

Who Controls the Agenda Controls the Legislature[†]

By S. NAGEEB ALI, B. DOUGLAS BERNHEIM, ALEXANDER W. BLOEDEL,
AND SILVIA CONSOLE BATTILANA*

We model legislative decision-making with an agenda setter who can propose policies sequentially, tailoring each proposal to the status quo that prevails after prior votes. Voters are sophisticated, and the agenda setter cannot commit to future proposals. Nevertheless, the agenda setter obtains her favorite outcome in every equilibrium regardless of the initial default policy. Central to our results is a new condition on preferences, manipulability, that holds in rich policy spaces, including spatial settings and distribution problems. Our findings therefore establish that, despite the sophistication of voters and the absence of commitment power, the agenda setter is effectively a dictator. (JEL D71, D72, D78)

A central goal of democratic institutions is to balance competing interests by distributing political power evenly among society's members. A recurring concern is that power is often highly concentrated in the hands of a select few leaders, who steer policy in their favor by manipulating institutional procedures. Our objective is to understand how seemingly prodemocratic institutions can yield this concentration of power.

Legislative institutions merit particular attention in this regard because of the central role they typically play in determining policy. Although legislative procedures vary, they often involve (i) a single agenda setter (e.g., committee chair or party leader), who acts as a gatekeeper for proposals, and (ii) a group of legislators who vote to approve or reject proposals. While the agenda setter controls which policies come up for a vote, she cannot unilaterally dictate policy because passage of her proposals requires majority support. Even so, the degree to which voting constrains the agenda setter is unclear.

*Ali: Department of Economics, Pennsylvania State University (email: nageeb@psu.edu); Bernheim: Department of Economics, Stanford University (email: bernheim@stanford.edu); Bloedel: Department of Economics, UCLA (email: abloedel@econ.ucla.edu); Battilana: Auctionomics (email: silviacb@auctionomics.com). Jeff Ely was the coeditor for this article. We thank the coeditor, anonymous referees, Attila Ambrus, Vincent Anesi, Renee Bowen, Navin Kartik, Andreas Kleiner, Barton Lee, David Levine, Elliot Lipnowski, Jay Lu, Steve Matthews, Debraj Ray, Ludvig Sinander, Bruno Strulovici, Juuso Toikka, Alex Wolitzky, Jan Zapal, and various conference and seminar audiences for useful comments. We would also like to thank George Orwell, who recognized the importance of agenda control (albeit in a somewhat different setting), for inspiring our title. We gratefully acknowledge financial support from NSF Grants SES-0137129 (Bernheim) and SES-1530639 (Ali and Bernheim). Bloedel gratefully acknowledges the hospitality and financial support of Caltech and Stanford, where parts of this research were completed.

[†]Go to <https://doi.org/10.1257/aer.20221578> to visit the article page for additional materials and author disclosure statements.

An important literature in political economy seeks to understand the power that flows from agenda control. In a seminal contribution, McKelvey (1976) observes that, in rich policy settings, the agenda setter can exploit cycles in the majority relation to obtain *any* desired outcome by appropriately sequencing proposals and is therefore unconstrained by the majority's will. However, this striking conclusion assumes legislators are nonstrategic and vote myopically, without accounting for subsequent modifications of the policy. Sophisticated legislators, who have learned from experience, anticipate the paths of proposals and votes, and accept or reject proposals based on the final outcomes to which those paths lead. Subsequent research concludes that the requirement of majority approval constrains the agenda setter's power when voters are sophisticated. Specifically, Shepsle and Weingast (1984) show that the agenda setter can only achieve policies that are not *covered* by the initial default option.¹ In rich policy settings, this constraint is often stringent (McKelvey 1986).

The prior literature generally assumes the agenda setter must commit to a sequence of proposals, which is set in stone regardless of which proposals pass. And yet, in many settings, nothing prevents the agenda setter from bringing different proposals to the floor depending on how previous votes turned out. The impact of the fixed-agenda assumption is complex. On the one hand, it attenuates the power of agenda setters because it precludes them from tailoring their proposals to the prevailing circumstances. On the other hand, it magnifies their power because it endows them with the ability to commit to proposals they might not want to make when the time arrives. Relative to this benchmark, it is unclear how the benefits that real-world agenda setters accrue from being able to flexibly tailor proposals stack up against the costs they incur from being unable to commit. This is the question we address.

We study a novel model of *real-time agenda control*. A single agenda setter and a group of voters choose a policy using the following procedure: for each of finitely many rounds, (i) the agenda setter can propose an alternative to the contemporaneous default policy (where the first-round default policy is given exogenously), and (ii) a vote is taken to determine whether the current proposal or the contemporaneous default policy will serve as the default policy in the next round. The policy that prevails in the final vote (for the terminal round) is implemented and determines payoffs. We investigate the subgame perfect equilibria of this game, assuming neither voters nor the agenda setter can commit to their future decisions.²

Real-time agenda control has stark implications for collective choice problems that satisfy the following property: a collective choice problem is *manipulable* if, for

¹ A policy x covers y if a majority of voters strictly prefer x to y , and every policy that is majority-preferred to x is also majority-preferred to y .

² One possible interpretation of this amendment process is that legislatures can revisit any bill specifying the policy for some future date after passing it, before that date arrives. Such reconsideration regularly occurs in practice. It is perhaps most visible in the context of sunset provisions, which specify a policy change as of a sunset date. Prior to the arrival of that date, the provisions are often extended. For example, sunset provisions for the Bush tax cuts of 2001 and 2003 effectively specified that tax rates would increase in 2010. Before 2010 arrived, those cuts were extended to 2012, effectively reducing the rates for 2010. Senator Diane Feinstein remarked that the sunset provisions in the original bill were "critical to my decision to support this legislation" because they "will allow us to revisit the components of this bill in the future." See Fahrenthold (2012). More generally, whenever Congress passes a "permanent" policy in year T , it implicitly specifies a policy for year $T' > T$. Any further modification of the policy between years T and T' then amounts to reconsideration of the year- T' policy.

every policy x other than the agenda setter's favorite, there is an alternative policy y that both she and a majority of voters strictly prefer to x . We show that canonical formulations of spatial and distributive politics satisfy manipulability. More precisely, the standard spatial model almost surely satisfies manipulability if the policy space has three or more dimensions. We also define a broad class of "distribution problems" that satisfy manipulability. This class includes divide-the-dollar problems as well as problems that mix non-zero-sum policies with transfers. Indeed, augmenting any collective choice problem with pork or transfers makes that problem manipulable. Our favored interpretation of manipulability, and its prevalence, is that the discordance of the majority's will in multidimensional problems inevitably creates opportunities for the agenda setter.

We establish the potency of real-time agenda control when the environment is manipulable. Our main finding, stated informally, is as follows:

Main Result: *If there are sufficiently many rounds, the agenda setter obtains her favorite policy in every equilibrium regardless of the initial default policy if and only if the collective choice problem is manipulable.*

Thus, for a wide range of collective choice problems, the agenda setter effectively dictates policy, despite voters' sophistication and her lack of commitment power. Manipulability is necessary and sufficient for this conclusion; absent manipulability, we show that equilibrium outcomes sometimes remain bounded away from the agenda setter's favorite.

Theorems 1–3 formalize our main conclusion under a range of technical conditions, including for finite and continuous policy spaces, and clarify how many rounds are "sufficient." The core argument involves a simple observation concerning "one-step" improvements: if the default option going into the terminal round is not the agenda setter's favorite, she will propose her favorite policy among those that both she and a majority of voters prefer to the default. Manipulability guarantees that such an improvement exists. Applying this logic iteratively implies that, if she can make proposals for t rounds, she can obtain the outcome generated by the t -fold iteration of this "favorite improvement" operator. At each stage, voters pass proposals that lead to this outcome because a majority prefer it to the outcome that would emerge otherwise. When t is sufficiently large, this iterative process yields a policy arbitrarily close to (if not exactly the same as) the agenda setter's favorite.

Though simple, this logic is extremely general. It applies for voting rules other than simple majority, so long as the analog of manipulability holds. The same conclusion also holds for other widely studied legislative procedures, such as the *closed-rule* or *successive procedure* as well as *open-rule* bargaining. More broadly, we obtain a *protocol-equivalence* result for the class of "generalized amendment" protocols: fixing a preference profile and voting rule, all of these protocols (and others) yield the same equilibrium outcome under real-time agenda control.

Our findings thereby illuminate the forces that contribute to the concentration of political power, and explain why voter sophistication may not be an effective safeguard against agenda control. Recent empirical findings highlight similar themes: Berry and Fowler (2016, 2018) observe that chairs of congressional committees have disproportionate influence on policymaking, and Fourniaies (2018) finds that

special interest groups make greater campaign contributions to legislators endowed with procedural authority.

Our analysis has the additional implication that agenda setters benefit from bundling policy choices with transfers and pork. Augmenting any collective choice problem with transfers renders it manipulable not only for simple majority rule but also for any voting rule that does not provide any voter with veto power. Our main results then imply that such bundling allows the agenda setter to obtain her favorite policy for any “veto-proof” voting rule.³ Moreover, we show that in these settings, the dictatorship result holds even when the process involves a relatively small number of rounds: for simple majority rule, three rounds suffice, and for general “veto-proof” voting rules, the number of rounds need not exceed the number of voters. Analogously, we find that the agenda setter may benefit from linking policy issues in order to make bargaining more multidimensional. Specifically, our analysis of spatial politics shows that the collective choice problem is generically manipulable if and only if the policy space has three or more dimensions. Thus, if the current legislative debate concerns a one- or two-dimensional policy decision, the agenda setter may benefit from bundling that decision with other policy issues—even “settled” ones for which the default option already coincides with her favorite policy—because the overall problem thereby becomes manipulable.

To isolate the role of sequential rationality constraints in real-time agenda control, we compare our model to the commitment benchmark in which the agenda setter can commit to any strategy in the dynamic game. Therein, we find that the agenda setter can obtain her favorite policy among those that are reachable from the initial default option through a finite chain of majority improvements, mirroring Miller’s (1977) classic characterization of outcomes achievable through general binary voting trees.⁴ Relative to this benchmark, our main results show that manipulability enables the agenda setter to attain her commitment payoff without having to commit. Absent manipulability, sequential rationality typically precludes her from achieving the commitment benchmark. In such cases, even a commitment to a fixed agenda, as in Shepsle and Weingast (1984), may leave her better-off.

As in the prior literature, we assume the agenda process involves a finite number of rounds. This assumption is appropriate when the purpose of negotiation is to solve a time-indexed collective choice problem, i.e., to select the policy that prevails at a given point in time. Problems of this form are ubiquitous. For example, when a legislature negotiates over the budget for a given fiscal year, it cannot continue those negotiations into the subsequent fiscal year. In such cases, there is both a *deadline* for meaningful deliberations (e.g., 11:59 PM on December 31) and an inherent constraint on the speed at which the legislature can consider new proposals. In combination, these considerations imply that the number of rounds is necessarily bounded.

³Previous studies have highlighted the detrimental effects of pork on legislative and democratic politics (e.g., Lizzeri and Persico 2001; Battaglini and Coate 2008; Maskin and Tirole 2019). Our analysis shows that the mere ability to use pork or transfers yields dictatorial power; in equilibrium, the agenda setter does not actually transfer benefits to any party.

⁴It follows from Miller (1977) that the commitment outcome is the agenda setter’s favorite policy whenever the latter is reachable from the initial default option through some majority chain. Combined with McKelvey’s (1976, 1979) results on global intransitivities, this observation implies that, for rich policy settings, voter sophistication does not constrain the power of an agenda setter who can commit to a plan specifying a proposal for every contingency.

We assume, for the sake of tractability, that this number is known in advance, but our results plainly extend to settings in which an initially unknown termination point becomes evident during the course of negotiations.

While we think it is reasonable to assume the existence of a deadline—insofar as negotiations over time-indexed actions are widespread—one could alternatively consider processes that allow negotiations to continue indefinitely. Diermeier and Fong (2011) and Anesi and Seidmann (2014) adopt that approach, analyzing infinite-horizon counterparts of our baseline model. They show that the limitless potential for reconsideration severely constrains agenda power. Read in the context of that work, our findings establish the critical importance of deadlines. We show, in effect, that a simple commitment to termination at a fixed point in time allows an agenda setter to achieve her favorite outcome (in manipulable environments). More broadly, we find that an agenda setter's preference over the length of negotiations is nonmonotonic: she prefers a moderate number of proposal rounds both to a single round and to an open-ended process with no limit on duration. Thus, our analysis implies that even if there is no natural deadline, a strategic agenda setter benefits from inventing excuses to establish one.

We are not the first to show that certain collective choice processes can produce dictatorial outcomes. Kalandrakis (2004) finds that bargaining with an endogenous status quo and changing proposers yields such results in a divide-the-dollar setting. Bernheim, Rangel, and Rayo (2006) analyze a model of pork barrel politics with changing proposers and show that the final proposer is effectively the dictator. The endogenous evolution of the default option is an essential feature of those frameworks. A few papers conclude that dictatorial power prevails under the opposite assumption of closed-rule negotiations, where accepting an offer results in its immediate implementation. Ali, Bernheim, and Fan (2019) find that a modest form of predictability about future bargaining power results in the first proposer obtaining the entire surplus in a closed-rule divide-the-dollar setting; Duggan and Ma (2023) consider settings in which a single agent makes all proposals and show that she has approximate dictatorial power.

These studies obtain results for specific policy spaces and legislative procedures, building primarily on the legislative bargaining literature. Our work differs in several important respects. It is instead rooted in the classical literature on agenda setting, to which we contribute by investigating the implications of real-time agenda control without commitment. Instead of focusing on a particular policy space, we identify manipulability—a property that any given space may or may not satisfy—as the necessary and sufficient condition for dictatorial power, and we also show that canonical models of spatial and distributive politics satisfy this property. Moreover, we demonstrate that the strategic logic behind our main result is robust, in that it applies to a wide range of legislative procedures, allowing for either evolving or fixed default options.

Section I illustrates the core logic of our results through a simple example. Section II describes the general model, and Section III contains our main results. Section IV explains why manipulability holds in spatial and distributive politics. Section V describes the commitment benchmark, investigates the implications of real-time agenda control for other legislative procedures, and elaborates on the role of deadlines. Section VI concludes. Omitted proofs are in the Appendix and online Appendix.

I. An Example

A legislature, comprised of an agenda setter and n voters (where n is odd), chooses a policy from $\{w, x, y, z\}$. The agenda setter has a strict preference relation that coincides with the listed order. Each voter has a complete and transitive preference relation that is also strict, but the profile of voter preferences results in a strict majority relation \succ_M that cycles. We depict the agenda setter's preferences and the majority relation in Figure 1.⁵

The legislature selects a policy through what is known in the literature as the *amendment procedure*, which operates as follows:

- (i) There is an initial default option;
- (ii) In each of finitely many rounds, the agenda setter proposes a new policy (the *proposal*), which is put to a vote against the prevailing default;
- (iii) In each nonterminal round, the policy that obtains a majority of the votes becomes the default for the subsequent round;
- (iv) The policy that obtains a majority of the votes in the final round is implemented.

The amendment procedure features prominently in practice and is the focus of considerable prior work. In much of the literature, the agenda setter lays out the sequence of proposals in advance, prior to any voting. We call this procedure a *fixed agenda* protocol because it does not permit the agenda setter to vary her proposals based on the concurrent default or prior votes. Our analysis contrasts this protocol with *real-time agenda control*, which allows the agenda setter to tailor proposals to the circumstances that arise but does not endow her with any commitment power. We use this example to illustrate the distinct implications of fixed agenda protocols and real-time agenda control.

Suppose the initial default is z , the agenda setter's least favorite policy. McKelvey (1976) points out that if voters are myopic, the agenda setter can obtain her favorite policy w by exploiting the cycles in the majority relation: she uses a fixed agenda where y is the first proposal, x is the second proposal, and w is the third and final proposal. Because voters are myopic, in each instance they anticipate no further revisions, so each proposal passes, and the process selects w . Shepsle and Weingast (1984) show that this conclusion does not hold if voters are sophisticated. Instead, the agenda setter can obtain only those policies that are *not covered* by the initial default option (as defined in footnote 1). In our example, the agenda setter's favorite uncovered policy is x , which she can obtain with the following fixed agenda: propose x in the first round and y in the second.⁶ Voter sophistication would therefore appear to limit the power of agenda control.

⁵This majority relation can arise whenever $n \geq 3$. For the three-voter case, suppose $z \succ_1 w \succ_1 x \succ_1 y, x \succ_2 y \succ_2 z \succ_2 w$, and $y \succ_3 z \succ_3 w \succ_3 x$.

⁶Although a majority of voters prefer z to x , sophisticated voters anticipate that rejection of x in the first round would lead to a final outcome of y , as y is majority-preferred to z .




FIGURE 1

Notes: Panel A shows the agenda setter's preferences. Panel B shows the majority relation. In each case, an arrow from policy p to p' denotes $p' \succ p$.

Our central insight is that giving the agenda setter the flexibility to make proposals in real time, so that she can tailor each proposal to the prevailing default option, unleashes the full power of agenda control and allows her to obtain her favorite policy even if voters are sophisticated. While it is intuitive that the agenda setter benefits from greater flexibility, note that we simultaneously remove her ability to commit, which could in principle limit her power by introducing sequential rationality constraints.

To illustrate how the agenda setter can exploit real-time agenda control, we construct an equilibrium for a three-round game that selects policy w . Consider the following strategy for the agenda setter: if the default option in any round is policy p , she proposes her *favorite improvement* to p —in other words, her favorite policy among those the majority prefers to p . We use $\phi(p)$ to denote this policy. On the equilibrium path (starting from an initial default of z), this strategy prescribes proposing y first, then x , and then w . Notice that this sequence coincides with the optimal agenda for myopic voters. But in this instance, voters approve each policy not out of myopia but rather because they (correctly) anticipate future play. We can verify this claim through backward induction.

$t = 3$: If the default option is p , the agenda setter proposes $\phi(p)$, which, by construction, results in $\phi(p)$.

$t = 2$: If the default option is p , the agenda setter proposes $\phi(p)$. Anticipating the behavior at $t = 3$, voters understand that approving this policy today ultimately results in $\phi^2(p)$ —the twofold iteration of the ϕ operator—whereas rejecting this policy results in $\phi(p)$. Since a majority of voters prefer $\phi^2(p)$ to $\phi(p)$, the proposal passes.

$t = 1$: Analogously, in the first period, the agenda setter proposes $y = \phi(z)$. Voters anticipate that approving this proposal ultimately results in $w = \phi^2(y)$, whereas rejecting it ultimately results in $x = \phi^2(z)$. Since a majority favor w over x , the proposal passes.

Thus, a majority of voters always finds it sequentially rational at each stage to approve the proposal that this strategy prescribes.

Because the agenda setter cannot make commitments, her behavior must also be sequentially rational. Indeed, in the final round, she proposes her favorite option among those that will pass. Given the equilibrium for the final round, her second-round proposal always achieves her favorite outcome among the feasible alternatives (see above). Likewise, given the equilibrium for the last two rounds, she cannot improve on her prescribed first-round proposal. Therefore, no deviation can make her strictly better-off.

Thus, there is a subgame perfect equilibrium in which the agenda setter obtains w . Our main result (Theorem 1) reaches a stronger conclusion: even though there are multiple equilibria, w is the *unique* subgame perfect equilibrium outcome regardless of the initial default so long as there are three or more rounds.⁷ Real-time agenda control therefore guarantees that the group will select the agenda setter's favorite policy.

In this example, the agenda setter's preferences and the majority relation jointly satisfy a condition we call *manipulability*: for every policy p other than the agenda setter's favorite, there is a policy p' that both she and a majority of voters strictly prefer to p . Our main results show that the agenda setter exercises dictatorial power if and only if this condition is satisfied: when it fails, then for some initial default options, the agenda setter cannot obtain her favorite policy in any equilibrium.

To understand why manipulability is necessary, consider the majority relation in Figure 2. The solid black arrows are the same as before, but the red dashed arrows are different. Policy x is now *unimprovable*: there is no other policy that the agenda setter and a majority of voters all prefer to x . As x is not the agenda setter's favorite option, manipulability fails. Our characterization result (Lemma 1) implies that if the initial default option is z , the agenda setter necessarily obtains x in every equilibrium regardless of the horizon. Intuitively, voters anticipate that if z remains the default option in the terminal round, sequential rationality will compel the agenda setter to propose x because x is her favorite policy among the options that will pass. But then x must also be the outcome of a two-stage game starting with a default of z : rejecting the first proposal leads to x , and x is unimprovable, so the majority will not support any proposal leading to an option the agenda setter would prefer. The same argument applies, recursively, to games of any length.

In this example, the agenda setter is stymied by her inability to commit: were she able to lock in a fixed agenda, as in Shepsle and Weingast (1984), she could achieve w by proposing w in the first round and y in the second. A majority of voters would then approve w in the first round because rejection would yield y . In our setting, the agenda setter cannot achieve this outcome because proposing y in the second round is not sequentially rational: if z remains the default option in the second round, she would instead propose x , and anticipating that behavior, voters would be unwilling to approve w in the first round.

These examples illustrate the role of manipulability in empowering the agenda setter to obtain her favorite policy without the need for commitment. Although manipulability may appear restrictive, we show in Section IV that it is satisfied in standard models of spatial and distributive politics.

⁷We impose the standard refinement that voters vote as if they are pivotal.




FIGURE 2

Notes: Panel A shows the agenda setter’s preferences. Panel B shows the majority relation, with red dashed arrows denoting differences from that in Figure 1.

II. Model

Our model consists of two components: (i) a (static) collective choice problem comprising the set of feasible policies and agents’ preferences over them and (ii) a dynamic procedure for selecting a policy. We describe each in turn.

Collective Choice Problem.—A group $N := \{1, \dots, n\}$ of voters (where $n > 1$ is odd) and a single nonvoting agenda setter (A) choose a policy from policy space X . This space is compact and metrizable; in most of our examples, it is either finite or a subset of a finite-dimensional Euclidean space. For each $i \in \{1, \dots, n, A\}$, \succsim_i denotes player i ’s preference relation over policies. Each relation is continuous and has a continuous utility representation $u_i : X \rightarrow \mathbb{R}$. If a majority of voters weakly (respectively, strictly) prefers x to y , we say that $x \succcurlyeq_M y$ (respectively, $x \succ_M y$). We use $X_A^* := \arg \max_{x \in X} u_A(x)$ to denote the set of the agenda setter’s favorite policies. Together, the policy space and preference profile constitute a *collective choice problem*, $\mathcal{C} := (X, \{\succsim_i\}_{i=1, \dots, n, A})$.

Legislative Procedure.—Our baseline analysis focuses on what the literature calls the *amendment procedure*. Voting takes place in rounds $t \in \{1, \dots, T\}$, where T is finite. Activity prior to round t determines a *default* policy x^{t-1} . The initial default, x^0 , is exogenous. In each round t , the agenda setter proposes a policy (the *proposal*) denoted $a^t \in X$, which can coincide with the existing default policy. The proposal a^t is then put to a vote against the default x^{t-1} . If a majority of voters vote in favor of the proposal (i.e., it “passes”), then it becomes the new default for the subsequent round: $x^t = a^t$. If the proposal does not pass, the default remains unchanged: $x^t = x^{t-1}$. The policy x^T that prevails (after voting) in round T determines payoffs.⁸

⁸Formally, in the pertinent literature (e.g., Shepsle and Weingast 1984), this bargaining framework is known as a “forward agenda” amendment procedure. The literature also considers “backward agenda” procedures wherein, after all amendments have been incorporated, the amended bill is put to a final up-or-down vote against the original default. Our analysis also applies to settings in which the legislature can consider a sequence of bills, each amendable through a backward agenda procedure, prior to the deadline (i.e., the date at which the policy that the bills concern is to take effect). Under this interpretation, each “round” of the procedure we study in this paper pertains to a distinct bill that, if passed, will be implemented unless it is subsequently supplanted by passage of another bill.

Solution Concept.—All players can condition their actions, both proposals and votes, on the history of prior actions. This assumption captures the idea that people take actions—and, in particular, set the agenda—in real time. A *history* h^t as of the beginning of round t records the initial default policy x^0 , the sequence of proposals (a^1, \dots, a^{t-1}) , and the sequence of voting profiles (v^1, \dots, v^{t-1}) in all prior rounds. It therefore identifies the default x^{t-1} prevailing at the beginning of round t . \mathcal{H}^t denotes the space of all round- t histories. A strategy for the agenda setter is a mapping $\sigma_A : \bigcup_{t=1}^T \mathcal{H}^t \rightarrow \Delta(X)$ specifying, for each history h^t , a distribution $\sigma_A(h^t) \in \Delta(X)$ over proposals a^t . A strategy for voter i is a mapping $\sigma_i : \bigcup_{t=1}^T \mathcal{H}^t \times X \rightarrow \Delta(\{y, n\})$ specifying, for each history h^t , a distribution over *yes* or *no* votes for each potential proposal a^t .

We study subgame perfect equilibria of this game. We also assume “as-if pivotal” voting: if passage (respectively, rejection) of the current proposal ultimately leads to continuation outcome x (respectively, y), then anyone who has a *strict* preference for x votes for the option that leads to x , and similarly for y .⁹ This standard assumption rules out unreasonable equilibria in which nonpivotal voters, who are technically indifferent because they cannot affect the outcome, vote contrary to their preferences.¹⁰ Henceforth, we use the term *equilibrium* to denote this solution concept.

III. The Power of Real-Time Agenda Control

We now turn to our main results concerning the agenda setter’s power. Section IIIA defines what it means for a collective choice problem to be manipulable. Section IIIB and Section IIIC consider finite and general policy spaces, respectively.

A. Improvability and Manipulability

We begin by describing policies that the agenda setter can improve upon with a single proposal round.

DEFINITION 1: *Policy x is Improvable if there exists a policy y such that $y \succ_A x$ and $y \succ_M x$; if no such policy exists, then policy x is Unimprovable.*

One can view the set of unimprovable policies as the core of a suitably defined cooperative game in which all decisive coalitions contain both the agenda setter and at least a simple majority of voters.

⁹There is no restriction on the behavior of voters who are indifferent between x and y . This definition applies only at histories where continuation outcomes do not depend on the composition of the current vote, conditional on which policy prevails. As will become apparent, our analysis either assumes strict preferences (Theorems 1 and 2), thereby ensuring that such equilibria exist, or allows for indifference while imposing a mild refinement (Theorem 3), which ensures that continuation outcomes depend only on the prevailing policy.

¹⁰When voters have strict preferences and the policy space is finite, subgame perfection with “as-if pivotal” (simultaneous) voting is outcome-equivalent to both (i) iterated deletion of weakly dominated strategies under simultaneous voting and (ii) mere subgame perfection when voting in each round occurs via “roll call” in a fixed sequential order (see Chapter 4 of Austen-Smith and Banks 2005 and references therein).

The set of unimprovable policies must include all of the agenda setter's favorites, X_A^* . For an important class of collective choice problems, everything else is improvable.

DEFINITION 2: *Collective choice problem \mathcal{C} is Manipulable if every $x \notin X_A^*$ is improvable.*

Manipulability is connected to intransitivity of the majority relation. If that relation is transitive, a Condorcet winner exists, and the choice problem is manipulable if and only if that policy is the agenda setter's favorite. Intransitivities make it easier for the agenda setter to find improvements that the majority will accept.¹¹

In the rest of this section, we identify the implications of real-time agenda control for manipulable and nonmanipulable collective choice problems. In Section IV, we demonstrate that collective choice problems belonging to some familiar and important classes are manipulable.

B. Dictatorial Power with Finite Alternatives

To convey the logic of agenda-setting power most transparently, we start by considering a finite policy space under the "generic" assumption that all players have strict preferences.¹²

DEFINITION 3: *Collective choice problem \mathcal{C} satisfies Generic Finite Alternatives if X is finite, and each \succ_i and \succ_A are antisymmetric.*

For such settings, we obtain the following result:

THEOREM 1: *Suppose the collective choice problem \mathcal{C} satisfies Generic Finite Alternatives. For any game with at least $|X| - 1$ rounds, the agenda setter obtains her favorite policy in every equilibrium regardless of the initial default if and only if \mathcal{C} is Manipulable.*

Theorem 1 articulates the power of real-time agenda control: with a manipulable policy space, the agenda setter always obtains her favorite policy. On its own, manipulability merely ensures that the agenda setter can find some improvement palatable to a majority. Indeed, when agendas are fixed in advance as in Shepsle and Weingast (1984) (so that proposals are not conditional on prior votes), the agenda setter can do no better than her favorite policy among those uncovered by the initial default option, even if the policy space is manipulable. It is therefore the combination of manipulability and real-time agenda control that yields dictatorial power.

¹¹ Manipulability is, however, distinct from the notion of global intransitivity in the majority relation (or "chaos") studied by the classical literature (e.g., McKelvey 1976). Global intransitivity stipulates that for any two policies x and y , there is a majority chain $\{a^k\}_{k=0}^{k=K}$ such that $x = a^0$, $y = a^K$, and $a^k \succ_M a^{k-1}$ for all $k \in \{1, \dots, K\}$. Manipulability plainly does not require global intransitivity and, unlike manipulability, global intransitivity does not depend on the agenda setter's preferences.

¹² For any finite policy space X , the set of utility profiles (u_1, \dots, u_n, u_A) representing strict preference profiles is both open-dense and of full measure in $\mathbb{R}^{|X| \times (n+1)}$.

The argument for Theorem 1 is elementary. Denote the set of policies that are majority-preferred to x by $M(x) := \{y \in X : y \succ_M x \text{ or } y = x\}$. We define the agenda setter’s favorite improvement mapping $\phi : X \rightarrow X$ by

$$(1) \quad \{\phi(x)\} := \arg \max_{y \in M(x)} u_A(y).$$

Given that the agenda setter has strict preferences, $\phi(\cdot)$ is well defined. We denote the fixed points of this mapping by $E := \{x \in X : x = \phi(x)\}$. Note that a policy x is unimprovable if and only if $x \in E$. We write the t -fold iteration of ϕ for any initial default option x^0 as $\phi^t(x^0)$. By definition of ϕ , for every default x^0 , (i) $\phi^{t+1}(x^0) \succ_A \phi^t(x^0)$, and (ii) if $T \geq |X| - 1$, $\phi^T(x^0)$ is an element of E (i.e., unimprovable).

We prove Theorem 1 by showing that equilibrium outcomes are characterized by iterations of the ϕ mapping, regardless of whether manipulability holds. Define the *equilibrium outcome correspondence* for a T -round game as $f_T : X \rightrightarrows X$, where $f_T(x^0)$ is the set of policies chosen with positive probability in any equilibrium given an initial default of x^0 .

LEMMA 1: *Suppose the collective choice problem \mathcal{C} satisfies Generic Finite Alternatives. For any game with T rounds and initial default policy x^0 , the equilibrium outcome correspondence satisfies $f_T(x^0) = \{\phi^T(x^0)\}$. Moreover,*

- (a) *There exists a pure-strategy equilibrium in which (i) the agenda setter always proposes $\phi(x)$ when the current default is x and (ii) each voter i votes to approve proposal y in round t if and only if $\phi^{T-t}(y) \succ_i \phi^{T-t}(x^{t-1})$.*
- (b) *For an initial default x^0 , $f_T(x^0) = \{x^0\}$ if and only if $x^0 \in E$.*
- (c) *If $T \geq |X| - 1$, then $\bigcup_{x^0 \in X} f_T(x^0) = E$.*

Lemma 1 states that the equilibrium outcome correspondence with T rounds coincides with the T -fold iteration of the agenda setter’s favorite improvement mapping, implying that all equilibria are outcome-equivalent. It also asserts the existence of a simple equilibrium in which the agenda setter follows a “greedy” strategy, always acting as if the current round is the last one.¹³ Finally, it records some useful implications: (b) the fixed points of the equilibrium outcome correspondence are the unimprovable policies, and (c) given sufficiently many rounds, every equilibrium outcome is unimprovable.

Perhaps surprisingly, the agenda setter’s strategy in the simple equilibrium above would implement the same outcome if voters were myopic as in McKelvey (1976)—that is, if they ignored the possibility of further amendments. But in our setting, voters approve each proposal *precisely* because they anticipate future revisions and

¹³In this equilibrium, voters break ties in favor of the agenda setter’s proposals. Under Generic Finite Alternatives, voters are indifferent between accepting and rejecting a proposal if and only if both choices lead to the same continuation outcome.

prefer the continuation path associated with the proposal. More specifically, the group of voters who approve each proposal along the equilibrium path are those who favor $\phi^T(x^0)$ to $\phi^{T-1}(x^0)$. Because the continuation outcomes for acceptance and rejection of the current proposal do not vary along the equilibrium path, *the same coalition of voters* supports each on-path proposal.

Theorem 1 is an immediate corollary of Lemma 1(c): the set of unimprovable policies E coincides with X_A^* if and only if \mathcal{C} is Manipulable. We therefore sketch the proof of Lemma 1 here (the full proof is in the Appendix):

- (i) With a single round, an equilibrium policy is an element of $\phi(X)$, where $\phi(X)$ is the image of X under ϕ : if the default option x^0 is improvable, then in equilibrium, the agenda setter proposes her favorite improvement $\phi(x^0)$, which passes.
- (ii) With two rounds, an equilibrium policy is an element of $\phi^2(X)$. If the initial default option x^0 prevails at the end of the first round, then by (i), the resulting policy is $\phi(x^0)$. If the latter policy is improvable, then there exist policies y such that $\phi(y) \succ_M \phi(x^0)$ (for example, $y = \phi(x^0)$). Crucially, in equilibrium, the agenda setter is guaranteed passage of any such proposal in the first round because voters anticipate, by step (i) above, that accepting y would lead to final outcome $\phi(y)$, while rejecting it would lead to $\phi(x^0)$. By definition, the agenda setter's favorite improvement over $\phi(x^0)$ is $\phi^2(x^0)$, so proposing any policy y for which $\phi(y) = \phi^2(x^0)$ is optimal for her. As described in Lemma 1(a), one such first-round proposal is $y = \phi(x^0)$.
- (iii) By induction, with T rounds, an equilibrium policy is an element of $\phi^T(X)$. As noted before, $\phi^T(X)$ must coincide with E if $T \geq |X| - 1$.

While the default evolves gradually in the simple equilibrium of Lemma 1(a), there are other equilibria with sudden transitions. Specifically, if $\phi^T(x^0)$ is unimprovable, there are equilibria where the agenda setter proposes it in the first round and it passes.¹⁴ Thus, if the policy space is manipulable and $T \geq |X| - 1$, the group may adopt the agenda setter's favorite policy immediately even if a majority does not prefer it to the initial default.

¹⁴However, if $\phi^T(x^0)$ is improvable, then subgame perfection requires gradualism: were voters to accept $\phi^T(x^0)$ in the first round, they would expect the agenda setter to further amend the policy to obtain additional gains for herself, contrary to the majority's interests.

Note that our explanation for Lemma 1 did not invoke any properties of majority rule. Consequently, with appropriate adjustments to the notions of *favorite improvement* and *manipulability*, these results generalize to arbitrary voting rules. In Section VB, we obtain similar results for legislative procedures that feature adjournment clauses that terminate deliberation, such as the successive procedure/closed-rule bargaining and open-rule bargaining. Therefore, the simple recursive logic applies to a broad range of legislative institutions.

Two caveats are in order. First, although manipulability is generic in rich multidimensional collective choice problems (see Section IV), the same statement does not hold under Generic Finite Alternatives: for any finite policy space X , the set of utility profiles (u_1, \dots, u_n, u_A) for which manipulability holds has strictly positive, but not full, Lebesgue measure in $\mathbb{R}^{|X| \times (n+1)}$. Second, as the cardinality of X increases, the above results require the number of rounds to increase without bound. We address both issues below.

C. Near-Dictatorial Power with Continuous Policy Spaces

Next, we extend our analysis to settings with continuous policy spaces using two distinct approaches. For the first, we take the view that the typical real-world collective choice problem offers an extremely large but finite number of alternatives and that the assumption of continuity is usually a convenient analytic approximation (e.g., for budgets, pennies are indivisible). Instead of studying the continuous case that approximates the settings of interest, we study the discrete settings that the continuous case approximates (i.e., those with large numbers of alternatives). For the second approach, we study continuous policy spaces directly but impose a mild equilibrium refinement. Both approaches yield the same conclusion: regardless of the initial default option, with sufficiently many rounds, the agenda setter’s payoff is arbitrarily close to its maximum.

Discretized Settings.—Consider a collective choice problem $\mathcal{C} := (X, \{\succ_i\}_{i=1, \dots, n, A})$ that need not satisfy Generic Finite Alternatives. Let $d(x, y)$ denote a metric on X ; for a subset $Y \subseteq X$, $d(x, Y) := \inf_{y \in Y} d(x, y)$ denotes the distance of x from Y . A **generic ϵ -grid** is a finite subset $X_\epsilon \subseteq X$ for which (i) $\max_{x \in X} d(x, X_\epsilon) < \epsilon$ and (ii) the preferences of voters and the agenda setter are antisymmetric within X_ϵ . We study “ambient” collective choice problems that admit generic ϵ -grids for every $\epsilon > 0$. As we establish in the Appendix (Lemma 3), such problems are characterized by the condition that all players have “thin” indifference curves. Formally, let $I_i(x) := \{y \in X : y \sim_i x\}$ denote player i ’s indifference curve through policy x .

DEFINITION 4: *Collective choice problem \mathcal{C} satisfies Thin Individual Indifference if $I_i(x) \setminus \{x\}$ has empty interior for every player i and $x \in X$.*

Thin Individual Indifference holds in most applications with continuous policies, including divide-the-dollar problems and any setting with strictly convex preferences. The assumption also features in McKelvey (1979) and Shepsle and Weingast (1984), who further assume that the policy space has no isolated points.

Our definition generalizes their condition and for finite policy spaces is equivalent to Generic Finite Alternatives.¹⁵

Loosely, we show that, under this assumption, in games with large numbers of rounds and options, the agenda setter “nearly” exercises dictatorial power in all equilibria if and only if the ambient collective choice problem is manipulable. Formally, defining $u_A^* := \max_{x \in X} u_A(x)$ for any (continuous) utility representation u_A of \succsim_A , we have¹⁶

THEOREM 2: *Suppose the collective choice problem \mathcal{C} satisfies Thin Individual Indifference. The following holds if and only if \mathcal{C} is Manipulable:*

For every $\delta > 0$, there exist $\epsilon_\delta > 0$ and $T_\delta \in \mathbb{N}$ such that, if the policy space is restricted to any generic ϵ -grid X_ϵ where $\epsilon < \epsilon_\delta$ and the game has at least T_δ rounds, then given any initial default in X_ϵ , the agenda setter’s payoff is no lower than $u_A^ - \delta$ in any equilibrium.*

We note three features of this result. First, it does not require the discretized collective choice problems, $\mathcal{C}_\epsilon := (X_\epsilon, \{\succsim_i\}_{i=1, \dots, n, A})$, to be manipulable. Second, the agenda setter achieves a payoff within $\delta > 0$ of her maximum among all policies in the ambient policy space X , not merely those in the grid X_ϵ . Third, the minimal horizon length T_δ and maximal discretization size ϵ_δ depend on the payoff approximation δ but are uniform across both the choice of the particular grid X_ϵ and the initial default within that grid. These features distinguish Theorem 2 from Theorem 1: even if \mathcal{C}_ϵ were manipulable, Theorem 1 would only establish that the agenda setter achieves her favorite option if the number of rounds is at least $|X_\epsilon| - 1$, which explodes as $\epsilon \rightarrow 0$. In contrast, Theorem 2 shows that, with T_δ rounds, the agenda setter obtains a payoff within δ of her maximum for *all* sufficiently fine grids.

The following is a sketch of the proof. First, we show that, if the ambient collective choice problem \mathcal{C} is manipulable, then policies that are *unimprovable within the grid* X_ϵ lie in a neighborhood of X_A^* that shrinks to X_A^* as $\epsilon \rightarrow 0$. Thus, even if the agenda setter cannot obtain her favorite policy in X_ϵ (let alone in X), she can make sequences of successful proposals that bring the policy *arbitrarily close* to her favorite. The second step shows that, as long as the grid is sufficiently fine, she can achieve these gains within a fixed number of rounds that does not depend on the particular grid. The essence of the argument is that, for any $\delta > 0$, there exists a minimal payoff improvement $\eta_\delta > 0$ such that, whenever the agenda setter’s payoff differs from that of her favorite policy by more than δ , she can find another policy that improves both her payoff and the payoffs of a majority of voters by at least η_δ .

¹⁵Generic Finite Alternatives implies that every $x \in X$ is isolated and $I_i(x) \setminus \{x\} = \emptyset$.

¹⁶The statement of Theorem 2 uses a cardinal measure of near-dictatorial power, but we can restate it in ordinal terms: if and only if \mathcal{C} is manipulable, the final policy itself must be close to the agenda setter’s favorite policies, X_A^* . That is, for every $\delta > 0$, there exist $\epsilon_\delta > 0$ and $T_\delta \in \mathbb{N}$ such that, for any sufficiently fine grid and long horizon, and for any initial default, all equilibrium outcomes y are close to the agenda setter’s favorite policies, in the sense that $d(y, X_A^*) < \delta$.

Using this observation, it is easy to determine a bound on the number of rounds that necessarily brings her payoff within δ of her maximum.¹⁷

An Equilibrium Refinement for Continuous Settings.—When considering settings with continuous policy spaces, we cannot assume away indifference. This inconvenient fact raises two issues. First, how do voters break ties when indifferent between two continuation paths? Second, how do we define “as-if pivotal” voting when future tie-breaking for other players, and hence continuation paths, may differ depending on the composition of the majority in the current round? A standard approach in the literature is to resolve both issues by restricting attention to pure strategy Markov perfect equilibria in which voters always break indifference in favor of the proposal (e.g., Baron and Ferejohn 1989; Diermeier and Fong 2011). While convenient, this tie-breaking convention potentially stacks the deck in the agenda setter’s favor. We therefore consider a weaker refinement: we allow voters to break ties *against* the agenda setter’s proposals, as long as they always resolve the same tie (between ultimate outcomes) the same way. Formally,

DEFINITION 5: *An equilibrium is non-capricious if it has the following properties:*

- (a) *The induced mapping from histories to continuation outcomes is deterministic and Markovian (it conditions on the history only through the prevailing default and number of remaining rounds).*
- (b) *For each voter i and pair of distinct policies x and y such that $x \sim_i y$, at every history-proposal pair for which x is the continuation outcome if the proposal is accepted and y is the continuation outcome if the proposal is rejected, voter i either (i) always votes for the proposal or (ii) always votes against the proposal.*

Part (a) slightly weakens the standard definition of Markov perfect equilibrium by allowing strategies, but not the continuation outcomes they induce, to depend on payoff-irrelevant features of the history. Part (b) is more important because it disciplines tie-breaking across histories. Suppose voter i is indifferent between policies x and y and that at history h (respectively, h'), accepting a proposal a (respectively, a') leads to policy x , while rejecting it leads to y . Then if i votes for (respectively, against) proposal a at history h , she must also vote for (respectively, against) proposal a' at history h' . In other words, the manner in which a voter breaks ties only depends on the resulting continuation outcomes. The logic of this restriction is that the particular history is “water under the bridge” and consequently should not affect the voter’s deliberations, even in cases of indifference.¹⁸

¹⁷The desired conclusion follows when the number of rounds exceeds $\lceil [u_A^* - \min_{x \in X} u_A(x)] / \eta_b \rceil$, which allows the agenda setter to migrate the policy from her least favorite to one that achieves a payoff within δ of u_A^* .

¹⁸In settings with generic finite alternatives, non-capriciousness is always satisfied because (i) all equilibria are outcome-equivalent to the specific pure strategy Markov perfect equilibrium from Lemma 1(a) and (ii) voters are never indifferent between *distinct* continuation outcomes (which means the tie-breaking restriction in Definition 5(b) has no bite). Thus, Theorem 3 is a proper generalization of Theorem 1.

We prove that a non-capricious equilibrium exists and that in all such equilibria, the agenda setter has near-dictatorial power whenever the collective choice problem is manipulable.

THEOREM 3: *For any collective choice problem \mathcal{C} :*

- (a) *There exists a non-capricious equilibrium.*
- (b) *The following holds if and only if \mathcal{C} is Manipulable: For every $\delta > 0$, there exists some $T_\delta \in \mathbb{N}$ such that if the game has $T \geq T_\delta$ rounds, then given any initial default, the agenda setter's equilibrium payoff is no lower than $u_A^* - \delta$ in any non-capricious equilibrium.¹⁹*

The general logic of our earlier results continues to govern the proof: once there are sufficiently many rounds, every non-capricious equilibrium outcome must be *nearly* unimprovable. Manipulability of the collective choice problem and continuity of the agenda setter's preferences then imply that the agenda setter's payoff is *nearly* maximized. The complete proof, which involves considerable technical detail, appears in the online Appendix. We illustrate its structure through a full analysis of the standard divide-the-dollar problem in the Appendix. That analysis highlights two additional features. First, even with a small number of rounds (in this case, three), the agenda setter may obtain her favorite policy.²⁰ Second, the main conclusion of Theorem 3 requires non-capricious tie-breaking: for the divide-the-dollar game, there is a Markovian equilibrium with *capricious* tie-breaking in which the agenda setter's power is more limited.

IV. Manipulable Collective Choice Problems

In this section, we demonstrate that the property driving our main results, manipulability, is (generically) satisfied in canonical models of spatial and distributive politics.

A. Spatial Politics

In the canonical spatial model, a policy consists of d continuous components. Formally, the policy space is $X = \mathbb{R}^d$, each player i has an ideal point x_i^* , and $u_i(x) = -(1/2)\|x - x_i^*\|^2$; i.e., players evaluate a policy based on its Euclidean distance from their ideal points.²¹ Given this specification of utilities, the profile of ideal points, $(x_i^*)_{i=1,\dots,n,A} \in \mathbb{R}^{d(n+1)}$, completely characterizes the preference profile.

¹⁹This result can be equivalently stated in ordinal terms: if and only if \mathcal{C} is manipulable, the final policy in any non-capricious equilibrium must itself be close to the agenda setter's favorite policies, X_A^* .

²⁰Theorem 6 generalizes this result to a broad class of distribution problems.

²¹Although we have assumed in Section II that the policy space is compact, it is convenient here to treat it as unbounded to simplify the statement of Theorem 4. However, the proof of that result (also sketched below) establishes the improbability of all policies aside from x_A^* in the *interior* of a compact and convex policy space $X \subseteq \mathbb{R}^d$. Policies on the boundary of such X are also improvable provided that all ideal points are interior, which is plausible when boundary policies represent extreme alternatives.

Our analysis invokes a property we call *non-coplanarity*. For a vector $x \in \mathbb{R}^d$ where $d \geq 3$, let $[x]_{abc} := (x_a, x_b, x_c) \in \mathbb{R}^3$ be the projection of x into the subspace spanned by any three of the dimensions a, b, c . Non-coplanarity entails the following property:

DEFINITION 6: *For $d \geq 3$, the profile of ideal points $(x_i^*)_{i=1, \dots, n, A}$ satisfies non-coplanarity if for every a, b, c , no four players’ projected ideal points, $[x_1^*]_{abc}, \dots, [x_n^*]_{abc}, [x_A^*]_{abc} \in \mathbb{R}^3$, are coplanar.*

When there are only three policy dimensions ($d = 3$), Definition 6 simply states that no four ideal points lie in the same plane. If there are more than three dimensions, it requires the same to be true for *all* three-dimensional projections—that is, when we only consider dimensions a, b, c and ignore the rest.

Our main result shows that spatial collective choice problems are manipulable whenever non-coplanarity is satisfied and, moreover, that this condition holds generically.

THEOREM 4: *Consider a collective choice problem \mathcal{C} with policy space $X = \mathbb{R}^d$, where $d \geq 3$ and players have Euclidean preferences with ideal points $(x_i^*)_{i=1, \dots, n, A}$.*

- (a) *If the profile $(x_i^*)_{i=1, \dots, n, A}$ satisfies non-coplanarity, then the collective choice problem \mathcal{C} is Manipulable.*
- (b) *The set of profiles for which non-coplanarity holds has full Lebesgue measure and is open-dense in $\mathbb{R}^{d(n+1)}$.*

Theorem 4 demonstrates that, when there are at least three policy dimensions, the spatial model generically satisfies manipulability; i.e., all policies other than the agenda setter’s ideal point, x_A^* , are improvable. Equivalently, for a cooperative game in which the decisive coalitions are those containing the agenda setter and some majority of voters, Theorem 4 states that, with three or more dimensions, the core of the spatial model generically contains the agenda setter’s ideal point and nothing else. Given the importance of the spatial model and the elementary geometric argument used to prove Theorem 4, this result may be of independent interest.²²

Together with our prior results (Theorems 1–3), Theorem 4 establishes that the agenda setter can exploit real-time agenda control to obtain her ideal point (exactly or approximately) whenever there are three or more policy dimensions. This conclusion holds even if voters’ preferences are largely congruent. Suppose voters’ ideal points are (relatively) close to each other and the agenda setter’s ideal point lies far outside their convex hull. As long as non-coplanarity holds, the agenda setter inevitably obtains her ideal point, *even if the initial default option lies within that convex hull*. In contrast, for fixed agenda models, McKelvey (1986) shows that an agenda setter can only achieve policies near the initial default, and the distance between the

²²Duggan and Ma (2023) offer a related result for “constrained core points,” which are similar to unimprovable policies. They show that with four or more dimensions, for the class of \mathbf{C}^2 -smooth and strictly concave utility functions, there are no interior constrained core points for a topologically generic class of utility functions.

initial and final policies shrinks to zero as voters' ideal points converge to a single point.

The existence of three policy dimensions is critical for Theorem 4. In the unidimensional case, all policies between x_A^* and the ideal point of the median voter are unimprovable (a consequence of the Median Voter Theorem). For the two-dimensional case, manipulability necessarily fails whenever the agenda setter's ideal point is outside the convex hull of voters' ideal points.²³ We illustrate this failure and elaborate further in the online Appendix.

These facts have implications for *policy bundling*. If the legislature faces a decision involving only one or two dimensions, the agenda setter benefits from introducing a third dimension—even if the associated default is already her ideal—because the collective choice problem thereby becomes manipulable, enabling her to achieve her optima in all three dimensions.

Next, we sketch the geometric argument for Theorem 4(a) in the three-dimensional case. The full proof appears in the online Appendix.

PROOF SKETCH FOR $d = 3$:

Consider a policy x that is not the agenda setter's favorite, x_A^* . We show that x is improvable using a two-step argument. First, we find a policy y near x such that a majority of voters strictly prefer y to x , and moving from x to y generates a *second-order* loss for the agenda setter. Second, we perturb y to some z such that the same majority of voters strictly prefer z to x , but moving from y to z generates a *first-order* gain for the agenda setter, so that the agenda setter also strictly prefers z to x . This argument then establishes that x is improvable.

Step 1: Constructing y . Let $S \subset \mathbb{R}^3$ denote the plane that is tangent to the agenda setter's indifference surface at the point x . As depicted in Figure 3 (left panel), S is orthogonal to the gradient $\nabla u_A(x) = x_A^* - x$. Denote the agenda setter's favorite point in S —henceforth, her *constrained ideal point*—by y_A^* , and observe that, by construction, y_A^* coincides with x . Similarly, let $y_i^* \in S$ denote each voter i 's constrained ideal point and note that the gradient $\nabla u_i(y_i^*) = x_i^* - y_i^*$ is orthogonal to S .

We claim the following:

$$(2) \quad \text{Under non-coplanarity, } \exists y \in S \text{ such that } y \succ_M y_A^*.$$

To prove (2), we make two preliminary observations: (i) for any line in S containing y_A^* , there are at most two voters $i \neq j$ such that y_i^* and y_j^* also lie on that line, and (ii) there can be at most one voter i for whom $y_i^* = y_A^*$. Figure 3 (right panel) illustrates the argument for (i): if there were a third voter $k \notin \{i, j\}$ for whom y_k^* were collinear with $\{y_A^*, y_i^*, y_j^*\}$, then, because all players' gradients at their constrained ideal points are orthogonal to the same plane S , the four unconstrained ideal points $\{x_A^*, x_i^*, x_j^*, x_k^*\}$ would be coplanar, contradicting the assumption of non-coplanarity. The argument for (ii) is similar: if there were two voters $i \neq j$ such that $y_i^* = y_j^* =$

²³This result contrasts with McKelvey's (1976) "chaos theorem," which shows that having two dimensions is sufficient for the majority relation to be globally intransitive (for generic ideal point configurations).




FIGURE 3. CONSTRUCTION OF PLANE S (LEFT) AND NON-COLLINEARITY OF CONSTRAINED IDEAL POINTS (RIGHT)

y_A^* , then the unconstrained ideal points $\{x_A^*, x_i^*, x_j^*\}$ would be collinear and hence coplanar with the ideal point of any other voter.

We now use these observations to deduce (2). Figure 4 (left panel) shows a head-on view of the plane S . By (i), for any line $L \subset S$ through y_A^* , there are at most two voters $i \neq j$ whose constrained ideal points lie on L . Therefore, at least $n - 2$ voters' constrained ideal points lie off of L . The pigeonhole principle implies that of these $n - 2$ constrained ideal points, a strict majority must lie “above” or “below” the line L ; our figure shows $(n - 1)/2$ points above the line. We can then shift y_A^* to some new policy y slightly “above” L so that all of those $(n - 1)/2$ voters strictly prefer y_A^* to y . Moreover, by fact (ii) above, we can pick the direction $y - y_A^*$ so that at least one of voters i and j also strictly prefers y to y_A^* ; in Figure 4 (left panel), this is voter j . Thus, a majority strictly prefers y to $y_A^* = x$. Furthermore, we can choose y arbitrarily close to y_A^* , so the agenda setter only incurs a second-order loss (because y_A^* is the agenda setter’s ideal policy in S).

Step 2: Constructing z . We construct z by perturbing y off of the plane S in the direction $\nabla u_A(x) = x_A^* - x$; see the right panel of Figure 4. Moving from y to z generates a first-order gain for the agenda setter, ensuring that $z \succ_A x$ (because the original move from x to y generates only a second-order loss). Moreover, we can choose the point z close enough to y to ensure that those voters who strictly prefer y to x (a majority) also strictly prefer z to x . ■

B. Distributive Politics

This section shows that collective choice problems involving “distributive politics” are generally manipulable and consequently that all problems become manipulable with the addition of pork or transfers. Moreover, settings with distributive politics satisfy a strong version of manipulability that encompasses all voting rules for which no voter has veto power.




FIGURE 4. PERTURBATIONS $x \rightarrow y \in S$ (LEFT) AND $y \rightarrow z \notin S$ (RIGHT)

We begin with a definition:

DEFINITION 7: *Collective choice problem $C = (X, \{\succsim_i\}_{i=1, \dots, n, A})$ is a distribution problem if it satisfies the following two properties for every policy $x \in X$ and player $i \in N \cup \{A\}$ (where we let u_i represent \succsim_i):*

- (a) **Scarcity:** *If $u_i(x) < \max_{z \in X} u_i(z)$, then there exists either some player $j \neq i$ such that $u_j(x) > \min_{z \in X} u_j(z)$ or some policy y such that $u_k(y) > u_k(x)$ for all players k .*
- (b) **Transferability:** *If $u_i(x) > \min_{z \in X} u_i(z)$, then there exists some policy y such that $u_j(y) > u_j(x)$ for all players $j \neq i$.*

Scarcity captures the notion that utility flows from a limited resource: if the resource is not being used to maximize player i 's payoff, then either it is being used to give some other player more than her minimal utility or there is waste, in which case some other allocation could make all players strictly better-off. Transferability captures the notion that the underlying resource is at least somewhat fungible: if player i enjoys surplus, we can redistribute some of that surplus to everyone else.²⁴ Notably, this definition does not require utility to be fully transferable.

As we show next, the class of distribution problems encompasses a wide range of possibilities, including two canonical cases: any collective choice problem augmented with transfers (including divide-the-dollar) and settings with pork-barrel politics.

EXAMPLE 1 (Divide-the-Dollar/Collective Choices with Transfers): *Consider any collective choice problem with policy space X and utility profile $(u_i)_{i=1, \dots, n, A}$. We*

²⁴Banks and Duggan (2006) call this notion "limited transferability."

augment this problem with monetary transfers. Assuming utility is quasi-linear in money and that each player has an outside option yielding a payoff of zero, the policy space for the resulting transferable-utility collective choice problem is

$$\mathcal{Y} = \left\{ y \in \mathbb{R}_+^{n+1} : \exists x \in X \text{ such that } \sum_{i=1, \dots, n, A} y_i = \sum_{i=1, \dots, n, A} u_i(x) \right\},$$

and the utility functions are $v_i(y) = y_i$. This formulation encompasses both the standard divide-the-dollar problem (e.g., Baron and Ferejohn 1989) as well as settings involving both production decisions and transfers.²⁵

EXAMPLE 2 (Pork Barrel Politics): Suppose there are finitely many public projects $k \in \mathcal{K}$, each of which generates an aggregate benefit $B^k > 0$ and aggregate cost $C^k > 0$. Some projects may be inefficient ($C^k > B^k$). A policy x specifies (i) the projects the group will implement (a subset $\mathcal{I} \subseteq \mathcal{K}$) and (ii) for each of those projects, the distribution of benefits and costs among the players (i.e., $b^k, c^k \in \mathbb{R}_+^{n+1}$ such that $\sum_{i=1, \dots, n, A} b_i^k = B^k$ and $\sum_{i=1, \dots, n, A} c_i^k = C^k$). Player i 's preferences correspond to $u_i(x) = \sum_{k \in \mathcal{I}} (b_i^k - c_i^k)$. Thus, costs and benefits are both perfectly transferable.²⁶

In addition to being ubiquitous, distribution problems are manipulable. The proof is simple: if a policy x is not one of the agenda setter's favorites (i.e., not in X_A^*), scarcity implies that either (i) some other policy y strongly Pareto dominates x or (ii) some voter i obtains more than her minimal utility from policy x , in which case transferability implies that there is a policy y such that both the agenda setter and all voters other than i strictly prefer y to x . In either case, x is obviously improvable; indeed, all players (with the possible exception of i) strictly prefer y to x .

As the preceding argument makes clear, distribution problems satisfy a strong version of manipulability that encompasses any voting rule for which no voter has veto power (rather than just majority rule). We formalize this point as follows. A general voting rule is a collection $\mathcal{D} \subseteq 2^N$ of all winning coalitions $D \subseteq N$, by which we mean that a proposal passes if and only if there exists some $D \in \mathcal{D}$ for which all voters $i \in D$ vote in favor of the proposal.²⁷ This class of voting rules includes quota rules, for which there is a quota q such that $\mathcal{D} = \{D \subseteq N : |D| \geq q\}$, as well as rules that treat voters asymmetrically. A voting rule \mathcal{D} is veto-proof if for every voter $i \in N$, there exists a winning coalition $D \in \mathcal{D}$ such that $D \subseteq N \setminus \{i\}$; in other words, voter i 's support is not necessary for a proposal to pass. We say that a collective choice problem is \mathcal{D} -Manipulable if for every $x \notin X_A^*$, there exists a policy y and a coalition $D \in \mathcal{D}$ such that $y \succ_A x$ and $y \succ_i x$ for every $i \in D$.

²⁵For simplicity, in this example, utility is fully transferable. As noted above, one can weaken this assumption.

²⁶This example blends the model of Baron (1991), who considers a single project with perfectly transferable benefits but a fixed distribution of costs, with that of Ferejohn, Fiorina, and McKelvey (1987); Bernheim, Rangel, and Rayo (2006); and others, who consider multiple projects with fixed distributions of both benefits and costs.

²⁷That is, the proposal passes with support from all voters in some $D \in \mathcal{D}$ regardless of how voters in $N \setminus D$ vote. This definition of winning coalitions has two implications. First, our voting rules are necessarily monotone, namely, $D \in \mathcal{D}$ implies that $D' \in \mathcal{D}$ for all supersets $D' \supseteq D$. Second, it is not problematic if both $D \in \mathcal{D}$ and its complement $N \setminus D \in \mathcal{D}$. For such rules, if all voters in D (respectively, $N \setminus D$) support the proposal and all voters in $N \setminus D$ (respectively, D) oppose it, our definition implies that the proposal passes.

The following result formalizes our observation that distribution problems are manipulable for a broad class of voting rules.

THEOREM 5: *Every distribution problem is \mathcal{D} -Manipulable for every veto-proof voting rule \mathcal{D} .*

We omit a formal proof, as the argument is identical to the one given above for simple-majority rule. Theorem 5, coupled with our prior results, highlights the broad power of agenda control: for any distribution problem and veto-proof voting rule, the agenda setter obtains a near-favorite policy in every (non-capricious) equilibrium, regardless of the initial default, provided there are sufficiently many rounds.²⁸

In fact, for distribution problems, we obtain an even starker result: the agenda setter can obtain her favorite policy (not merely an approximation) even if the game is short, where the required number of rounds depends on the voting rule. Recalling that $u_A^* := \max_{x \in X} u_A(x)$, we have,

THEOREM 6: *Suppose C is a distribution problem satisfying Thin Individual Indifference. Then,*

- (a) *If the voting rule is a quota rule with $q < n$, the agenda setter obtains payoff u_A^* in every non-capricious equilibrium regardless of the initial default for any game with at least $\lceil n/(n - q) \rceil$ rounds.*
- (b) *If the voting rule is veto-proof, the agenda setter obtains payoff u_A^* in every non-capricious equilibrium regardless of the initial default for any game with at least n rounds.*

Theorem 6(a) implies that three rounds suffice for the agenda setter to obtain her favorite policy under any quota rule requiring no more than two-thirds majority. For more demanding quotas, more rounds are required. Theorem 6(b) tells us that n rounds suffice for all veto-proof voting rules (because the agenda setter can appropriate the surplus of at least one voter in each round).²⁹

Theorems 5 and 6 have two broad implications for legislative bargaining. First, because any collective choice problem becomes a distribution problem when bundled with transfers (as in Example 1), our analysis highlights how the introduction of distributional policy instruments can augment an agenda setter's power. Second, equilibrium outcomes need not maximize total surplus. In Example 2, the agenda setter secures a policy that includes all projects, which maximizes total benefits, along with transfers that off-load all costs onto the voters. Plainly, such policies typically involve excess spending relative to the utilitarian optimum.

²⁸As we asserted in Section IIIB, Theorems 1–3 extend, with obvious (minor) adjustments, to general voting rules \mathcal{D} and collective choice problems that are \mathcal{D} -manipulable.

²⁹Notably, the number of rounds required in Theorem 6(a) coincides with the Nakamura number for the given quota rule (Austen-Smith and Banks 2002, pp. 74–82). This coincidence invites the conjecture that, more generally, for distribution problems with arbitrary veto-proof voting rules, the Nakamura number for the given rule equals the number of rounds required for the agenda setter to obtain her favorite policy from all initial defaults. If this conjecture is correct, then it should be possible to tighten the n -round bound in Theorem 6(b).

V. Commitment, Procedures, and Deadlines

This section presents results that either clarify or extend our main findings. Section VA explains how commitments to agendas change the attainable outcomes. Section VB considers other legislative protocols and voting rules. Section VC clarifies the role of a finite horizon and highlights the resulting deadline effect.

A. The Commitment Benchmark

We have studied agenda control without commitment. If the agenda setter could commit to a strategy in the dynamic game, she would do weakly better. Theorems 1–3 imply that if the collective choice problem is manipulable, the agenda setter gains little or nothing from commitment. In this section, we make the same comparison without imposing manipulability. This exercise shows how sequential rationality constrains the agenda setter in nonmanipulable problems, and also connects our work to prior research that assumes she can make commitments.

For simplicity, we restrict attention to Generic Finite Alternatives and assume that $T \geq |X| - 1$. For any positive integer K and policies x and y , we say that y is *K-reachable from x* if there is a sequence of policies $\{a^k\}_{k=0}^K$ such that (i) $y = a^K$ and $x = a^0$ and (ii) $a^k \succ_M a^{k-1}$ for all $k \in \{1, \dots, K\}$. We say that y is *reachable from x* if it is *K-reachable from x* for some K . With commitment, the agenda setter can achieve all such policies.

Fact 1: *If the agenda setter can commit to any strategy, then she can obtain her favorite policy among those that are reachable from x^0 .*

The logic is as follows. Suppose that from an initial default option x^0 , the agenda setter’s favorite reachable policy is y and that $\{a^k\}_{k=0}^K$ is the proposal sequence that reaches it. The agenda setter obtains y by committing to any K -round strategy with the following property: if the default option in some round is a^k , she proposes $a^{\min\{K, k+1\}}$. Clearly, voters necessarily approve the final proposal; approval of every prior proposal follows recursively.

Fact 1 is familiar from the literature on *binary voting trees*—in other words, multistage voting games in which voters decide (through majority voting) to move “left” or “right” in each round, where the resulting path determines the final policy.³⁰ This connection is not coincidental, as the frameworks are closely related.³¹

³⁰ See Austen-Smith and Banks (2005, chap. 4) for a formal definition, as well as a survey of the classic literature (e.g., Black 1958; Farquharson 1969; Miller 1977).

³¹ To appreciate the connection, consider a setting in which x^0 is the default. First, note that we can induce any binary voting tree for which x^0 is a feasible outcome by choosing an appropriate agenda-setter strategy. Formally, we relabel the branches of the voting tree so that the “all left” path leads to x^0 . We then construct the desired strategy by identifying x^0 with the initial default and rightward (respectively, leftward) moves in the tree with acceptances (respectively, rejections) of proposals, where the proposals are inferred from the policies associated with the tree’s terminal nodes. By construction, the resulting strategy is pure and conditions only on the history of past proposals and defaults but not on voting profiles beyond the chosen defaults. Conversely, it is easy to see that any agenda-setter strategy with these properties induces a binary voting tree for which x^0 is a feasible outcome. In light of this equivalence, Fact 1 follows indirectly from the well-known theorem stating that a policy $y \in Y \subseteq X$ is implementable by a binary voting tree with outcomes in Y if and only if y is in the majority-preference “top cycle set” for Y (e.g., Austen-Smith and Banks 2005, Theorem 4.3). To see why, note that any y reachable from x^0 is in

Next, we consider commitments to fixed agendas as in Shepsle and Weingast (1984). Such agendas are equivalent to *default-independent* strategies that prescribe the same proposal for a given round regardless of how the game unfolds. We restate their main result as follows.

Fact 2: *If the agenda setter can commit but is restricted to default-independent strategies, she obtains her favorite policy among those that are 2-reachable from x^0 .*

A comparison of Facts 1 and 2 reveals how this restriction on possible commitments limits the agenda setter's power. The notion of 2-reachability is equivalent to Shepsle and Weingast's concept of being *uncovered* by the initial default, and consequently, the logic of Fact 2 is familiar.³²

We now compare real-time agenda control to these benchmarks. We say that y is *credibly reachable* from x if there is a sequence $\{a^k\}_{k=0}^K$ running from x to y such that $a^k = \phi(a^{k-1})$, where ϕ is the agenda setter's favorite improvement (defined in equation (1)). In other words, each proposal in the chain that reaches y from x is the agenda setter's favorite among policies that are majority-preferred to the preceding proposal. Lemma 1(a) implies the following fact.

Fact 3: *If the agenda setter cannot commit, then she obtains her favorite policy among those that are credibly reachable from x^0 .*

Theorem 1 establishes that commitment has no value if the collective choice problem is manipulable. In that case, the agenda setter's *favorite* policy is not only reachable but also credibly reachable, from all x^0 .³³ However, if the collective choice problem is not manipulable, the agenda setter may do strictly better with commitment, even to a default-independent strategy, as seen in the second example of Section I. Facts 1–3 imply only that commitment to general strategies weakly outperforms both alternative protocols. All weak or strict rankings over these three modes of commitment are feasible as long as they are compatible with this implication.

the top cycle for some $Y \subseteq X$ containing both x^0 and y (e.g., let Y collect the elements of the majority preference chain connecting x^0 to y); because there is a binary tree for Y that delivers y , the equivalence implies that there is a strategy with default x^0 that delivers y as well. The theorem also implies that, if the agenda setter can commit only to pure strategies that do not condition on voting profiles beyond the chosen defaults, she can obtain *only* policies that are reachable from x^0 . We conjecture that the same conclusion holds even if the agenda setter can commit to *any* strategy. If this conjecture is true, then with commitment to any strategy, the agenda setter obtains *precisely* her favorite reachable policy.

³² Suppose y is 2-reachable from x^0 via a sequence $\{a^k\}_{k=0}^2$. Then the agenda setter can achieve y by committing to propose y in the first round and a^1 in every subsequent round. Now suppose the fixed agenda (a^1, \dots, a^T) achieves y . We claim that y is 2-reachable. Let f^t be the equilibrium continuation outcome if a^t passes in round t . Because y is the eventual outcome, we must have $y = f^\tau$ for some τ , and for all $f^t \neq y$, $y \succ_M f^t$. Were y not 2-reachable from x^0 , then it would have to be the case that $x^0 \succ_M y$ (otherwise, the sequence $\{x^0, y, y\}$ would reach y) and for all $f^t \neq x^0$, $x^0 \succ_M f^t$ (otherwise, the sequence $\{x^0, f^t, y\}$ would reach y for some $f^t \succ_M x^0$). But then none of the proposals would pass, a contradiction.

³³ Manipulability allows the agenda setter to obtain her favorite policy only if the number of rounds is sufficiently large. With a small number of rounds, the agenda setter could potentially benefit from commitment because it allows her to exploit a larger class of majority-preference chains.

B. General Legislative Procedures

Legislatures sometimes use alternatives to the amendment procedure studied in previous sections. The best-known alternative is the *successive procedure* (also *closed-rule bargaining*): all proposals include *adjournment provisions* specifying that their acceptance ends deliberation. Another is *open-rule bargaining*: in any round, the agenda setter can “move” the prevailing default; if the motion passes, the legislature adjourns.³⁴ Legislatures also differ with respect to voting rules (e.g., majority versus supermajority requirements).

This section analyzes the implications of real-time agenda control for these alternative procedures. We develop a general framework that allows for an arbitrary voting rule and a general *adjournment protocol*, including as special cases our baseline framework and both alternatives mentioned above. We obtain the following result: *for every preference profile and voting rule, (essentially) all adjournment protocols result in the same equilibrium outcome.* In other words, real-time agenda control nullifies the effect of adjournment provisions, rendering moot the distinction between these various protocols.

We extend the framework of Section II as follows. The definition of a collective choice problem \mathcal{C} is unchanged, except we allow for an even number of voters, n . For simplicity, we focus on settings with Generic Finite Alternatives. Policy selection takes place over finitely many rounds $t = 1, \dots, T$. The agenda setter (respectively, voters) has exclusive proposal (respectively, approval) power. Here, we allow for a wider class of voting rules and adjournment protocols, which we call *generalized amendment procedures*:

- (a) The *voting rule* is defined (as in Section IVB) by a collection $\mathcal{D} \subseteq 2^N$ of winning coalitions. A proposal passes if and only if all voters in some coalition $D \in \mathcal{D}$ approve it. We impose no structure on \mathcal{D} (cf. footnote 27).
- (b) The *adjournment protocol* is defined as follows. In round t , the agenda setter can propose an alternative $\hat{a}^t = (a^t, i) \in X \times \{0, 1\}$, where a^t denotes the policy to supersede the prevailing default x^{t-1} and i denotes the presence or absence of an *adjournment provision*. If $i = 0$, passage makes policy a^t the default in round $t + 1$, as in our baseline model. If $i = 1$, passage ends deliberation and results in the implementation of a^t . In either case, rejection means that x^{t-1} remains the default in round $t + 1$. We allow for the possibility that deliberation changes the set of feasible proposals: for a generic history h , the agenda setter can propose an element of $X(h) \subseteq X \times \{0, 1\}$.

A generalized amendment procedure is *rich* if, at every history h , either $X(h) \subseteq X \times \{0\}$ or $X(h) \subseteq X \times \{1\}$ (or both). In other words, richness rules out protocols where some policy x is available only without an adjournment provision,

³⁴The literature on legislative bargaining has focused on the closed- and open-rule procedures since Baron and Ferejohn (1989), while the literature on agenda setting with fixed agendas has largely focused on the amendment (or Anglo-Saxon) and successive (or Euro-Latin) procedures since Black (1958); Farquharson (1969); and Miller (1977). While the literature models the closed-rule and successive procedures differently, they are essentially equivalent in that, under both procedures, the first accepted proposal is implemented.

while some other policy y is only available with one. Our baseline model and the other procedures mentioned above are rich generalized amendment procedures.³⁵

We show that, for any fixed preference profile, all rich generalized amendment procedures with the same voting rule yield equivalent equilibrium outcomes. To state the formal result, we extend our notion of improvability and favorite improvements to arbitrary voting rules. First, given any policy x , we define the set of policies that some winning coalition prefers to x :

$$M_{\mathcal{D}}(x) := \{y \in X : y = x \text{ or } \exists D \in \mathcal{D} \text{ such that for every } i \in D, y \succ_i x\}.$$

A policy x is \mathcal{D} -improvable if there exists a policy $y \in M_{\mathcal{D}}(x)$ such that $y \succ_A x$; otherwise, policy x is \mathcal{D} -unimprovable. Let $\phi_{\mathcal{D}} : X \rightarrow X$ denote the agenda setter's favorite \mathcal{D} -improvement:

$$\{\phi_{\mathcal{D}}(x)\} := \arg \max_{y \in M_{\mathcal{D}}(x)} u_A(y).$$

The set of \mathcal{D} -unimprovable policies is $E_{\mathcal{D}} := \{x \in X : x = \phi_{\mathcal{D}}(x)\}$. Using this notation, we state our *protocol-equivalence* result:

THEOREM 7: *Suppose the collective choice problem \mathcal{C} satisfies Generic Finite Alternatives and the generalized amendment procedure is rich. For any game with T rounds and initial default policy x^0 , the unique equilibrium outcome is $\phi_{\mathcal{D}}^T(x^0)$. Consequently, for $T \geq |X| - 1$, a policy is an equilibrium outcome if and only if it is an element of $E_{\mathcal{D}}$.*

Thus, with real-time agenda control, equilibrium outcomes do not depend on the adjournment protocol. This result has two noteworthy implications. First, as long as the collective choice problem is \mathcal{D} -manipulable (as defined in Section IVB), the agenda setter is effectively a dictator regardless of the adjournment protocol. Formally, Theorem 7 implies

COROLLARY 1: *Suppose the collective choice problem \mathcal{C} satisfies Generic Finite Alternatives and the generalized amendment procedure is rich. For any game with at least $|X| - 1$ rounds, the agenda setter obtains her favorite policy in every equilibrium regardless of the initial default if and only if \mathcal{C} is \mathcal{D} -Manipulable.*

Second, Theorem 7 contrasts with known results on fixed agendas. In that context, the agenda setter's power depends on the adjournment protocol.³⁶ Specifically, commitment to a fixed successive (or closed-rule) agenda allows her to obtain her

³⁵For the amendment procedure, $X(\cdot) = X \times \{0\}$. For the successive/closed-rule procedure, $X(\cdot) = X \times \{1\}$. For the open-rule procedure, $X(h) = [X \times \{0\}] \cup [\{x(h)\} \times \{1\}]$, where $x(h)$ denotes the prevailing default at history h . Note that the open-rule procedure involves history-dependent feasible sets.

³⁶This theme emerges in Farquharson (1969); Miller (1977); and McKelvey and Niemi (1978); see Chapter 4 of Austen-Smith and Banks (2005) for a survey. More recent work includes Apesteguia, Ballester, and Masatlioglu (2014) and Barberà and Gerber (2017).

favorite policy among those reachable from the initial default (Miller 1977),³⁷ whereas commitment to a fixed amendment agenda only allows her to obtain her favorite 2-reachable policy (Shepsle and Weingast 1984). Theorem 7 shows that this distinction disappears when the agenda setter selects proposals in real time without commitment.

To prove Theorem 7, we adjust the argument for Lemma 1 to account for adjournment provisions; we omit a formal proof but describe the adjustment. Richness of the generalized amendment procedure guarantees that, at every round- t history, if the default option is x^{t-1} , the agenda setter can propose at least one of the following: (i) her favorite \mathcal{D} -improvement $\phi_{\mathcal{D}}(x^{t-1})$ without an adjournment provision or (ii) the “eventual outcome” $\phi_{\mathcal{D}}^{T-t+1}(x^{t-1})$ with an adjournment provision. These options yield the same outcome in any one-round subgame ($t = T$). Therefore, by the backward-induction logic of Lemma 1, both proposals lead to the outcome $\phi_{\mathcal{D}}^{T-t+1}(x^{t-1})$ regardless of how many rounds remain. Theorem 7 then follows from our observation that $\bigcup_{x^0 \in X} \phi_{\mathcal{D}}^T(x^0) = E_{\mathcal{D}}$ for $T \geq |X| - 1$.³⁸

C. The Role of Deadlines

We think of our framework as representing negotiations, starting at date 0, over the policy that will prevail at date τ .³⁹ Negotiations obviously cannot continue past the implementation date. Given the inherent frictions arising either from institutional constraints or simply from speed-of-light latency considerations, we treat each round of bargaining as requiring at least $\Delta > 0$ units of time. Consequently, there can be at most $T = \lfloor \tau/\Delta \rfloor$ rounds of deliberation. Hence, we follow the prior literature on agenda setting by modeling finite-round processes.⁴⁰ The finite deadline effectively provides the agenda setter with a bit of commitment power: the process ends with a take-it-or-leave-it offer. In this section, we investigate the role of this deadline.

In contrast to our analysis, Diermeier and Fong (2011) and Anesi and Seidmann (2014) model agenda control with an infinite horizon. For concreteness, we focus on the latter analysis, which differs from ours in one key respect: there is no exogenous terminal round T . Instead, bargaining endogenously terminates only when the agenda setter either (a) proposes the prevailing default option or (b) makes a proposal that is rejected. Payoffs are undiscounted and determined by the policy implemented at termination; a nonterminating path is the worst outcome for all

³⁷This observation essentially restates Fact 1 for binary voting trees. The logic is as follows: if policy y is reachable from default x through the sequence $\{a^k\}_{k=0}^K$, the agenda setter can obtain y in $T = K$ rounds through the fixed agenda where the first proposal is a^K , the second is a^{K-1} , and so on, and each proposal includes an adjournment provision.

³⁸The following example further illustrates the role that richness plays in the proof of Theorem 7. Suppose the options are $\{w, x, y, z\}$ and that preferences are as depicted in Figure 1. With our baseline amendment procedure, the agenda setter obtains her favorite policy w in every equilibrium provided there are at least three rounds, regardless of the initial default. Now consider the nonrich generalized amendment procedure with simple majority rule for which, at every history, the agenda setter can propose policies w, y, z only with adjournment provisions and policy x only without one. As the reader may verify, starting from initial default $x^0 = z$, regardless of the number of rounds, $(y, 1)$ is the outcome (the agenda setter proposes it and it passes), contrary to Theorem 7. The logic of the proof fails because $y = \phi(z)$ is not available without adjournment and $x = \phi^2(z)$ is not available with adjournment.

³⁹The legislature presumably undertakes many such negotiations (one for each future date) in parallel.

⁴⁰The classical literature on agenda setting discussed in Section VA studies fixed agendas with a finite sequence of proposals or binary voting trees of finite depth.

players.⁴¹ The solution concept is pure strategy Markov perfect equilibrium with as-if pivotal voting (henceforth, MPE), meaning in this context that strategies are stationary and condition only on the prevailing default.

To illustrate the implications of this protocol, we revisit our introductory example from Section I (see Figure 1). We showed in that section that, assuming a finite horizon, the agenda setter obtains her favorite policy (w) starting from any initial default if there are three or more rounds. In contrast, with the infinite-horizon protocol (and its termination rule), the agenda setter can do no better than y when starting from an initial default of z or y . We sketch the logic by considering each default option.

- **Default Option of w :** *In every MPE, w is implemented.* The agenda setter can guarantee w by proposing it, and any other outcome (a different policy or a nonterminating cycle) is worse for the agenda setter.
- **Default Option of x :** *In every MPE, the agenda setter proposes w and it passes.* Voters predict that passing w enacts it (by the preceding logic). Because $w \succ_M x$, a majority of voters approve the proposal.
- **Default Option of y :** *In every MPE, y is implemented.* A majority of voters will not approve either w or x : by the preceding logic, they predict that passage of either will result in the implementation of w , and $y \succ_M w$.
- **Default Option of z :** *In every MPE, the agenda setter proposes y and it passes.* Voters predict that passing y leads to its enactment, and $y \succ_M z$, so a majority of voters approve the proposal. As with a default of y , a majority of voters will not approve either w or x .

When T is finite, the agenda setter can propose x in the terminal round *with a de facto commitment not to amend it further in the future*, even though x is improvable. In contrast, with an infinite horizon, the preceding discussion reveals that the agenda setter can never make a similar commitment. The freedom to reconsider policies indefinitely generates additional sequential rationality constraints, significantly weakening the agenda setter's power.

Diermeier and Fong (2011) and Anesi and Seidmann (2014) show generally that, in infinite-horizon settings with Generic Finite Alternatives, (a) the set of MPE outcomes corresponds to the *von Neumann-Morgenstern stable set*, V , and (b) in every MPE with initial default x^0 , the agenda setter obtains her favorite policy y among those in V satisfying $y \succ_M x^0$.⁴² The stable set necessarily includes all unimprovable policies (i.e., $E \subseteq V$) but is typically larger. Consequently, (a) more policies can arise in equilibrium with an infinite horizon than with a long finite horizon, and (b) the agenda setter is weakly worse off with an infinite horizon than with a single proposal round.

⁴¹ Instead of adopting this termination rule, Diermeier and Fong (2011) study the patient limit of an infinite-horizon model in which players maximize the discounted payoff from the infinite sequence of equilibrium default policies. We group them and Anesi and Seidmann (2014) together, as their models have identical equilibrium outcomes in settings with Generic Finite Alternatives.

⁴² A policy x is *dominated* by a policy y if $y \succ_A x$ and $y \succ_M x$. A set $V \subseteq X$ is *stable* if no $x \in V$ is dominated by another $y \in V$ ("internal stability"), while every $x \notin V$ is dominated by some $y \in V$ ("external stability"). Diermeier and Fong (2012) show that there exists a unique stable set in the present context.

We view the finite- and infinite-horizon models as having different domains of applicability. A leading interpretation of the infinite-horizon model is that it represents legislative decision-making with an uncertain deadline: deliberations end in round t with probability $(1 - \beta)$, the realization occurring after that round’s proposal and votes; analysis focuses on the limiting case of $\beta \rightarrow 1$.⁴³ For such settings, the deadline is both uncertain and unbounded. As articulated earlier, our perspective is that there often is a known upper bound on the number of rounds, particularly for negotiations over time-indexed policies. Even when the deadline is uncertain *ex ante*, our results apply if it becomes known during deliberations. Theorem 3 implies that revelation of the deadline T_δ rounds in advance allows the agenda setter to obtain a payoff within δ of her maximum. More starkly, Theorem 6 implies that in (essentially) any distribution problem, three rounds of advance notice concerning the deadline allows the agenda setter to obtain her favorite policy. Both of these conditions strike us as modest, particularly when negotiations are relatively frictionless: if players learn the deadline $\epsilon > 0$ units of time in advance, then with sufficiently short proposal rounds, there will be at least three rounds left and potentially many more.

Setting aside the question of applicability, our analysis allows us to characterize the effect of the number of proposal rounds on the agenda setter’s power. To that end, let $U_T(x^0)$ denote the agenda setter’s equilibrium payoff in the finite-horizon game with T rounds and initial default x^0 , and let $U_\infty(x^0)$ denote that payoff in the infinite-horizon game.⁴⁴ We obtain the following characterization by analyzing properties of the unimprovable and stable sets.

THEOREM 8: *Suppose the collective choice problem \mathcal{C} satisfies Generic Finite Alternatives. For every $x^0 \in X$ and $T' > T \geq 1$, we have*

$$U_{T'}(x^0) \geq U_T(x^0) \geq U_\infty(x^0).$$

Moreover, exactly one of the following two statements holds:

- (a) *There exists some $x^0 \in X$ such that $U_T(x^0) > U_1(x^0) > U_\infty(x^0)$ for all $T \geq 2$.*

⁴³An alternative interpretation is that the infinite-horizon model with discounting and no termination, as in Diermeier and Fong (2011), captures settings in which the legislature chooses policies for a sequence of calendar dates $t \in \{1, 2, \dots\}$. Specifically, the winning option for round t , x^t , becomes the policy for that period and serves as the default for $t + 1$. Accordingly, policies do not vary over time (i.e., $x^\tau = x^t$ for all $\tau \geq t + 1$) unless there are further amendments. Under this interpretation, the “legislative session” at each calendar date t consists of a single proposal round. In contrast, our view is that legislatures can negotiate over *policy trajectories* specifying continuation paths (x^t, x^{t+1}, \dots) of time-indexed policies for all future dates. Examples include phase-in and sunset provisions. In other words, default trajectories are not necessarily constant as the preceding perspective assumes. We also take the view that each time-indexed session should consist of multiple proposal rounds rather than one. Modeling dynamic collective choice in this manner effectively makes the problem separable across periods, in which case our separate solutions for all of the time-indexed-policy selection problems collectively provide a solution to the full dynamic collective choice problem. See Section 6 of Bernheim, Rangel, and Rayo (2006) for an elaboration of this perspective.

⁴⁴Formally, Lemma 1 implies that $U_T(x^0) = u_A(\phi^T(x^0))$, and the results of Diermeier and Fong (2011, 2012) and Anesi and Seidmann (2014) imply that $U_\infty(x^0) = u_A(\psi(x^0; V))$, where $\psi(x^0; V)$ is the agenda setter’s favorite improvement on x^0 among policies in the stable set, V .

(b) For all $x^0 \in X$ and $T \geq 2$, $U_T(x^0) = U_1(x^0) = U_\infty(x^0)$.

Theorem 8 shows that the agenda setter either (a) benefits most from having multiple (but finite) rounds and least from having infinite rounds or (b) is indifferent about the number of rounds. Thus, her payoff is nonmonotone in the number of rounds, except when she cannot benefit from an ability to revisit any one-round proposal. This nonmonotonicity suggests that an agenda setter may benefit from creating a deadline even if one does not arise naturally. She might accomplish this objective by creating a “crisis” to instill urgency or by bundling the policy issue of interest with a separate time-indexed matter (e.g., a deadline for raising the debt ceiling).

VI. Conclusion

We have shown that agenda setters have dictatorial power in collective choice problems with two features. The first is that the agenda setter proposes policies in real time without commitment, tailoring her current proposal to the prevailing default option. The second is a widely satisfied manipulability condition that ensures the existence of one-step improvements.

Our analysis contributes to a literature that seeks to understand why legislative institutions concentrate political power in the hands of agenda setters and why majority rule may not be an effective safeguard. To this end, our results also highlight how the agenda setter benefits from bundling policies with transfers and pork or by linking unrelated policy issues. Finally, we have shown that when the agenda setter makes proposals in real time, many bargaining protocols have equivalent implications for equilibrium outcomes.

While our analysis addresses several important questions, it leaves others unanswered. Inasmuch as real-time agenda control often confers dictatorial power both when voters are myopic and when they are sophisticated, one might expect the same conclusion to follow if voters are partially sophisticated. We suspect that the validity of this conjecture depends on the form of partial sophistication being modeled and whether it is heterogeneous across voters. Separately, it would be worthwhile to investigate settings in which a committee controls the agenda and to explore the implications of the committee’s rules for the ultimate legislative outcome. Finally, voters may form voting blocs so that an agenda setter cannot play them off against each other; it would be useful to determine the degree to which such blocs attenuate the agenda setter’s power. We hope to address these and other questions in future work.

APPENDIX A

A1. Proof of Lemma 1

We prove that $f_T(x^0) = \{\phi^T(x^0)\}$ by induction.

Base Step: If $T = 1$, then $f_T(x^0) = \{\phi(x^0)\}$ for all $x^0 \in X$.

Rejection of any proposal results in x^0 being chosen. Therefore, in every equilibrium, any proposal $y \succ_M x^0$ passes with probability 1. Thus, if x^0 is

improvable, proposing $\phi(x^0)$ with probability 1 is uniquely optimal for the agenda setter and hence must occur in every equilibrium. If x^0 is unimprovable, then any proposal that the agenda setter makes in equilibrium results in x^0 being chosen. In both cases, the result follows.

Inductive Step: Given any $T \in \mathbb{N}$ and $x^0 \in X$, if $f_{T-1}(x) = \{\phi^{T-1}(x)\}$ for all $x \in X$, then $f_T(x^0) = \{\phi^T(x^0)\}$.

Consider the procedure with T rounds and initial default x^0 . By subgame perfection, $f_{T-1}(x)$ is the set of outcomes arising with positive probability in any equilibrium in any subgame with $T - 1$ rounds in which x is the prevailing default after the first round. Therefore, by the inductive hypothesis, passage of proposal y in the first round results in outcome $\phi^{T-1}(y)$ and rejection of that proposal results in outcome $\phi^{T-1}(x^0)$. Thus, in every equilibrium, any proposal y where $\phi^{T-1}(y) \succ_M \phi^{T-1}(x^0)$ will pass with probability 1. Therefore,

- if $\phi^{T-1}(x^0)$ is improvable, then in every equilibrium, the agenda setter proposes some y whose continuation outcome $\phi^{T-1}(y)$ coincides with $\phi(\phi^{T-1}(x^0))$, which is $\phi^T(x^0)$. Note that $\phi(x^0)$ is one such proposal because ϕ and ϕ^{T-1} commute. If multiple such proposals exist, then she may randomize among them.
- if $\phi^{T-1}(x^0)$ is unimprovable, then it is an element of E and therefore, $\phi^{T-1}(x^0) = \phi^T(x^0)$. Thus, any proposal that the agenda setter makes in equilibrium results in $\phi^{T-1}(x^0)$ being chosen; again, there may be multiple such proposals.

In either case, $f_T(x^0) = \{\phi^T(x^0)\}$.

Lemma 1(a) follows immediately from the above; (b) follows immediately from the above and the definition of E . For (c), observe that the inclusion $\bigcup_{x^0 \in X} f_T(x^0) \supseteq E$ follows immediately from (b), while the opposite inclusion $\bigcup_{x^0 \in X} f_T(x^0) \subseteq E$ follows from the above, together with the fact that $\phi^T(x^0) \in E$ for $T \geq |X| - 1$. ■

A2. Details and Proofs for Theorem 2

Our argument proceeds in three steps. First, we present a *uniform improvement* lemma that plays a critical role in showing that the agenda setter can obtain a payoff within δ of her highest payoff, u_A^* , with a uniform bound on the number of rounds. The second step formalizes the assertion that Thin Individual Indifference characterizes collective choice problems that admit arbitrarily fine generic ϵ -grids. With these steps in place, we then prove Theorem 2.

A Uniform Improvement Lemma.—For each $\delta > 0$, define $\Gamma_\delta := \{x \in X \mid u_A^* \geq u_A(x) + \delta\}$. We say that policies in Γ_δ are δ -suboptimal for the agenda setter and policies in $X \setminus \Gamma_\delta$ are δ -optimal for her. For each $x \in X$ and $\eta > 0$, define

$$Q(x, \eta) := \{y \in X \mid u_A(y) \geq u_A(x) + \eta \text{ and } \exists \text{ majority } S \subseteq N$$

$$\text{such that } u_i(y) \geq u_i(x) + \eta \forall i \in S\}$$

to be the set of policies that lead to a utility improvement of at least η for some winning coalition. If $Q(x, \eta) \neq \emptyset$, then we say that x is η -improvable.

LEMMA 2: *Suppose the collective choice problem \mathcal{C} is Manipulable. Then for every $\delta > 0$, there exists $\eta_\delta > 0$ such that $Q(x, \eta_\delta) \neq \emptyset$ for all $x \in \Gamma_\delta$.*

Manipulability implies that any policy that is δ -suboptimal for the agenda setter must be improvable, but does not specify how much the agenda setter and a winning coalition of voters gain from that improvement. Lemma 2 asserts that for each $\delta > 0$, there is a uniform threshold η_δ such that any policy that is δ -suboptimal for the agenda setter must also be η_δ -improvable. This uniformity will be important for establishing the uniform bounds on the number of rounds needed for the agenda setter to achieve δ -optimality in Theorem 2.

PROOF OF LEMMA 2:

Let $\delta > 0$ be given. Suppose that $\Gamma_\delta \neq \emptyset$, for otherwise the lemma is vacuously true. Define the map $\eta^* : \Gamma_\delta \rightarrow \mathbb{R}_+$ by

$$(A1) \quad \eta^*(x) := \sup_{\eta \in \mathbb{R}_+} \eta \quad \text{s.t.} \quad Q(x, \eta) \neq \emptyset.$$

The supremum in (A1) is attained because the correspondence $(x, \eta) \rightarrow Q(x, \eta)$ is upper hemicontinuous and compact-valued, and $x \in Q(x, 0)$. Moreover, because \mathcal{C} is Manipulable, for each $x \in \Gamma_\delta$, there exists some $\eta_x > 0$ such that $Q(x, \eta_x) \neq \emptyset$. It follows that $\eta^*(\Gamma_\delta) \subseteq (0, \infty)$.

We claim that η^* is lower semicontinuous. Let $x^* \in \Gamma_\delta$ be given and take any sequence $\{x_n\} \subset \Gamma_\delta$ satisfying $x_n \rightarrow x^*$. Because preferences are continuous, for every $\epsilon > 0$, there exists some $N_\epsilon > 0$ such that $n \geq N_\epsilon$ implies $|u_i(x_n) - u_i(x^*)| < \epsilon$ for all players i . Letting $\epsilon \in (0, \eta^*(x^*))$, which is possible because $\eta^*(x^*) > 0$, we therefore have $Q(x_n, \eta^*(x^*) - \epsilon) \neq \emptyset$ for $n \geq N_\epsilon$. Hence, $\eta^*(x_n) \geq \eta^*(x^*) - \epsilon$ for $n \geq N_\epsilon$. Sending $\epsilon \rightarrow 0$, we obtain $\liminf_{n \rightarrow \infty} \eta^*(x_n) \geq \eta^*(x^*)$, which establishes the claim.

To conclude the proof, note that $\eta_\delta := \min_{x \in \Gamma_\delta} \eta^*(x)$ is well defined because Γ_δ is compact and η^* is lower semicontinuous, is strictly positive because $\eta^*(\Gamma_\delta) \subseteq (0, \infty)$, and satisfies $Q(x, \eta_\delta) \neq \emptyset$ for all $x \in \Gamma_\delta$ by construction (recall that the supremum in (A1) is attained). ■

Generic ϵ -Grids and Thin Individual Indifference.—Here, we formalize the assertion that Thin Individual Indifference (Definition 4) characterizes collective choice problems that admit arbitrarily fine generic ϵ -grids. The following formalizes what it means to admit arbitrarily fine grids:

DEFINITION 8: *Collective choice problem $\mathcal{C} = (X, \{\succsim_i\}_{i=1, \dots, n, A})$ is Finitely Approximable if, for every $x \in X$ and $\epsilon > 0$, there exists a generic ϵ -grid X_ϵ such that $x \in X_\epsilon$.*

Definition 8 requires not only that a generic ϵ -grid $X_\epsilon \subseteq X$ exists for every $\epsilon > 0$ but also that such a grid can be constructed so as to contain any prespecified point x in the ambient policy space X . Lemma 3 shows that finite approximability is characterized by Thin Individual Indifference.

LEMMA 3: *A collective choice problem \mathcal{C} is finitely approximable if and only if it satisfies Thin Individual Indifference.*

As Lemma 3’s proof is technical and involved, we relegate it to the online Appendix.

Proof of Theorem 2.—For any $\epsilon > 0$ and generic ϵ -grid X_ϵ , we denote the corresponding discretized collective choice problem by $\mathcal{C}_\epsilon := (X_\epsilon, \{\succ_i\}_{i=1, \dots, n, A})$. We define two maps analogous to the definitions in Section IIIB. The agenda setter’s favorite improvement within grid X_ϵ , denoted by $\phi(\cdot; X_\epsilon) : X_\epsilon \rightarrow X_\epsilon$, is

$$(A2) \quad \{\phi(x; X_\epsilon)\} := \arg \max_{y \in M(x) \cap X_\epsilon} u_A(y),$$

where, as in Section IIIB, $M(x) := \{y \in X : y \succ_M x \text{ or } y = x\}$. The second map is $f_T(\cdot; X_\epsilon) : X_\epsilon \rightrightarrows X_\epsilon$, which denotes the equilibrium outcome correspondence (as defined in Section IIIB) for \mathcal{C}_ϵ . With these definitions in hand, we prove each direction of Theorem 2 in turn.

Sufficiency of Manipulability for Approximate Dictatorial Power: Suppose that \mathcal{C} is Manipulable. As \mathcal{C} satisfies Thin Individual Indifference, Lemma 3 assures that for each $\epsilon > 0$, there exists a generic ϵ -grid X_ϵ .

Let $\delta > 0$ be given, and let $\eta_\delta > 0$ be as defined in Lemma 2. Let $\epsilon_\delta > 0$ be such that

$$(A3) \quad \max_{i \in N \cup \{A\}} \max_{x \in X} \max_{y \in B_{\epsilon_\delta}(x)} |u_i(x) - u_i(y)| \leq \frac{\eta_\delta}{2},$$

where $B_{\epsilon_\delta}(x) := \{y \in X : d(y, x) < \epsilon_\delta\}$, noting that such an $\epsilon_\delta > 0$ exists because each u_i is uniformly continuous (being that X is compact). For each $\epsilon < \epsilon_\delta$, $X_\epsilon \cap B_{\epsilon_\delta}(x) \neq \emptyset$ for all $x \in X$ by construction. Therefore, (A3) implies that

$$(A4) \quad \max_{i \in N \cup \{A\}} \max_{x \in X} \min_{y \in X_\epsilon} |u_i(x) - u_i(y)| \leq \frac{\eta_\delta}{2}.$$

Henceforth, we consider $\epsilon < \epsilon_\delta$.

We first claim, building on Lemma 2, that every policy in X_ϵ that is δ -suboptimal for the agenda setter is $\eta_\delta/2$ -improvable in X_ϵ ; namely, there exists an alternative in X_ϵ that leads to a utility increase of at least $\eta_\delta/2$ for herself and some majority of voters. Formally,

$$(A5) \quad \text{For every } x \in X_\epsilon \cap \Gamma_\delta, \quad Q(x, \eta_\delta/2) \cap X_\epsilon \neq \emptyset.$$

To see why (A5) is true, let $x \in X_\epsilon \cap \Gamma_\delta$ be given. First, by Lemma 2, there exists some $y' \in X$ such that $u_i(y') \geq u_i(x) + \eta_\delta$ for every $i \in S \cup \{A\}$, where $S \subseteq N$ contains some majority of voters. Second, (A4) assures that there exists some $y \in X_\epsilon$ such that $|u_i(y') - u_i(y)| \leq \eta_\delta/2$ for every i . Combining these observations, we conclude that $u_i(y) \geq u_i(x) + \eta_\delta/2$ for all $i \in S \cup \{A\}$.

Because $y \succ_M x$ above, an important implication of (A5) is that

$$(A6) \quad \text{for every } x \in X_\epsilon \cap \Gamma_\delta, \quad u_A(\phi(x; X_\epsilon)) \geq u_A(x) + \frac{\eta_\delta}{2}.$$

We use this fact to prove the theorem: there exists some (uniform) $T_\delta \in \mathbb{N}$ such that, if there are $T \geq T_\delta$ rounds, then the agenda setter’s payoff is no lower than $u_A^* - \delta$ in every equilibrium for any generic ϵ -grid X_ϵ with $\epsilon < \epsilon_\delta$.⁴⁵ To put it formally, there exists T_δ such that for every $T \geq T_\delta$ and $\epsilon < \epsilon_\delta$,

$$\bigcup_{x^0 \in X_\epsilon} f_T(x^0; X_\epsilon) \subseteq X_\epsilon \setminus \Gamma_\delta.$$

If $x^0 \in X_\epsilon \setminus \Gamma_\delta$, the fact that $\bigcup_{x^0 \in X_\epsilon} f_T(x^0; X_\epsilon) \subseteq X_\epsilon \setminus \Gamma_\delta$ for all $T \geq 1$ follows from applying Lemma 1 to C_ϵ , noting that $\phi(x; X_\epsilon) \succ_A x$ for every x . Thus, we consider $x^0 \in X_\epsilon \cap \Gamma_\delta$. We denote the payoff difference between the agenda setter’s favorite and least favorite policies by $\Delta := u_A^* - \min_{y \in X} u_A(y)$, which is well defined and finite because u_A is continuous and X is compact. Correspondingly, define $T_\delta := \lceil 2\Delta/\eta_\delta \rceil \in \mathbb{N}$. Suppose, toward a contradiction, that $y := f_{T_\delta}(x^0; X_\epsilon) \in X_\epsilon \cap \Gamma_\delta$. Then, it follows that

$$\begin{aligned} u_A(\phi(y; X_\epsilon)) - u_A(x^0) &\geq u_A(y) - u_A(x^0) + \frac{\eta_\delta}{2} \\ &= u_A(\phi^{T_\delta}(x^0; X_\epsilon)) - u_A(x^0) + \frac{\eta_\delta}{2} \\ &= \sum_{t=1}^{T_\delta} \left[u_A(\phi^t(x^0; X_\epsilon)) - u_A(\phi^{t-1}(x^0; X_\epsilon)) \right] + \frac{\eta_\delta}{2} \\ &\geq T_\delta \cdot \frac{\eta_\delta}{2} + \frac{\eta_\delta}{2} \\ &\geq \Delta + \frac{\eta_\delta}{2}, \end{aligned}$$

where the first line is by (A6), the second line is by Lemma 1 applied to C_ϵ , the third line is an identity, the fourth line is by another application of (A6) to each term in the sum (noting that $\phi^{T_\delta}(x^0; X_\epsilon) \in X_\epsilon \cap \Gamma_\delta$ implies that $\phi^t(x^0; X_\epsilon) \in X_\epsilon \cap \Gamma_\delta$ for all $t < T_\delta$), and the final line is by definition of T_δ . However, given that $\eta_\delta > 0$, this inequality contradicts the definition of Δ . We conclude that $y \in X_\epsilon \setminus \Gamma_\delta$, as desired.

⁴⁵We note that this statement does not follow from Lemma 1. Lemma 1 implies that if there are $T \geq |X_\epsilon| - 1$ rounds, then $\bigcup_{x^0 \in X_\epsilon} f_T(x^0; X_\epsilon) = E(X_\epsilon)$, where $E(X_\epsilon) := \{x \in X_\epsilon : x = \phi(x; X_\epsilon)\}$ denotes the set of unimprovable policies in C_ϵ . We know from (A5) that $E(X_\epsilon) \subseteq X_\epsilon \setminus \Gamma_\delta$, i.e., any policy that is unimprovable in C_ϵ must be δ -optimal for the agenda setter. It would then follow that the agenda setter’s payoff is at least $u_A^* - \delta$ when there are $T \geq |X_\epsilon| - 1$ rounds; as $\epsilon \rightarrow 0$, this argument would then require the number of rounds to grow without bound.

Necessity of Manipulability for Approximate Dictatorial Power: Suppose that \mathcal{C} is not Manipulable. Then, there exists an unimprovable policy x and $\delta > 0$ such that $u_A(x) < u_A^* - \delta$. As \mathcal{C} satisfies Thin Individual Indifference, Lemma 3 assures that there exists an $\epsilon_\delta > 0$ such that, for all $\epsilon \in (0, \epsilon_\delta)$, there is a generic ϵ -grid X_ϵ for which $x \in X_\epsilon$. Observe that x must also be unimprovable in the corresponding discretized collective choice problem \mathcal{C}_ϵ . Applying Lemma 1 to this discretized problem reveals that for every number of rounds, the equilibrium outcome starting from initial default $x^0 = x$ is x itself: $f_T(x; X_\epsilon) = \{x\}$ for every $T \in \mathbb{N}$. The agenda setter then attains a payoff of $u_A(x) < u_A^* - \delta$, failing to achieve approximate dictatorial power regardless of the number of rounds. ■

A3. The Divide-the-Dollar Problem

Herein, we specialize to the standard “divide-the-dollar” problem, in which the policy space is $X = \Delta^{n+1}$ and a policy $x := (x_1, \dots, x_n, x_{n+1})$ is a division of the dollar; the first n indices are the shares of the n voters, and x_{n+1} is that of the agenda setter. Each player has selfish risk-neutral preferences, so $u_i(x) = x_i$. The legislature begins with an initial default option x^0 and, as in our baseline analysis, uses simple majority rule in each of finitely many rounds.

In this context, we elucidate two features of our general analysis. First, we construct a non-capricious equilibrium in which the agenda setter appropriates the entire dollar whenever there are three or more rounds. Second, we highlight how our dictatorial power result (Theorem 3) does not apply to equilibria with capricious tiebreaking: regardless of the number of rounds, there exists a pure strategy Markov perfect equilibrium with capricious tiebreaking in which the agenda setter fails to appropriate the entire dollar.

A Non-Capricious Equilibrium: To describe a non-capricious equilibrium, we adapt the agenda setter’s favorite improvement operator ϕ from Section IIIB to this setting (which features indifferences). For default policy x , let $\beta(x)$ denote the policy that sets the $(n - 1)/2$ largest elements (among the first n elements) to zero and reallocates that portion of the dollar to the agenda setter; in the event of ties, $\beta(x)$ selects the group of voters with this size with the lowest player indices. More precisely, let $G^0(x) := \emptyset$, and define $G^k(x)$ inductively for $k \in \{1, \dots, n\}$ as follows:

$$G^k(x) := G^{k-1}(x) \cup \left\{ j \in N : j = \min \left\{ \arg \max_{i \in N \setminus G^{k-1}(x)} x_i \right\} \right\}.$$

Observe that $G^k(x)$ identifies the k voters who have the highest shares in default policy x (and breaks ties in favor of those with lower player labels). We define the policy $\beta(x)$ as

$$(\beta(x))_i := \begin{cases} 0, & \text{if } i \in G^{(n-1)/2}(x); \\ x_i, & \text{if } i \in N \setminus G^{(n-1)/2}(x); \\ x_{n+1} + \sum_{j \in G^{(n-1)/2}(x)} x_j, & \text{if } i = n + 1. \end{cases}$$

This operator adapts the favorite improvement operator ϕ to this setting: among policies that a majority of voters weakly prefer to x , $\beta(x)$ is one of the agenda setter's favorites. Observe that for any policy x , $\beta^2(x)$ extracts the shares from all but one voter—the one who has the lowest share in policy x —and $\beta^3(x)$ yields the agenda setter the entire dollar.

We now construct a non-capricious equilibrium in which the agenda setter obtains the entire dollar if there are $T \geq 3$ rounds. Consider the following strategy profile: in each round $t \in \{1, \dots, T\}$, if the prevailing default is x , then (i) the agenda setter proposes $\beta(x)$, and (ii) each voter $i \in N$ votes in favor of a proposal y if and only if $\beta^{T-t}(y) \succsim_i \beta^{T-t}(x)$; namely, she weakly prefers the continuation outcome from acceptance to the continuation outcome from rejection. As no player has a strictly profitable deviation and the strategy profile is pure and Markovian, this defines a non-capricious equilibrium by construction.

We illustrate the path of play in this equilibrium using the following example:

EXAMPLE 3: *Suppose that there are three voters and the default option x^0 is such that $x_1^0 > x_2^0 > x_3^0 > 0$. Then, $\beta(x^0) = (0, x_2^0, x_3^0, 1 - x_2^0 - x_3^0)$, $\beta^2(x^0) = (0, 0, x_3^0, 1 - x_3^0)$, and $\beta^3(x^0) = (0, 0, 0, 1)$. Observe that voters 1 and 2 approve each equilibrium-path proposal: they are indifferent between $\beta^2(x^0)$ and $\beta^3(x^0)$ and thus anticipate that rejecting either the first or second on-path offer nevertheless results in both of them obtaining zero surplus.*

The above construction demonstrates a particular non-capricious equilibrium in which the agenda setter obtains the entire dollar within three rounds. Theorem 6 further implies that *all* non-capricious equilibria of this distribution problem share this property.

A Capricious Equilibrium: We now show, using a setting with $n = 3$ voters, that the dictatorial power conclusion of Theorem 3 does not apply to equilibria with capricious tiebreaking.

Consider a strategy profile that differs from that described above only with respect to voters' tiebreaking rule: voters now resolve indifference in favor of the agenda setter's proposal if and only if it is in the final or penultimate round (i.e., $t \in \{T - 1, T\}$) and otherwise break ties in favor of the prevailing default option. Observe that this strategy profile satisfies Definition 5(a), as it is pure and Markovian. However, it violates Definition 5(b) because, for any pair of continuation outcomes, the tie-breaking decision conditions on the current round.

We claim that, for *any* initial default x^0 and number of rounds $T \geq 2$, this strategy profile (i) results in the outcome $\beta^2(x^0)$ and (ii) is an equilibrium. Consequently, given any default x^0 in which all voters have positive shares (as in Example 3) and any number of rounds, the agenda setter fails to appropriate the entire dollar in this (capricious) equilibrium.

To see that the outcome is $\beta^2(x^0)$, observe that—as in Example 3—the outcome is $\beta(x^0)$ if $T = 1$ and $\beta^2(x^0)$ if $T = 2$. Suppose now that $T = 3$. Voters anticipate that accepting the initial on-path proposal $\beta(x^0)$ results in outcome $\beta^3(x^0)$, whereas rejecting it leads to $\beta^2(x^0)$. As the two voters with the largest shares (and lowest

indices) in x^0 both receive zero shares under both outcomes,⁴⁶ our capricious tie-breaking rule stipulates that they both vote against the initial proposal—unlike in Example 3—resulting in an on-path outcome of $\beta^2(x^0)$. It is then easy to see by induction that the same outcome arises for all $T \geq 3$.

We now argue that this strategy profile is an equilibrium. It suffices to consider the agenda setter's incentives in rounds $t \leq T - 2$.⁴⁷ Consider round $T - 2$, and let the prevailing default x be given. Suppose, toward a contradiction, that the agenda setter has a strictly profitable deviation from proposing some $y \neq \beta(x)$. By the argument in the preceding paragraph, (i) the continuation outcome is $\beta^2(x)$ if the agenda setter follows her strategy of proposing $\beta(x)$, and (ii) players anticipate that acceptance of y results in outcome $\beta^2(y)$, whereas its rejection leads to $\beta^2(x)$. By the supposition and properties (i)–(ii), it must be that some voter—call her i —is strictly worse off under $\beta^2(y)$ than under $\beta^2(x)$. Property (ii) then implies that voter i must vote against this proposal. Moreover, by the construction of β , both continuation outcomes result in at least two voters obtaining zero utility. Although the identities of these two voters may differ across these outcomes, as there are three voters in total, the pigeonhole principle implies that at least one voter—call her j —obtains a zero share in *both* $\beta^2(y)$ and $\beta^2(x)$. Our capricious tiebreaking rule and property (ii) then stipulate that voter j rejects the proposal y . Hence, we have found two distinct voters, i and j , who both vote against proposal y , implying that it does not pass. Therefore, the agenda setter does not find this deviation to be strictly profitable, leading to the desired contradiction. An analogous argument also applies for all rounds $t \leq T - 2$.⁴⁸

REFERENCES

- Ali, S. Nageeb, B. Douglas Bernheim, and Xiaochen Fan. 2019. "Predictability and Power in Legislative Bargaining." *Review of Economic Studies* 86 (2): 500–525.
- Anesi, Vincent, and Daniel J. Seidmann. 2014. "Bargaining over an Endogenous Agenda." *Theoretical Economics* 9 (2): 445–82.
- Apesteguía, Jose, Miguel A. Ballester, and Yusufcan Masatlioglu. 2014. "A Foundation for Strategic Agenda Voting." *Games and Economic Behavior* 87: 91–99.
- Austen-Smith, David, and Jeffrey S. Banks. 2002. *Positive Political Theory I: Collective Preference*. Ann Arbor, MI: University of Michigan Press.
- Austen-Smith, David, and Jeffrey S. Banks. 2005. *Positive Political Theory II: Strategy and Structure*. Ann Arbor, MI: University of Michigan Press.
- Banks, Jeffrey S., and John Duggan. 2006. "A General Bargaining Model of Legislative Policy-Making." *Quarterly Journal of Political Science* 1 (1): 49–85.
- Barberà, Salvador, and Anke Gerber. 2017. "Sequential Voting and Agenda Manipulation." *Theoretical Economics* 12 (1): 211–47.
- Baron, David P. 1991. "Majoritarian Incentives, Pork Barrel Programs, and Procedural Control." *American Journal of Political Science* 35 (1): 57–90.

⁴⁶This pair consists of voters 1 and 2 for the initial default in Example 3 but may be different for other initial defaults.

⁴⁷From the non-capricious equilibrium construction above, it is easy to see that the continuation strategies in any round- T or round- $(T - 1)$ subgame constitute an equilibrium therein. Moreover, voters have no profitable deviations in any round by construction.

⁴⁸The reader may wonder why the agenda setter cannot induce voters to break ties in favor of her proposal in rounds $t \leq T - 2$ by offering $\epsilon > 0$ more to each voter, as in models of legislative bargaining based on the closed-rule procedure. The issue is that here, unlike closed-rule bargaining, such policies are left open to further reconsideration; voters (correctly) anticipate that bargaining in rounds $T - 1$ and T leave at most one voter with a nonzero surplus.

- Baron, David P., and John A. Ferejohn.** 1989. "Bargaining in Legislatures." *American Political Science Review* 83 (4): 1181–1206.
- Battaglini, Marco, and Stephen Coate.** 2008. "A Dynamic Theory of Public Spending, Taxation, and Debt." *American Economic Review* 98 (1): 201–236.
- Bernheim, B. Douglas, Antonio Rangel, and Luis Rayo.** 2006. "The Power of the Last Word in Legislative Policy Making." *Econometrica* 74 (5): 1161–90.
- Berry, Christopher R., and Anthony Fowler.** 2016. "Cardinals or Clerics? Congressional Committees and the Distribution of Pork." *American Journal of Political Science* 60 (3): 692–708.
- Berry, Christopher R., and Anthony Fowler.** 2018. "Congressional Committees, Legislative Influence, and the Hegemony of Chairs." *Journal of Public Economics* 158: 1–11.
- Black, Duncan.** 1958. *The Theory of Committees and Elections*. New York, NY: Springer.
- Diermeier, Daniel, and Pohan Fong.** 2011. "Legislative Bargaining with Reconsideration." *Quarterly Journal of Economics* 126 (2): 947–85.
- Diermeier, Daniel, and Pohan Fong.** 2012. "Characterization of the von Neumann–Morgenstern Stable Set in a Non-Cooperative Model of Dynamic Policy-Making with a Persistent Agenda Setter." *Games and Economic Behavior* 76 (1): 349–53.
- Duggan, John, and Zizhen Ma.** 2023. "Extreme Agenda Setting Power in Dynamic Bargaining Games." Unpublished.
- Fahrenthold, David A.** 2012. "In Congress, sunset clauses are commonly passed but rarely followed through." *Washington Post*, December 15. https://www.washingtonpost.com/politics/in-congress-sunset-clauses-are-commonly-passed-but-rarely-followed-through/2012/12/15/9d8e3ee0-43b5-11e2-8e70-e1993528222d_story.html.
- Farquharson, Robin.** 1969. *Theory of Voting*. New Haven, CT: Yale University Press.
- Ferejohn, John A., Morris Fiorina, and Richard D. McKelvey.** 1987. "Sophisticated Voting and Agenda Independence in the Distributive Politics Setting." *American Journal of Political Science* 31 (1): 169–93.
- Fournaies, Alexander.** 2018. "When Are Agenda Setters Valuable?" *American Journal of Political Science* 62 (1): 176–91.
- Kalandrakis, Tasos.** 2004. "A Three-Player Dynamic Majoritarian Bargaining Game." *Journal of Economic Theory* 116 (2): 294–322.
- Lizzeri, Alessandro, and Nicola Persico.** 2001. "The Provision of Public Goods under Alternative Electoral Incentives." *American Economic Review* 91 (1): 225–39.
- Maskin, Eric, and Jean Tirole.** 2019. "Pandering and Pork-Barrel Politics." *Journal of Public Economics* 176: 79–93.
- McKelvey, Richard D.** 1976. "Intransitivities in Multidimensional Voting Models and Some Implications for Agenda Control." *Journal of Economic Theory* 12 (3): 472–82.
- McKelvey, Richard D.** 1979. "General Conditions for Global Intransitivities in Formal Voting Models." *Econometrica* 47 (5): 1085–1112.
- McKelvey, Richard D.** 1986. "Covering, Dominance, and Institution-Free Properties of Social Choice." *American Journal of Political Science* 30 (2): 283–314.
- McKelvey, Richard D., and Richard G. Niemi.** 1978. "A Multistage Game Representation of Sophisticated Voting for Binary Procedures." *Journal of Economic Theory* 18 (1): 1–22.
- Miller, Nicholas R.** 1977. "Graph-Theoretical Approaches to the Theory of Voting." *American Journal of Political Science* 21 (4): 769–803.
- Shepsle, Kenneth A., and Barry R. Weingast.** 1984. "Uncovered Sets and Sophisticated Voting Outcomes with Implications for Agenda Institutions." *American Journal of Political Science* 28 (1): 49–74.