-round implementation for particular realisations.

Together the two assumptions bound the nonparametric skewness $S := (\mathbb{E}(\theta) - \theta_m) / \text{Std}(\theta)$ of F . In particular, for $\theta_m > 0.5$:

$$\frac{1}{\text{Std}(\theta)} \int_0^{2\theta_m-1} F(v) dv \leq -S \leq \frac{\text{Std}(\theta)}{\theta_m + \mathbb{E}(\theta)},$$

where the lower and upper bounds follow from Assumptions 1 (see Lemma 1) and 3, respectively. Analogous bounds hold for $\theta_m < 0.5$. The upper bound reflects the quadratic loss: the extreme type weighs a squared distance against a second moment, so the admissible skewness rises with the standard deviation. It is therefore not a universal cap on S : a tightly concentrated F must be nearly symmetric, whereas a dispersed one may be substantially skewed.

5 Discussion and Outlook

We demonstrate that incorporating random proposers into an iterative majority voting process efficiently implements the socially optimal alternative within a few rounds. Furthermore, if individual utility-maximizing alternatives are sufficiently symmetrically distributed, or if the initial status quo or the first proposer's type is not extreme, the VRP procedure selects the Condorcet winner in only two rounds.

¹⁰The two distributional conditions are related but not nested. Assumption 3 compares the most extreme type with the lottery generated by the full distribution F , whereas Assumption 2 compares the threshold types with lotteries generated by distributions censored to symmetric windows about the median. The underlying lotteries differ, so neither assumption is an instance of the other.

We use a simple framework with three key assumptions. First, the distribution of peaks of the individual preferences is public knowledge. Second, there is a continuum of voters, which ensures that no single voter can be selected as proposer more than once and that proposals carry no private information, thereby eliminating signaling incentives. Third, our analysis assumes single-peaked preferences.

Our results also suggest that the central moderation mechanism may extend, at least approximately, when the distribution of preference peaks is not publicly known. In such an environment, the relevant benchmark would no longer be the known median θ_m , but the median of the posterior distribution over preference peaks. A proposer who is uncertain about the location of the Condorcet winner may still have an incentive to propose an alternative close to this posterior median in order to insure herself against an unfavorable proposal in a later round. The main difference is that proposals and voting outcomes may now convey information about the distribution of preferences, so that the relevant posterior and its median can evolve across rounds. The qualitative force toward moderation may therefore survive, while exact implementation of the Condorcet winner may require additional rounds and depend on the informativeness of the observed history. Formally characterizing this interaction between learning, proposal incentives, and convergence under VRP is a natural direction for future research.

A related extension concerns finite electorates. With finitely many voters, the realized median is uncertain, individual voters may be pivotal, and a proposal can reveal information about both the proposer and the composition of the electorate. Consequently, the exact continuum results need not hold: proposal-making may involve signaling (Banks, 1990, 1993), voting may reflect both pivotal and informational considerations (Piketty, 2000; Razin, 2003), and equilibrium outcomes may be history dependent. Nevertheless, when the electorate is large and the realized median is sufficiently concentrated around the population median, the continuum model suggests that VRP may retain an approximate implementation property, except in histories in which the relevant majority comparisons are close.

Future work could establish finite-population bounds on the probability of implementing a Condorcet winner, study how rapidly these bounds approach the continuum benchmark, and determine whether alternative disclosure or proposer-selection rules can limit the additional signaling incentives.

Finally, our analysis assumes single-peaked preferences. It is well known that a Condorcet winner might not exist when preferences violate this assumption and pairwise voting procedures might exhibit cycles. Voting with Random Proposers offers a probabilistic approach to resolving such cycles. Future research should build on the work of [Airiau and Endriss \(2009\)](#), who show that the proposal dynamic resembles a Markov chain, and explore the convergence and strategic incentives in more general preference domains.

References

- Airiau, S., & Endriss, U. (2009). Iterated majority voting. In F. Rossi & A. Tsoukias (Eds.), *Algorithmic decision theory* (pp. 38–49). Berlin, Heidelberg: Springer.
- Ali, S. N., Bernheim, B. D., Bloedel, A. W., & Console Battilana, S. (2023). Who controls the agenda controls the legislature. *American Economic Review*, 113(11), 3090–3128.
- Anesi, V., & Seidmann, D. (2014). Bargaining over an endogenous agenda. *Theoretical Economics*, 9(2), 445–482.
- Austen-Smith, D. (1987). Sophisticated sincerity: Voting over endogeneous agendas. *American Political Science Review*, 81(4), 1323–1330.
- Banks, J. S. (1990). Equilibrium behavior in crisis bargaining games. *American Journal of Political Science*, 599–614.
- Banks, J. S. (1993). Two-sided uncertainty in the monopoly agenda setter model. *Journal of Public Economics*, 50(3), 429–444.
- Banks, J. S. (2002). Strategic aspects of political systems. In R. Aumann & S. Hart (Eds.), *Handbook of game theory with economic applications* (1st ed., Vol. 3, p. 2203–2228). Elsevier.
- Banks, J. S., & Duggan, J. (2000). A bargaining model of collective choice. *American Political Science Review*, 94(1), 73–88.
- Barberà, S., & Gerber, A. (2017). Sequential voting and agenda manipulation. *Theoretical Economics*, 12(1), 211–247.
- Barberà, S., & Jackson, M. (1994). A characterization of strategy-proof social choice functions for economies with pure public goods. *Social Choice and Welfare*, 11(3), 241–252.
- Barberà, S., & Gerber, A. (2022). Deciding on what to decide. *International Economic Review*, 63(1), 37–61.
- Baron, D. (1991). A spatial bargaining theory of government formation in parliamentary systems. *American Political Science Review*, 85(1), 137–164.
- Baron, D., & Ferejohn, J. A. (1989). Bargaining in legislatures. *The American Political*

- Science Review*, 83(4), 1181–1206.
- Baron, D., & Kalai, E. (1993). The simplest equilibrium of a majority-rule division game. *Journal of Economic Theory*, 61(2), 290–301.
- Bell, C. (1981). A random voting graph almost surely has a Hamiltonian cycle when the number of alternatives is large. *Econometrica*, 49, 1597–1603.
- Black, D. (1948). On the rationale of group decision-making. *Source: Journal of Political Economy*, 56(1), 23–34. Retrieved from <https://www.jstor.org/stable/1825026>
- Border, K. C., & Jordan, J. S. (1983). Straightforward elections, unanimity and phantom voters. *The Review of Economic Studies*, 50(1), 153–170.
- Callander, S. (2011). Searching for good policies. *American Political Science Review*, 105(4), 643–662.
- Eraslan, H., Evdokimov, K., & Zápál, J. (2020). Dynamic legislative bargaining. (Available at SSRN: <https://ssrn.com/abstract=3634333> or via DOI: <http://dx.doi.org/10.2139/ssrn.3634333>)
- Eraslan, H., & Evdokimov, K. S. (2019). Legislative and multilateral bargaining. *Annual Review of Economics*, 11, 443–472.
- Eraslan, H., Evdokimov, K. S., & Zápál, J. (2022). Dynamic legislative bargaining. In E. Karagözoğlu & K. B. Hyndman (Eds.), *Bargaining: Current research and future directions* (pp. 151–175). Cham: Springer International Publishing.
- Garg, N., Kamble, V., Goel, A., Marn, D., & Munagala, K. (2019). Iterative local voting for collective decision-making in continuous spaces. *Journal of Artificial Intelligence Research*, 64, 315–355.
- Gersbach, H. (2009). Democratic mechanisms. *Journal of the European Economic Association*, 7, 1436–1469.
- Goel, A., & Lee, D. (2012). Triadic consensus: A randomized algorithm for voting in a crowd. In P. W. Goldberg (Ed.), *Internet and network economics* (pp. 434–447). Berlin, Heidelberg: Springer Berlin Heidelberg.

- Horan, S. (2021). Agendas in legislative decision-making. *Theoretical Economics*, 16, 235–274.
- Kalai, A., & Kalai, E. (2001). Strategic polarization. *Journal of Mathematical Psychology*, 45(4), 656–663.
- Kalandrakis, T. (2016). Pareto efficiency in the dynamic one-dimensional bargaining model. *Journal of Theoretical Politics*, 28(4), 525–536.
- Klaus, B., & Protopapas, P. (2020). On strategy-proofness and single-peakedness: median-voting over intervals. *International Journal of Game Theory*, 49, 1059 – 1080.
- McKelvey, R. (1979). General conditions for global intransitivities in formal voting models. *Econometrica*, 47, 1085–1112.
- Miller, N. (1977). Graph theoretic approaches to the theory of voting. *American Journal of Political Science*, 24, 68–96.
- Moulin, H. (1980). On strategy-proofness and single peakedness. *Public Choice*, 35(4), 437–455.
- Nakamura, Y. (2020). *Non-manipulable agenda setting for voting with single-peaked preferences* (Vol. 72; Working Paper No. 2). Yokohama City University.
- Piketty, T. (2000). Voting as communicating. *The Review of Economic Studies*, 67(1), 169–191.
- Rasch, B. (2000). Parliamentary floor voting procedures and agenda setting in Europe. *Legislative Studies Quarterly*, 25(1), 3–23.
- Razin, R. (2003). Signaling and election motivations in a voting model with common values and responsive candidates. *Econometrica*, 71(4), 1083–1119.
- Riboni, A., & Ruge-Murcia, F. J. (2010). Monetary policy by committee: Consensus, chairman dominance, or simple majority? *The Quarterly Journal of Economics*, 125(1), 363–416.
- Rosenthal, H., & Zame, W. R. (2022). Sequential referenda with sophisticated voters. *Journal of Public Economics*, 212, 104702.

- Rubinstein, A. (1979). A note about the “nowhere denseness” of societies having an equilibrium under majority rule. *Econometrica*, 47, 511–514.
- Schofield, N. (1983). Generic instability of majority rule. *Review of Economic Studies*, 50, 695–705.
- van Zwet, W. R. (1979). Mean, median, mode II. *Statistica Neerlandica*, 33(1), 1-5.
- Zapal, J. (2016). Markovian equilibria in dynamic spatial legislative bargaining: Existence with three players. *Games and Economic Behavior*, 98(C), 235-242.

Notation

Variable	Explanation
$[0, 1]$	Policy space
θ	Type of an agent in $[0, 1]$, representing his/her peak of preferences
F	Cumulative distribution function of types
f	Density function of types
θ_m	Median of the type distribution
$\mathbb{E}(\theta)$	Mean of the type distribution
$\text{Var}(\theta)$	Variance of the type distribution
$\text{Beta}(\alpha, \beta)$	Beta distribution with parameters α and β
$u_\theta(x)$	Utility of type θ for policy x
$c(x)$	Reflection of x around θ_m , truncated to $[0, 1]$
$\underline{\theta}$	Lower threshold type
$\bar{\theta}$	Upper threshold type
T	Fixed number of voting rounds
q^t	Status quo in round t ; $q^{t+1} = w^t$
p^t	Proposal made in round t by the randomly drawn proposer
w^t	Winner of the pairwise majority vote in round t
θ_{s^t}	Type of the proposer drawn in round t
$q(h)$	Difference between the upper- and lower-tail masses at distance h from θ_m
$s(h)$	Sum of the two tail masses at distance h from θ_m

Variable	Explanation
H	Maximum symmetric window width, $H = \min\{\theta_m, 1 - \theta_m\}$
X_h	Round- T proposer type censored to $[\theta_m - h, \theta_m + h]$
S	Nonparametric skewness measure, $(\mathbb{E}(\theta) - \theta_m)/\text{Std}(\theta)$

A Appendix: Proof of Lemma 1

Since θ_m is a constant, we express the difference between the mean and the median using the tail-integral representation:

$$\begin{aligned}\mathbb{E}(\theta) - \theta_m &= \mathbb{E}(\theta - \theta_m) = \int_0^\infty \Pr(\theta - \theta_m > h) dh - \int_0^\infty \Pr(\theta - \theta_m < -h) dh \\ &= \int_0^{1-\theta_m} [1 - F(\theta_m + h)] dh - \int_0^{\theta_m} F(\theta_m - h) dh,\end{aligned}$$

where the limits tighten because $1 - F(\theta_m + h) = 0$ for $h > 1 - \theta_m$ and $F(\theta_m - h) = 0$ for $h > \theta_m$, given $\theta \in [0, 1]$.

Consider first a distribution with $\theta_m \leq \frac{1}{2}$, so that $\theta_m \leq 1 - \theta_m$. We split the upward integral at $h = \theta_m$ and group its lower part with the downward integral:

$$\mathbb{E}(\theta) - \theta_m = \int_0^{\theta_m} [1 - F(\theta_m + h) - F(\theta_m - h)] dh + \int_{\theta_m}^{1-\theta_m} [1 - F(\theta_m + h)] dh.$$

Since the first integral is nonnegative by Assumption 1, we have

$$\mathbb{E}(\theta) - \theta_m \geq \int_{\theta_m}^{1-\theta_m} [1 - F(\theta_m + h)] dh = \int_{2\theta_m}^1 [1 - F(v)] dv \geq 0,$$

where the equality results from substituting $v = \theta_m + h$ and is the final result. The case $\theta_m > \frac{1}{2}$ is handled the same way, with the roles of the two tails reversed. Since

$$\mathbb{E}(\theta) - \theta_m = \int_0^{1-\theta_m} [1 - F(\theta_m + h) - F(\theta_m - h)] dh - \int_{1-\theta_m}^{\theta_m} F(\theta_m - h) dh,$$

and the first integral is nonpositive by Assumption 1,

$$\mathbb{E}(\theta) - \theta_m \leq - \int_{1-\theta_m}^{\theta_m} F(\theta_m - h) dh = - \int_0^{2\theta_m-1} F(v) dv \leq 0,$$

where the equality results from substituting $v = \theta_m - h$ and is the final result.

B Appendix: Proof of Theorem 1

Recall that $c(x)$ denotes the farthest positioned alternative in $[0, 1]$ that is at least as preferred as x . It is given by the symmetric point of x with respect to the Condorcet winner, θ_m , i.e., $c(x) := \min\{\max\{2\theta_m - x, 0\}, 1\}$. We solve the sequential game through backward induction starting at round T .

Optimal proposal in round T

In the last round T the winner is implemented immediately, so that every voter's continuation utility from an alternative coincides with her current utility from it. Sincere and sophisticated voting therefore coincide for *every* voter in round T , irrespective of the alternatives on the ballot, and each voter supports whichever of p^T and q^T is closer to her peak. Since preferences are single-peaked and the round is a pairwise majority vote, the winning alternative $w(p^T, q^T)$ for $q^T \leq \theta_m$ is given by

$$(1) \quad w(p^T, q^T) = \begin{cases} p^T & \text{if } p^T \in [q^T, c(q^T)], \\ q^T & \text{if } p^T \in [0, q^T) \cup (c(q^T), 1]. \end{cases}$$

The peak of the preference of the proposer at round T is given by θ_{s^T} . Consider the optimization problem of the proposer in round T :

$$\arg \max_{p^T \in [q^T, c(q^T)]} -(p^T - \theta_{s^T})^2.$$

Note that the proposer is not constrained about the position of the proposal in the interval $[0, 1]$ by the voting procedure, but he/she would anticipate that proposing an alternative in $[0, q^T)$ or $(c(q^T), 1]$ would lose the pairwise vote and cannot be utility-improving. Thus, the

optimal proposal and final winner are given by

$$p^T \in \begin{cases} [0, q^T] & \text{if } \theta_{s^T} \in [0, q^T], \\ \{\theta_{s^T}\} & \text{if } \theta_{s^T} \in [q^T, c(q^T)], \\ \{c(q^T)\} & \text{else.} \end{cases} \quad w(p^T, q^T) = \begin{cases} q^T & \text{if } \theta_{s^T} \in [0, q^T], \\ \theta_{s^T} & \text{if } \theta_{s^T} \in [q^T, c(q^T)], \\ c(q^T) & \text{else.} \end{cases}$$

Whenever $\theta_{s^T} \in [0, q^T]$, every proposal in $[0, q^T]$ induces the same winner q^T and is therefore optimal. We select $p^T = \theta_{s^T}$. The symmetric argument applies when $q^T > \theta_m$, yielding:

$$p^T = \begin{cases} \min\{\theta_{s^T}, c(q^T)\} & \text{if } q^T \leq \theta_m, \\ \max\{\theta_{s^T}, c(q^T)\} & \text{if } q^T > \theta_m. \end{cases} \quad w(p^T, q^T) = \begin{cases} \min\{q^T, c(q^T)\} & \text{if } \theta_{s^T} < \min\{q^T, c(q^T)\}, \\ \max\{q^T, c(q^T)\} & \text{if } \theta_{s^T} > \max\{q^T, c(q^T)\}, \\ \theta_{s^T} & \text{else.} \end{cases}$$

Optimal proposal in round $T - 1$

We now consider the optimal proposal and voting at round $T - 1$. We organise the proof in the following way. First, we derive the winner of round $T - 1$ that would maximize the expected utility of the proposer, which, as we show, is either the proposer's own peak or the Condorcet winner. Then, we determine the optimal proposals in round $T - 1$. Finally, we show that the majority of the strategic voters would vote in favour of the proposal in Lemma 2.

We now state the expected utility of the proposer $\theta_{s^{T-1}}$ when proposing w in round $T - 1$. To ease the notation we denote the proposer's type $\theta_{s^{T-1}} = \theta_s$ in this subsection:

$$\text{EU}_s(w) = \begin{cases} \int_0^w u_{\theta_s}(w) f(v) dv + \int_w^{c(w)} u_{\theta_s}(v) f(v) dv + \int_{c(w)}^1 u_{\theta_s}(c(w)) f(v) dv & \text{if } w \leq \theta_m, \\ \int_0^{c(w)} u_{\theta_s}(c(w)) f(v) dv + \int_{c(w)}^w u_{\theta_s}(v) f(v) dv + \int_w^1 u_{\theta_s}(w) f(v) dv & \text{if } w > \theta_m, \end{cases}$$

which can be simplified as follows:

$$\text{EU}_s(w) = \begin{cases} u_{\theta_s}(w)F(w) + \int_w^{c(w)} u_{\theta_s}(v)f(v)dv + u_{\theta_s}(c(w))(1 - F(c(w))) & \text{if } w \leq \theta_m, \\ u_{\theta_s}(c(w))F(c(w)) + \int_{c(w)}^w u_{\theta_s}(v)f(v)dv + u_{\theta_s}(w)(1 - F(w)) & \text{if } w > \theta_m. \end{cases}$$

Each of the three terms in the expected utility function corresponds to the possible winners in round T , depending on the proposer's type in round T . In the case when $w \leq \theta_m$ and for $\theta_{sT} \in [0, w^{T-1}]$, the final winner is $w^T = q^T = w^{T-1}$. For $\theta_{sT} \in [w^{T-1}, c(w^{T-1})]$, the final winner will be θ_{sT} . Finally, for $\theta_{sT} \in (c(w^{T-1}), 1]$, the final winner is $c(w^{T-1})$. The analogous argument holds for the case when $w \geq \theta_m$.

We proceed by finding the utility-maximizing winner considering the following cases.

Case 1: $w^{T-1} \in [0, 2\theta_m - 1]$. Thus, $c(w^{T-1}) = 1$. The first order condition is given by

$$u_{\theta_s}(w)'F(w) + u_{\theta_s}(w)f(w) - u_{\theta_s}(w)f(w) = u_{\theta_s}(w)'F(w) = 0.$$

Observe that $\text{EU}_s(w)$ is monotonically increasing for $\theta_s \geq \max\{2\theta_m - 1, 0\}$. If $\theta_s < \max\{2\theta_m - 1, 0\}$, we have $u_{\theta_s}(w)' = 0$ when $w^{T-1} = \theta_s$. Thus, $\text{EU}_s(w)$ has a maximum at $w^{T-1} = \theta_s$ as $u_{\theta_s}(w)' = 2(\theta_s - w)$ and

$$w^{T-1} = \min\{2\theta_m - 1, \theta_s\}.$$

Case 2: $w^{T-1} \in (\max\{0, 2\theta_m - 1\}, \min\{2\theta_m, 1\})$. Let $h := |w^{T-1} - \theta_m|$ and $H := \min\{\theta_m, 1 - \theta_m\}$, so that $h \in [0, H]$. On this range $c(w^{T-1}) = 2\theta_m - w^{T-1}$. For $h \in [0, H]$ both $\theta_m - h$ and $\theta_m + h$ lie in $[0, 1]$ with $c(\theta_m - h) = \theta_m + h$ and $c(\theta_m + h) = \theta_m - h$, so by the characterisation of the final round the winner is the round- T proposer's peak truncated to $[\theta_m - h, \theta_m + h]$ for both $w^{T-1} = \theta_m - h$ and $w^{T-1} = \theta_m + h$. Both proposals thus induce the same winner lottery $X_h := \min\{\max\{\theta_{sT}, \theta_m - h\}, \theta_m + h\}$, which equals θ_{sT} on $[\theta_m - h, \theta_m + h]$ and is pushed to the nearer endpoint outside it, while $w^{T-1} = \theta_m$ ($h = 0$) yields the certain outcome θ_m .

It therefore suffices to compare $u_{\theta_s}(\theta_m)$ with $\text{EU}_s(w^{T-1}) = -\mathbb{E}[(\theta_s - X_h)^2]$:

$$\text{EU}_s(\theta_m) - \text{EU}_s(w^{T-1}) = \mathbb{E}[(X_h - \theta_m)^2] - 2(\theta_s - \theta_m) \mathbb{E}[X_h - \theta_m].$$

Since $X_h = \min\{\max\{\theta_{s^T}, \theta_m - h\}, \theta_m + h\}$, its centred value $X_h - \theta_m$ equals $-h$ when $\theta_{s^T} < \theta_m - h$, equals $\theta_{s^T} - \theta_m$ when $\theta_{s^T} \in [\theta_m - h, \theta_m + h]$, and equals $+h$ when $\theta_{s^T} > \theta_m + h$.

Hence

$$\begin{aligned} \mathbb{E}[X_h - \theta_m] &= -h F(\theta_m - h) + \int_{\theta_m - h}^{\theta_m + h} (v - \theta_m) f(v) dv + h (1 - F(\theta_m + h)), \\ \mathbb{E}[(X_h - \theta_m)^2] &= h^2 F(\theta_m - h) + \int_{\theta_m - h}^{\theta_m + h} (v - \theta_m)^2 f(v) dv + h^2 (1 - F(\theta_m + h)). \end{aligned}$$

Integrating the middle term of the first line by parts, $\int_{\theta_m - h}^{\theta_m + h} (v - \theta_m) f(v) dv = [(v - \theta_m) F(v)]_{\theta_m - h}^{\theta_m + h} - \int_{\theta_m - h}^{\theta_m + h} F(v) dv = h F(\theta_m + h) + h F(\theta_m - h) - \int_{\theta_m - h}^{\theta_m + h} F(v) dv$, and the boundary terms cancel those already present, leaving

$$\mathbb{E}[X_h - \theta_m] = h - \int_{\theta_m - h}^{\theta_m + h} F(v) dv.$$

Substituting $v = \theta_m + r$ and folding the two half-ranges, $\int_{\theta_m - h}^{\theta_m + h} F(v) dv = \int_0^h (F(\theta_m + r) + F(\theta_m - r)) dr$, while $h = \int_0^h dr$; therefore

$$\mathbb{E}[X_h - \theta_m] = \int_0^h (1 - F(\theta_m + r) - F(\theta_m - r)) dr = \int_0^h q(r) dr.$$

Treating the second line identically—by parts, $\int_{\theta_m - h}^{\theta_m + h} (v - \theta_m)^2 f(v) dv = [(v - \theta_m)^2 F(v)]_{\theta_m - h}^{\theta_m + h} - 2 \int_{\theta_m - h}^{\theta_m + h} (v - \theta_m) F(v) dv$, whose boundary terms again cancel—gives

$$\mathbb{E}[(X_h - \theta_m)^2] = h^2 - 2 \int_{\theta_m - h}^{\theta_m + h} (v - \theta_m) F(v) dv = 2 \int_0^h r (1 - F(\theta_m + r) + F(\theta_m - r)) dr = 2 \int_0^h r s(r) dr,$$

where the last line substitutes $v = \theta_m + r$, uses $h^2 = 2 \int_0^h r dr$, and folds the two half-ranges

of $\int_{\theta_m-h}^{\theta_m+h} (v-\theta_m)F(v) dv = \int_0^h r(F(\theta_m+r) - F(\theta_m-r)) dr$. Substituting these moments gives

$$\text{EU}_s(\theta_m) - \text{EU}_s(w^{T-1}) = 2 \left(\int_0^h r s(r) dr - (\theta_s - \theta_m) \int_0^h q(r) dr \right), \quad \theta_s \in [0, 1], \quad h \in [0, H].$$

Since $\theta_s \in [\underline{\theta}, \bar{\theta}]$, we have $\theta_s - \theta_m \in [-(\theta_m - \underline{\theta}), \bar{\theta} - \theta_m]$, and therefore

$$-(\theta_s - \theta_m) \int_0^h q(r) dr \geq \begin{cases} -(\bar{\theta} - \theta_m) \int_0^h q(r) dr, & \text{if } \int_0^h q(r) dr \geq 0, \\ (\theta_m - \underline{\theta}) \int_0^h q(r) dr, & \text{if } \int_0^h q(r) dr \leq 0. \end{cases}$$

In either case, Assumption 2 gives $\int_0^h r s(r) dr - (\theta_s - \theta_m) \int_0^h q(r) dr \geq 0$ for every $h \in [0, H]$, hence $\text{EU}_s(\theta_m) - \text{EU}_s(w^{T-1}) \geq 0$. Therefore, for every $\theta_s \in [\underline{\theta}, \bar{\theta}]$, the median peak yields weakly higher expected utility than every winner in the middle range.

Case 3: $w^{T-1} \in [2\theta_m, 1]$. Thus, by definition $c(w^{T-1}) = 0$. The FOC is given by

$$u_{\theta_s}(w)'(1 - F(w)) + u_{\theta_s}(w)f(w) - u_{\theta_s}(w)f(w) = u_{\theta_s}(w)'(1 - F(w)) = 0.$$

Observe that $\text{EU}_s(w)$ is monotonically decreasing for $\theta_s \leq \min\{2\theta_m, 1\}$. If $\theta_s > \min\{2\theta_m, 1\}$, the expected utility is maximized at $w^{T-1} = \theta_s$. Thus,

$$w^{T-1} = \max\{2\theta_m, \theta_s\}.$$

The utility-maximizing winner of round $T-1$ therefore lies in $\{\min\{2\theta_m-1, \theta_s\}, \theta_m, \max\{2\theta_m, \theta_s\}\}$.

We organise the rest of the proof in a series of lemmas:

1. we identify the interval of types of proposers who have a higher expected utility from θ_s winning round $T-1$ instead of θ_m when $\theta_s < 2\theta_m - 1$ or $\theta_s > 2\theta_m$,
2. we show that all such types of proposers prefer a winner within the interval over θ_m ,
3. we determine the voting behaviour that sustains the resulting proposal, and show that

a decisive majority votes identically whether sincerely or sophisticatedly.

Lemma 2. *If $\theta_s \in [0, \underline{\theta}) \cup (\bar{\theta}, 1]$, then $EU_s(\theta_s) > u_{\theta_s}(\theta_m)$. If $\theta_s \in [\underline{\theta}, \bar{\theta}]$, then $EU_s(\theta_s) \leq u_{\theta_s}(\theta_m)$.*

Proof. Suppose first that $\theta_m > \frac{1}{2}$ and define

$$D_L(\theta_s) := EU_s(\theta_s) - u_{\theta_s}(\theta_m) = \int_{\theta_s}^1 u_{\theta_s}(v) f(v) dv - u_{\theta_s}(\theta_m)$$

for $\theta_s \leq 2\theta_m - 1$. Differentiating gives

$$D'_L(\theta_s) = 2 \left(\theta_s F(\theta_s) + \int_{\theta_s}^1 v f(v) dv - \theta_m \right).$$

At $\theta_s = 0$ the expression in parentheses is nonpositive since $\mathbb{E}(\theta) \leq \theta_m$ whenever $\theta_m > \frac{1}{2}$, by Assumption 1 and Lemma 1. It is nondecreasing in θ_s , and at $2\theta_m - 1$ it equals

$$\int_0^{1-\theta_m} (1 - F(\theta_m + h) - F(\theta_m - h)) dh \leq 0$$

by Assumption 1. Hence D_L is nonincreasing on $[0, 2\theta_m - 1]$. By the definition of $\underline{\theta}$ as the boundary of the strict-preference region, $D_L(\theta_s) > 0$ for $\theta_s \in [0, \underline{\theta})$ and $D_L(\theta_s) \leq 0$ for $\theta_s \in [\underline{\theta}, 2\theta_m - 1]$.

For $\theta_s \in [2\theta_m - 1, 1]$, the proposer's own peak belongs to the middle range, because $2\theta_m > 1$. Case 2 and Assumption 2 therefore give $EU_s(\theta_s) \leq u_{\theta_s}(\theta_m)$ for every interior type. This proves the claim when $\theta_m > \frac{1}{2}$.

Suppose next that $\theta_m < \frac{1}{2}$. The symmetric upper-tail difference

$$D_R(\theta_s) := EU_s(\theta_s) - u_{\theta_s}(\theta_m) = \int_0^{\theta_s} u_{\theta_s}(v) f(v) dv - u_{\theta_s}(\theta_m)$$

is nondecreasing on $[2\theta_m, 1]$, by the symmetric calculation and Assumption 1. Hence $D_R(\theta_s) > 0$ for $\theta_s \in (\bar{\theta}, 1]$ and $D_R(\theta_s) \leq 0$ for $\theta_s \in [2\theta_m, \bar{\theta}]$. For $\theta_s \in [0, 2\theta_m]$, the comparison follows

from Case 2 and Assumption 2. At $\theta_m = \frac{1}{2}$, Case 2 and continuity cover the entire policy interval. This completes the proof. $\square$

Next, we show that sufficiently extreme proposer types prefer every winner between their peak and the corresponding threshold to the median peak.

Lemma 3. *For every $\theta_s \in [0, \underline{\theta}]$, $EU_s(w) \geq u_{\theta_s}(\theta_m)$ for all $w \in [\theta_s, \underline{\theta}]$, and, symmetrically, for every $\theta_s \in (\bar{\theta}, 1]$, $EU_s(w) \geq u_{\theta_s}(\theta_m)$ for all $w \in [\bar{\theta}, \theta_s]$.*

Proof. Suppose first that $\theta_m > \frac{1}{2}$. If $\underline{\theta} = 0$, the lower-tail claim is vacuous. Let $\underline{\theta} > 0$ and fix $\theta_s \in [0, \underline{\theta}]$. Case 1 implies that $EU_s(w)$ is nonincreasing in w on $[\theta_s, \underline{\theta}]$, so it suffices to compare $\underline{\theta}$ with θ_m . Define

$$K_L(\theta_s) := EU_{\theta_s}(\underline{\theta}) - u_{\theta_s}(\theta_m).$$

Differentiating with respect to θ_s gives

$$K'_L(\theta_s) = 2 \left(\underline{\theta} F(\underline{\theta}) + \int_{\underline{\theta}}^1 v f(v) dv - \theta_m \right).$$

The expression in parentheses is nonpositive. Indeed, it is nondecreasing in $\underline{\theta} \leq 2\theta_m - 1$, and its value at $2\theta_m - 1$ is nonpositive by the calculation in the proof of Lemma 2. Since $K_L(\underline{\theta}) = 0$, it follows that $K_L(\theta_s) \geq 0$ for every $\theta_s \in [0, \underline{\theta}]$. Together with the monotonicity in w , this proves the lower-tail claim.

The upper-tail claim when $\theta_m < \frac{1}{2}$ is symmetric. If $\bar{\theta} = 1$, it is vacuous. Otherwise, Case 3 shows that $EU_s(w)$ is nondecreasing on $[\bar{\theta}, \theta_s]$. The symmetric analogue of K_L is nondecreasing in θ_s and vanishes at $\bar{\theta}$, so it is nonnegative for every $\theta_s \in (\bar{\theta}, 1]$. $\square$

It remains to translate the utility-maximizing winner into a proposal. If $q^{T-1} \geq \underline{\theta}$, no proposal below $\underline{\theta}$ can win against the standing status quo. Thus a lower-tail winner is attainable only if both the status quo and the proposer's type are below $\underline{\theta}$. Symmetrically,

an upper-tail winner is attainable only if both are above $\bar{\theta}$. In every other case, the best attainable winner is θ_m . Hence an optimal proposal for the proposer in round $T - 1$ is:

$$(2) \quad p^{T-1} = \begin{cases} \theta_{s^{T-1}} & \text{if } \max\{q^{T-1}, \theta_{s^{T-1}}\} < \underline{\theta} \text{ or } \min\{q^{T-1}, \theta_{s^{T-1}}\} > \bar{\theta}, \\ \theta_m & \text{else.} \end{cases}$$

It remains to verify that this proposal is sustained by the voting behaviour selected by the pivotal-voting refinement, and that the same winner is obtained under sincere voting.

Interior-type bound. For every $\theta \in [\underline{\theta}, \bar{\theta}]$ and every $w \in [0, 1]$,

$$\text{EU}_\theta(w) \leq u_\theta(\theta_m).$$

To see this, first suppose that $\theta_m > \frac{1}{2}$. For a winner in the middle range, the claim follows from Case 2 and Assumption 2. For $w \leq 2\theta_m - 1$, Case 1 shows that the maximum over this range is attained at $w = \theta$ when $\theta \leq 2\theta_m - 1$, and at $w = 2\theta_m - 1$ otherwise. The first comparison follows from Lemma 2; the second follows from Case 2 at the boundary, by continuity. Since $2\theta_m > 1$, there is no upper outer range. The case $\theta_m < \frac{1}{2}$ is symmetric, and when $\theta_m = \frac{1}{2}$ Case 2 and continuity cover the whole interval.

Lemma 4. *Under the proposal strategy in Equation (2), sophisticated voting satisfies the pivotal-voting refinement. Under the same proposal strategy, sincere and sophisticated voting induce the same winner in every round.*

Proof. Suppose first that $p^{T-1} = \theta_m$. If $\theta_m > \frac{1}{2}$, the set $[\underline{\theta}, 1]$ has mass $1 - F(\underline{\theta}) > \frac{1}{2}$. By the interior-type bound, every voter in this set weakly prefers the certain continuation outcome θ_m to the round- T lottery induced by q^{T-1} . Hence θ_m wins under sophisticated voting. Under sincere voting, θ_m also wins because it is the Condorcet winner. Thus the winner is the same under the two voting behaviours, although their supporting coalitions need not coincide. The case $\theta_m < \frac{1}{2}$ is symmetric, using the decisive set $[0, \bar{\theta}]$; at $\theta_m = \frac{1}{2}$ the

threshold interval is the whole policy space.

Next suppose that $p^{T-1} = \theta_{s^{T-1}}$. If $\theta_m > \frac{1}{2}$, both q^{T-1} and $\theta_{s^{T-1}}$ lie below $\underline{\theta}$. Let

$$a := \max\{q^{T-1}, \theta_{s^{T-1}}\}, \quad b := \min\{q^{T-1}, \theta_{s^{T-1}}\}.$$

Every voter with type at least a weakly prefers the round- T lottery induced by a to the lottery induced by b : for every final-round proposer type, the former continuation outcome is weakly higher, and the two outcomes differ only when the higher one is closer to such a voter. This set has mass $1 - F(a) > \frac{1}{2}$, so a wins under sophisticated voting. The same set sincerely supports a against b , and therefore the same alternative wins under sincere voting. The argument for $\theta_m < \frac{1}{2}$ is symmetric, with the lower of the two alternatives supported by the decisive set of types below it. The same reasoning applies recursively in every earlier round because the proposal rule preserves the relevant tail interval until θ_m is proposed. $\square$

Optimal proposal in rounds $t \in \{1, 2, \dots, T-2\}$

If $\underline{\theta} = 0$ and $\bar{\theta} = 1$, the preceding analysis implies that every proposer selects the Condorcet winner. Suppose instead that $\theta_m > \frac{1}{2}$ and $\underline{\theta} > 0$; the case $\theta_m < \frac{1}{2}$ is symmetric. If the standing status quo is at least $\underline{\theta}$, no lower-tail proposal can win, and the continuation outcome is θ_m . If the status quo is below $\underline{\theta}$, a proposer compares a lower-tail continuation with the certain outcome θ_m .

For a proposer of type θ_{st} , the expected-utility difference generated by the boundary winner $\underline{\theta}$ in round $t \leq T-2$ is

$$\begin{aligned} & \text{EU}_{\theta_{st}}(\underline{\theta}) - u_{\theta_{st}}(\theta_m) \\ &= u_{\theta_{st}}(\underline{\theta})F(\underline{\theta})^{T-t} + F(\underline{\theta})^{T-t-1} \int_{\underline{\theta}}^1 u_{\theta_{st}}(v)f(v)dv - u_{\theta_{st}}(\theta_m)F(\underline{\theta})^{T-t-1} \\ &= F(\underline{\theta})^{T-t-1} \left(u_{\theta_{st}}(\underline{\theta})F(\underline{\theta}) + \int_{\underline{\theta}}^1 u_{\theta_{st}}(v)f(v)dv - u_{\theta_{st}}(\theta_m) \right). \end{aligned}$$

The term in parentheses is the round- $(T-1)$ comparison at $\underline{\theta}$. It is nonnegative for $\theta_{st} \in [0, \underline{\theta})$

by Lemma 3. It is nonpositive for $\theta_{s^t} \in [\underline{\theta}, \theta_m]$ because, as established in the proof of Lemma 3, it is nonincreasing in the proposer's type and vanishes at $\underline{\theta}$. Together with the monotonicity of the lower-tail continuation in the winner, this gives the same proposal rule as in Equation (2). The upper-tail argument is symmetric.

Hence the rule in Equation (2) selects an optimal proposal in every round $t \leq T - 1$. Finally, the argument of Lemma 4 applies recursively in every earlier round: sincere and sophisticated voting induce the same winner at each voting stage.

Lemma 5. *Let Assumptions 1 and 2 hold. Then $EU_{2\theta_m-1}(2\theta_m-1) \leq u_{2\theta_m-1}(\theta_m)$ if $\theta_m > \frac{1}{2}$, and $EU_{2\theta_m}(2\theta_m) \leq u_{2\theta_m}(\theta_m)$ if $\theta_m < \frac{1}{2}$.*

Proof. Write

$$Q_H := \int_0^H q(r) dr, \quad R_H := \int_0^H r s(r) dr.$$

Suppose first that $\theta_m > \frac{1}{2}$ and let $H := 1 - \theta_m$. The moment calculation in Case 2, evaluated at $\theta_s = w = 2\theta_m - 1$, gives $EU_{2\theta_m-1}(2\theta_m-1) - u_{2\theta_m-1}(\theta_m) = -2(R_H + HQ_H)$. Assumption 1 gives $Q_H \leq 0$. Moreover, $\underline{\theta} \leq 2\theta_m - 1$, so $\theta_m - \underline{\theta} \geq H$. By Assumption 2, $R_H \geq -(\theta_m - \underline{\theta})Q_H \geq -HQ_H$, which proves the first inequality.

Suppose next that $\theta_m < \frac{1}{2}$ and let $H := \theta_m$. At $\theta_s = w = 2\theta_m$, the same moment calculation gives $EU_{2\theta_m}(2\theta_m) - u_{2\theta_m}(\theta_m) = 2(HQ_H - R_H)$. Now Assumption 1 gives $Q_H \geq 0$, while $\bar{\theta} - \theta_m \geq H$. Assumption 2 therefore implies $R_H \geq (\bar{\theta} - \theta_m)Q_H \geq HQ_H$, which proves the second inequality. $\square$

C Appendix: Proof of Theorem 2

We prove the two parts of the theorem in turn.

Part (i). Let $q^1 \notin [\underline{\theta}, \bar{\theta}]$. Consider $\theta_m > \frac{1}{2}$, so that $\bar{\theta} = 1$ and hence $q^1 < \underline{\theta}$. Whenever the Condorcet winner is proposed in a round $t \leq T - 1$, it wins that round and is implemented. The only history in which θ_m is not proposed in the first $T - 1$ rounds is therefore the history

in which every proposer in those rounds has a type in $[0, \underline{\theta}]$. Along that history the status quo remains in the lower-tail interval. Since proposer types are drawn independently from F , this event has probability $F(\underline{\theta})^{T-1}$, and hence

$$\text{Prob}(w^T = \theta_m) = 1 - F(\underline{\theta})^{T-1}.$$

The symmetric argument for $\theta_m < \frac{1}{2}$ yields $1 - (1 - F(\bar{\theta}))^{T-1}$. If $q^1 \in [\underline{\theta}, \bar{\theta}]$, the first proposer selects θ_m , which is implemented. When $\theta_m = \frac{1}{2}$, the threshold interval is $[0, 1]$, so only this latter case arises.

Part (ii). We first show that Assumption 3 is equivalent to $\underline{\theta} = 0$ when $\theta_m > \frac{1}{2}$, and to $\bar{\theta} = 1$ when $\theta_m < \frac{1}{2}$. In either case the other threshold is already at the boundary, so $[\underline{\theta}, \bar{\theta}] = [0, 1]$. Equation (2) then prescribes θ_m in round 1, and Lemma 4 implies that it wins under sophisticated and sincere voting alike. Hence $q^2 = \theta_m$ and θ_m is implemented when $T = 2$. At $\theta_m = \frac{1}{2}$, the threshold interval is $[0, 1]$ by definition. Moreover, Assumption 1 gives $\mathbb{E}(\theta) \geq \theta_m$, so $\theta_m^2 - \mathbb{E}(\theta)^2 \leq 0$ and the first branch of Assumption 3 holds automatically.

Case $\theta_m > \frac{1}{2}$. Let

$$D_L(\theta_s) := \text{EU}_{\theta_s}(\theta_s) - u_{\theta_s}(\theta_m).$$

As established in the proof of Lemma 2, D_L is continuous and nonincreasing on $[0, 2\theta_m - 1]$. Lemma 5 gives $D_L(2\theta_m - 1) \leq 0$. By the definition of $\underline{\theta}$ as the boundary of the strict-preference region,

$$\underline{\theta} = 0 \iff D_L(0) \leq 0.$$

Since $c(0) = 1$ and $u_0(0) = 0$,

$$D_L(0) = \text{EU}_0(0) - u_0(\theta_m) = \int_0^1 u_0(v) f(v) dv - u_0(\theta_m).$$

Thus $\underline{\theta} = 0$ if and only if $\int_0^1 v^2 f(v) dv \geq \theta_m^2$, equivalently, $\text{Var}(\theta) \geq \theta_m^2 - \mathbb{E}(\theta)^2$, which is Assumption 3 in this case.

Case $\theta_m < \frac{1}{2}$. Define

$$D_R(\theta_s) := \mathbb{E}U_{\theta_s}(\theta_s) - u_{\theta_s}(\theta_m).$$

The symmetric calculation in the proof of Lemma 2 shows that D_R is continuous and non-decreasing on $[2\theta_m, 1]$, while Lemma 5 gives $D_R(2\theta_m) \leq 0$. Hence

$$\bar{\theta} = 1 \iff D_R(1) \leq 0 \iff \int_0^1 u_1(v) f(v) dv \leq u_1(\theta_m).$$

Substituting the utility function and simplifying gives

$$\begin{aligned} \text{Var}(\theta) + \mathbb{E}(\theta)^2 - \theta_m^2 &\geq 2(\mathbb{E}(\theta) - \theta_m), \\ \text{Var}(\theta) &\geq (\mathbb{E}(\theta) - \theta_m)(2 - \mathbb{E}(\theta) - \theta_m) \\ &= (1 - \theta_m)^2 - (1 - \mathbb{E}(\theta))^2, \end{aligned}$$

which is Assumption 3 for $\theta_m < \frac{1}{2}$.

It remains to establish the realisation-level sufficient condition. Suppose that $q^1 \in [\underline{\theta}, \bar{\theta}]$ or $\theta_{s^1} \in [\underline{\theta}, \bar{\theta}]$. Equation (2) prescribes the proposer's own type only if

$$\max\{q^1, \theta_{s^1}\} < \underline{\theta} \quad \text{or} \quad \min\{q^1, \theta_{s^1}\} > \bar{\theta}.$$

Either membership condition makes both inequalities false, so $p^1 = \theta_m$. By Lemma 4, θ_m wins round 1 under sophisticated and sincere voting alike, and therefore $q^2 = \theta_m$. Since $c(\theta_m) = \theta_m$, the final-round winner is θ_m . Hence $w^2 = \theta_m$ for every such realisation.