No Way Out:
Dual Channels of Manipulation in Agenda Institutions

Vincent Anesi[*] Tianxiong Yu[†]

March 11, 2026

Abstract

A large body of literature in political economy emphasizes the importance of limiting opportunities for manipulation of legislative institutions by self-interested actors. We show that the very conditions that shield institutions from agenda manipulation are precisely those that expose them to capture by special interests. This result holds in a highly general dynamic framework that encompasses a broad range of empirically relevant agenda institutions and policy-making environments, including those with policy uncertainty and experimentation.

[*]University of Luxembourg, *Email*: `vincent.anesi@uni.lu`
[†]University of Luxembourg, *Email*: `tianxiong.yu@uni.lu`

1 Introduction

This paper revisits a foundational question in political economy: Under what conditions can legislative policymaking institutions and their outcomes be insulated from manipulation by self-interested political actors? A classical strand of the literature has examined how policy choices within legislatures, and more generally committees, can be vulnerable to the influence of those individuals who control the agenda-setting process. The central lessons are now familiar. First, strategic agenda-setting can generally steer group decisions toward preferred alternatives,[1] undermining the idea of a coherent "will of the people" (Riker, 1982). Second, under conditions that arise rather naturally in political contexts, a unique policy outcome, the Condorcet winner, emerges and prevails regardless of the agenda procedure.[2] Yet agenda-setters are not the only actors capable of influencing the outcomes of agenda institutions. An equally influential literature shows that special-interest groups play a central role in shaping legislative outcomes.[3] This body of literature has evolved largely independently from that on agenda manipulation. The aim of this paper is to uncover a relationship between these two forms of manipulability: specifically, that the very conditions that shield institutions from agenda manipulation are precisely those that expose them to capture by special interests.

To establish the generality of this result, we analyze a broad class of policy-making environments that capture many situations of interest. In particular, this class captures the following empirically relevant features:

(i) Agenda institutions: As is standard in the literature on legislative agenda manipulation, committee members in our framework select policies by voting over sequential binary agendas. Common examples of such procedures — widely used in legislative bodies as well as in various other contexts (e.g., academic committees, international negotiations, boardrooms) — include amendment agendas, successive-elimination agendas, and issue-by-issue voting (Austen-Smith and Banks, 2005). Our framework encompasses the entire family of

[1] E.g., Farquharson (1969), Plott and Levine (1978), Shepsle and Weingast (1984), Ordeshook and Schwartz (1987), Miller (1995), and Barberà and Gerber (2017).

[2] E.g., Black (1958), Ferejohn and Grether (1974), Grandmont (1978), and Gans and Smart (1996).

[3] E.g., Denzau and Munger (1986), Snyder (1991), Groseclose and Snyder (1996), Grossman and Helpman (2001).

these binary agendas.[4]

(ii) Policy uncertainty and experimentation: Uncertainty over the consequences of policy choices is a pervasive feature of real-world decision-making. In legislative settings, such uncertainty typically arises from incomplete information about policy outcomes and the future political environment. Its empirical relevance is well established (e.g., Baker et al., 2022), and it has been a central focus in the formal analysis of legislative institutions for decades (Gilligan and Krehbiel, 1987). In the presence of such uncertainty, policy decisions often serve as experiments: by observing the outcomes of their choices, policymakers learn about the impacts and desirability of the policies they implement.

Our framework captures this dynamic dimension of learning. Committee members may be uncertain about the effects of different policies, form initial assessments of their potential benefits, and update these assessments over time by experimenting with policies and observing the resulting outcomes. The model imposes virtually no restriction on committee members' learning processes, thus accommodating a wide range of rationality levels — from fully Bayesian updating to complete naiveté. Moreover, it places only minimal restrictions on how revised assessments translate into new policy preferences. These preferences may reflect short- or long-term incentives, strategic considerations, behavioral traits, and more.

Our framework thus complements a broad and growing literature on policy experimentation, which has so far focused on institutional settings other than agenda voting — such as voting in two-armed bandit settings (Strulovici, 2010; Freer et al., 2020), accountability (Bils and Izzo, 2024), electoral competition (e.g., Callander, 2011; Hwang, 2023; Izzo, 2023), multilateral bargaining (Anesi and Bowen, 2021), and organizations with endogenous memberships (Gieczewski and Kosterina, 2024), to name a few.

Although the notion of agenda independence is typically applied to static contexts in the existing literature (e.g., Ferejohn et al., 1987), it extends naturally to our dynamic

[4]Note that this use of the term "agenda" differs from that found in another strand of literature, where it refers to the policy dimensions being addressed (e.g., Chen and Eraslan, 2017; and Barberà and Gerber, 2022).

environments: A policy-making environment is *agenda-independent* if no change in the agenda can alter the path of policy choices; otherwise, it is *agenda-manipulable.*

Existing work has identified conditions under which environments can be agenda-independent. Our manipulation theorem shows however that under those very conditions, the environment becomes *prone to special-interest capture* in the following strong sense: for *any* given policy, a special-interest group advocating for it can increase the number of times it is implemented up to *any* given period by influencing the committee's decision in only *some* arbitrary periods, even just one. Thus, agenda-independence guarantees that for any policy, a resourceful group with strong preferences for that policy will always find it beneficial to intervene and capture the committee's decision-making in any given period, regardless of how forward-looking the group may be. As we show, this is not always the case when the environment is instead agenda-manipulable: under policy uncertainty, pushing a policy today can change the committee's learning process in ways that make the policy significantly *less* likely to be selected again in the future.

We emphasize that our criterion for proneness to capture is deliberately stringent, building in a bias *against* capture. One can easily imagine that far less would suffice to motivate a special-interest group to undertake capture activities. This makes however our manipulation theorem as strong as possible: despite the demanding nature of the criterion, every policy-making environment is vulnerable to manipulation by self-interested political actors, whether agenda-setters, special interests, or both.

This finding resonates with a large body of political science research that emphasizes the importance of limiting opportunities for manipulation in order to protect democratic values (e.g., Riker, 1986), ensure policy quality and equity (e.g., Stigler, 1971; Ferejohn et al., 1987), and sustain effective, credible governance (e.g., Shepsle and Weingast, 1987). Given the generality of the environments to which our framework applies, our result suggests that many existing institutions are unlikely to fully achieve these goals.

2 Framework

2.1 Policy-Making Environments

General description. We present a general model of dynamic policymaking by a committee (such as a legislature or an academic committee) consisting of $n \geq 2$ members, denoted by $N \equiv \{1, \ldots, n\}$. In each of an infinite number of discrete periods, $t = 1, 2, \ldots$, the committee selects a policy x^t from a finite set $X \equiv \{x_1, \ldots, x_L\}$, where $L \geq 2$. The potential utility benefits that committee members expect to derive from policies depend on the value of a payoff-relevant state of the world $\omega = (\omega_1, \ldots, \omega_L)$. Each component ω_k of the state may itself be multidimensional, so as to capture for example agent-specific dimensions of policy outcomes. Although the value of the vector ω is initially unknown, it is common knowledge that each of its components ω_k belongs to some finite set Ω_k. Let $\Omega \equiv \Omega_1 \times \cdots \times \Omega_L$ denote the set of possible states.

Despite their uncertainty about the state of the world, committee members can form assessments of the potential quality of each policy, and revise those assessments over time by experimenting with different policies and observing the resulting outcomes. We begin by informally describing this process of learning and policy making, before formally detailing each of its steps below — Figure 1 provides a graphical illustration.

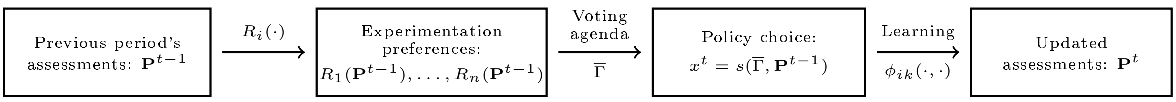

Figure 1: Policy making and learning in period t.

Each member begins every period t with her own "assessment" of each policy in X. This assessment can for example be thought of as an index of the policy's expected performance, quality, or desirability. It determines their "experimentation preferences" over X, i.e., their preferences regarding which policy should be implemented in period t. As we will see below, the model's assumptions concerning these preferences are general enough to capture a wide variety of incentives — e.g., static or dynamic, sincere or strategic, etc.

The preference profile of the n committee members is then aggregated into a collective policy choice x^t using a voting procedure, which will be modeled as a binary voting agenda. The implementation of this policy x^t is then followed by the realization of a policy outcome, the observation of which leads members to update their assessments of policy x^t. Holding these updated assessments, the committee transitions to period $t+1$, in which the same process is repeated.

Policy assessments. Formally, the assessment of any policy x_k by committee member i is represented by a real number p_{ik}. For example, it could correspond to the Gittins index, as in standard experimentation problems in the literature. We assume that the set of possible assessments P that any member can hold about any policy is countable.

Given an assessment p_{ik} for each member $i \in N$ and each policy $x_k \in X$, we obtain the $(n \times L)$ *assessment matrix* $\mathbf{P} \equiv [p_{ik}]$, whose ith row $p_{i\cdot} = (p_{i1}, \ldots, p_{iL})$ describes member i's assessments of each policy in X, and whose kth column $p_{\cdot k} = (p_{1k}, \ldots, p_{nk})$ describes the assessments of all n members regarding policy x_k. The (countable) set of all possible assessment matrices is denoted by $\mathcal{P}$.

As described in Figure 1, the committee begins each period t with an assessment matrix $\mathbf{P}^{t-1}$, inherited from policy experiments conducted in earlier periods, which is then revised into a new assessment matrix $\mathbf{P}^t$, based on the policy experiment carried out in period t. We elaborate on this updating process below. The initial assessment matrix $\mathbf{P}^0$ is exogenously chosen by Nature according to some arbitrary probability distribution with full support on $\mathcal{P}$.

Experimentation preferences. Every assessment matrix $\mathbf{P}$ gives rise to preferences, for each committee member $i \in N$, over which policy they would like to see implemented in the current period. To avoid imposing unnecessary restrictions on these preferences, we simply assume that they are represented by a function R_i which, for each $i \in N$, assigns to every $\mathbf{P}$ in $\mathcal{P}$ a binary relation $R_i(\mathbf{P})$ on X. We only make two mild assumptions about R_i. First, for every $\mathbf{P} \in \mathcal{P}$, the relation $R_i(\mathbf{P})$ is complete and asymmetric. This rules out indifference over any pair of alternatives in X which, given the finiteness X, would be nongeneric. Second, R_i satisfies the following independence-of-irrelevant-alternatives condition: for any pair of alternatives x_k and x_ℓ in X, and assessment matrices $\mathbf{P}$ and $\mathbf{Q}$ such that $p_{\cdot k} = q_{\cdot k}$ and $p_{\cdot \ell} = q_{\cdot \ell}$, we have $x_k R_i(\mathbf{P}) x_\ell$ if and only if $x_k R_i(\mathbf{Q}) x_\ell$.[5]

This formulation allows for a wide variety of situations of interest. The policy experimentation incentives described by the function R_i may reflect short-term or long-term considerations, exploration-exploitation tradeoffs, behavioral traits, etc. In particular, note that the preference relation of member i may depend not only on her own assessments but also on the assessments of other members. For instance, this can reflect strategic anticipations by member i about the future choices of other committee members in the equilibrium of a game, as illustrated in Example 1 below.

We also note that the functions R_i exhibit a form of Markovian stationarity, since they are time-independent. This assumption is made purely for expositional ease; none of our results would be affected if these functions were instead of the form $R_i^t(\cdot)$.

Voting agendas and outcomes. Once endowed with a profile of experimentation preferences $\big(R_1(\mathbf{P}), \ldots, R_n(\mathbf{P})\big)$, the committee must aggregate these preferences into a collective choice of a policy from X. Following most of the formal literature on the manipulation of legislative institutions, we focus on a family of empirically important institutions; namely, those characterized by majoritarian committee voting over binary agendas. The most commonly encountered examples are amendment agendas, successive elimination agendas, and issue-by-issue voting, in which voting proceeds sequentially, with each decision in the sequence involving a pair of (possibly composite) alternatives (e.g., Austen-Smith and Banks, 2005). Such procedures are typically observed in legislative settings, but they are also commonly used in academic committees, international organizations, judicial panels, and corporate boards.

Although the concept of a *(binary) agenda* (or *voting tree*) is relatively simple and intuitive, the formalism required for a precise definition can be somewhat cumbersome.

[5]This is especially appropriate in settings, which we focus on here, where the committee's alternatives correspond to distinct policy instruments or directions. It would not apply if the alternatives were different values of the same instrument (e.g., a tax rate or tariff), since trying one option would then reveal information about the others (e.g., as in Callander, 2011).

To avoid introducing additional notation that will not be used later, we provide here an informal definition, relegating the formal mathematical machinery to the Appendix.[6] An agenda is modeled as a binary decision tree that specifies the order in which pairwise votes are conducted. Each internal node in the tree corresponds to a decision between two alternatives (or intermediate winners of earlier votes), and each such decision leads to a new branch depending on the majority choice. The terminal nodes of the tree are labeled with the policy alternatives in X, while the internal nodes (called "decision nodes") represent the sequence of votes that take place. Every decision node has exactly two immediate successors, ensuring that each stage of voting involves a binary comparison. This tree structure captures the procedural constraints imposed by the institution: which alternatives are compared, and in what order. The family of all binary agendas is denoted by $\mathcal{A}$.

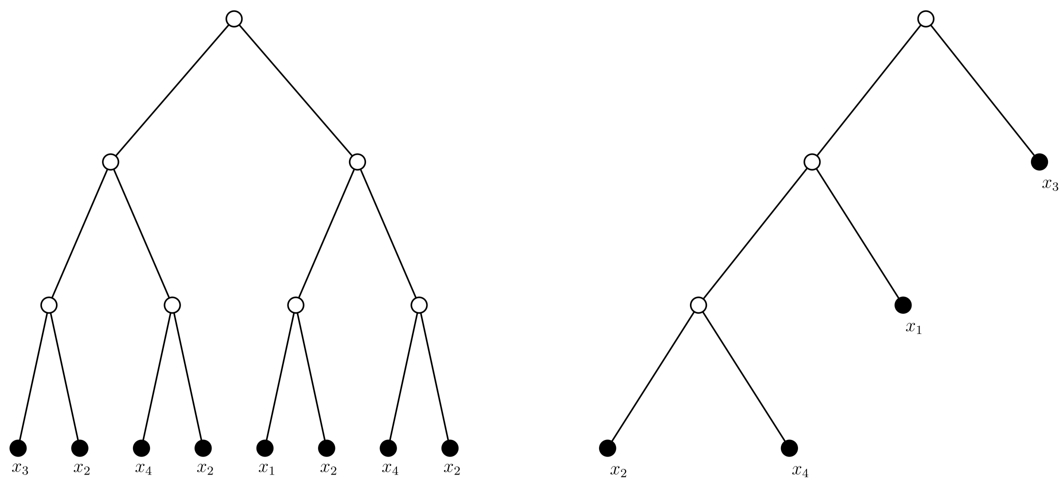

Figure 2: An amendment agenda (left) and a successive elimination agenda (right).

Figure 2 presents two standard examples of agendas in $\mathcal{A}$, both with voting sequence $(x_3, x_1,\ x_4, x_2)$: an amendment agenda and a successive-elimination agenda. The former describes a sequential decision-making process where one policy (or "proposal") can be amended step by step until a final decision is made: here, the initial proposal x_3 is first

[6]Formal treatments can also be found in standard textbooks (e.g., Austen-Smith and Banks, 2005) and surveys (e.g., Miller, 1995). The seminal reference is Farquharson (1969).

pitted against x_1; the winner of this pairwise vote is then pitted against x_4; the winner of this second pairwise vote is pitted against x_2; and the winner of this final vote is the committee's final choice. In the successive-elimination agenda shown in Figure 2, the first proposal x_3 is offered as the final choice in a yes-no vote; if it is accepted, it becomes the committee's final choice; if it is rejected, the second policy in the sequence, x_1, is considered; and so on.

As in most of the existing literature, we assume in the main text that n is odd and that each pairwise vote is decided by simple majority voting. This ensures that the committee's social preference relation is a tournament (i.e., it is both complete and asymmetric). This assumption greatly simplifies the exposition, guaranteeing the existence of a unique policy outcome for any agenda and any preference profile, without requiring specific assumptions about how ties between pairs of policies are broken. However, as we elaborate in Section 4, our results also apply to even-sized committees and/or non-majoritarian institutions.

The final outcome of an agenda depends on the order in which votes occur and the committee members' preferences on X. To determine the collective choice from a given agenda $\Gamma \in \mathcal{A}$ and preference profile $\big(R_1(\mathbf{P}), \ldots, R_n(\mathbf{P})\big)$ induced by an assessment matrix $\mathbf{P}$, we use the standard concept of sophisticated voting (Shepsle and Weingast, 1984) — defined formally in Section A of the Appendix. This involves backward induction through the voting tree: at each decision node, committee members anticipate the outcomes of future votes and vote so as to induce their most preferred policy as the final choice. By recursively applying majority rule at each decision node — starting from the bottom of the tree and working upwards — we can determine the unique policy that emerges from the full voting process. This policy is referred to as the *sophisticated outcome* of Γ,[7] and is denoted by $s(\Gamma, \mathbf{P})$.

Returning to the dynamic policy-experimentation framework described in Figure 1, we assume for (and only for) expositional ease that the committee uses the same agenda, denoted by $\overline{\Gamma}$ in every period. (None of our results would be affected if we instead allowed

[7]It is well known that this is the unique policy outcome obtained from the iterated elimination of dominated strategies in the corresponding noncooperative game form (Moulin, 1979).

for a time-dependent sequence $\{\overline{\Gamma}^t\}$ of agendas in $\mathcal{A}$.) Coupled with the assessment matrix $\mathbf{P}^{t-1}$ inherited from the previous period, the fixed agenda $\overline{\Gamma}$ generates a policy choice $x^t = s(\overline{\Gamma}, \mathbf{P}^{t-1}) \in X$ in each period t. This is the policy chosen by the committee in period t.

Learning. To complete the description of the sequence of events in period t, it remains to describe how committee members' assessments evolve during that period. After implementing a policy, say x_k, members observe a policy outcome that may lead them to revise their assessment of the quality of x_k. This policy outcome may stochastically depend on the state of nature, and more specifically, on its kth component ω_k. Each member i is therefore led to revise her beliefs about ω_k, and consequently her assessment of x_k, according to an individual updating process, which may be Bayesian but could also reflect forms of bounded rationality. We only assume that each member i's updating process is governed by a stochastic transition function $\phi_{ik}\colon P \times \Omega_k \to \Delta(P)$. Thus, if member i holds assessment $p_{ik}^{t-1} \in P$ at the beginning of period t, then conditional on the true state regarding x_k being $\omega_k \in \Omega_k$, her updated assessment p_{ik}^t at the end of period t is drawn according to the probability distribution $\phi_{ik}(p_{ik}^{t-1}, \omega_k)$ on P.[8] As for policies $x_\ell \neq x_k$ that have not been selected in this period, their assessments remain unchanged (i.e., $p_{\cdot\ell}^t = p_{\cdot\ell}^{t-1}$), since members acquire no new information about them.

As we have noted repeatedly, the framework described above is flexible. Depending on the policy instruments, committee members' preferences and incentives, their levels of rationality, the agenda institutions, or the informational structures that one seeks to capture, the parametric configuration of the model can be adjusted to account for a broad range of situations of interest. In what follows, we refer to any such parametric configuration as a *policy-making environment.*

Example 1: Two-armed Bandit Environments. Before proceeding further, it is useful to illustrate how the general framework above can accommodate two-armed bandit

[8] Our results would continue to hold if the transition of committee member i's assessment depended not only on their own assessment of policy x_k, but also on the assessments of the other members; that is, if the function $\phi_{ik}(\cdot, \omega_k)$ depended on the entire column $p_{\cdot k}^{t-1}$ of $\mathbf{P}^{t-1}$.

models, which are commonly used in the literature on policy experimentation. This example illustrates, in particular, how the committee members' experimentation preferences can result from their equilibrium continuation values in a dynamic game. Consider an environment in which there are only two feasible policies: a "safe policy" x_1 and a "risky reform" x_2. Suppose that $\Omega_1 = \{\text{Safe}\}$ and $\Omega_2 = \{\text{Good}, \text{Bad}\}$. There are fixed parameters $\bar{p}, \lambda \in (0,1)$ such that the set of possible assessment matrices $\mathcal{P}$ consists of all $(n \times 2)$ matrices of the form

$$\mathbf{P} = \begin{bmatrix} 1 & p \\ 1 & p \\ \vdots & \vdots \\ 1 & p \end{bmatrix} ,$$

for some

$$p \in P \equiv \left\{ \frac{\bar{p}(1-\lambda)^l}{\bar{p}(1-\lambda)^l + 1 - \bar{p}} : l = 0, 1, \ldots \right\} \cup \{1\} .$$

Each committee member $i \in N$ has experimentation preferences over $\{x_1, x_2\}$ determined by a fixed cutoff $q_i \in (0, \bar{p}) \setminus P$, such that for any $\mathbf{P} \in \mathcal{P}$, $x_1 R_i(\mathbf{P}) x_2$ if and only if $p < q_i$. Assume that $q_1 < \cdots < q_n$; so that letting $m \equiv (n+1)/2$, q_m is the median of the q_i's. Accordingly, the risky reform is majority-preferred to the safe policy if and only if $p > q_m$. As there are only two feasible policies, there is a unique agenda $\overline{\Gamma}$, whose sophisticated outcome is

$$s(\overline{\Gamma}, \mathbf{P}) = \begin{cases} x_1 & \text{if } p < q_m , \\ x_2 & \text{otherwise,} \end{cases}$$

for all $\mathbf{P} \in \mathcal{P}$.

The transition functions governing the learning process are identical for each member i. As there is only one possible assessment of the safe policy, the probability distribution $\phi_{i1}(1, \text{Safe})$ must assign probability one to the value 1. The transition function for the risky policy is defined by the following equalities: $\phi_{i2}(p, \text{Good})(1) = \lambda$,

$$\phi_{i2}(p, \text{Good}) \left(\frac{p(1-\lambda)}{1-p\lambda} \right) = 1 - \lambda \;\text{ and }\; \phi_{i2}(p, \text{Bad}) \left(\frac{p(1-\lambda)}{1-p\lambda} \right) = 1 ,$$

for all $p \in P$.

The environment described above captures policy experimentation in the unique Markov perfect equilibrium of the model developed by Anesi and Bowen (2021). In a collective-choice variant on the standard two-armed bandit framework, a committee repeatedly chooses between a safe policy x_1 with known, constant payoffs and a risky reform x_2 whose payoffs depend on an unknown, fixed state, which is either "good" or "bad." Committee members hold a common prior belief about the state of the reform and update their belief p over time through Bayesian learning, based solely on the observed outcomes of the implemented policy. The risky reform yields stochastic payoffs. Specifically, when the risky reform is implemented, no payoff is received unless a signal arrives. This signal arrives according to a state-dependent Poisson process: in each period, a "good news" signal arrives with probability $\lambda > 0$ only in the good state, while no signals are ever received in the bad state. Thus, the absence of signals is itself informative, gradually leading members to revise their belief downward when no good news occurs; and the first arrival of good news reveals to all members that the reform is good, causing their belief about the state to update to one.

The reform, if good, benefits all members. This is why the set of relevant states for x_2 is simply $\Omega_2 = \{\text{Good}, \text{Bad}\}$.[9] However, they differ in how much they value a good reform relative to the safe policy, with member m holding the median valuation. Each member seeks to maximize expected discounted payoffs over an infinite horizon. Anesi and Bowen (2021) show that the dynamic voting game, in which collective choice is determined by simple majority rule in each period, has a unique Markov perfect equilibrium. At the beginning of each period, committee members hold a common prior p^{t-1}, updated through past policy experimentation. This belief determines each member i's equilibrium continuation values for choosing the risky reform versus the safe policy. These continuation values, in turn, shape her experimentation preferences in that period: $x_1 R_i(\mathbf{P}^{t-1})\, x_2$ if and

[9] This stands in contrast to Strulovici (2010), where committee members are either "winners" for whom the reform is beneficial, or "losers" for whom it is detrimental; but all winners receive the same expected payoff from the reform. This framework can also be readily accommodated within ours. In this case, the state space would be given by $\Omega_2 = \{\text{Good}, \text{Bad}\}^n$, where the ith component of $\omega_2 \in \Omega_2$ would indicate whether member i is a winner ($\omega_{2i} = \text{Good}$) or a loser ($\omega_{2i} = \text{Bad}$).

only if the continuation value of implementing x_1 exceeds that of implementing x_2. In fact, the authors show that there exists an equilibrium cutoff q_i such that member i prefers to implement the risky reform if and only if $p^{t-1} > q_i$. Logically, members who place a higher value on a good reform have a lower cutoff. It then follows from majority rule that the risky reform is adopted if and only if $p^{t-1} > q_m$, as prescribed by $s(\overline{\Gamma}, \mathbf{P}^{t-1})$ above. $\square$

2.2 Manipulation

Having defined a broad class of policy-making environments under preference uncertainty, we now turn to the conceptualization of manipulability in such environments.

Agenda manipulation. Our formalization of agenda manipulation and independence follows the seminal papers on agendas and the control of policy outcomes. As in that literature, we say that the outcome of a policy-making environment is "agenda-independent" if it is impossible for an agenda-setter to manipulate the committee's decision by appropriately selecting the agenda in any period; more specifically, if it is impossible to alter the (stochastic) sequence of policies $x^t = s(\overline{\Gamma}, \mathbf{P}^{t-1})$ by modifying the default agenda $\overline{\Gamma}$ into another agenda from $\mathcal{A}$; and that it is "agenda-manipulable" otherwise. The following formal definition is a natural extension of that of Ferejohn et al. (1987) to dynamic settings with uncertainty.

Definition 1. *The outcome of a policy-making environment is* agenda-independent *if* $s(\Gamma, \mathbf{P}) = s(\overline{\Gamma}, \mathbf{P})$*, for all* $\mathbf{P} \in \mathcal{P}$ *and all* $\Gamma \in \mathcal{A}$*. Otherwise, it is* agenda-manipulable.

This definition implicitly assumes that agenda-setters are unbound by any institutional constraints in their choices of agendas. They can freely choose any agenda from $\mathcal{A}$. It is common in the literature on agenda manipulation to restrict agenda choices to the empirically relevant amendment agendas — e.g., Miller (1980) and Banks (1985). As we explain in Section 4, such a restriction would have no impact on our results.

Special-interest capture. The manipulation of policy outcomes through capture clearly presumes a special-interest group possessing both the necessary resources and sufficiently strong policy preferences to be willing to bear the costs of such activities. In a static environment, where the impact of policies is known and there is no learning, there is no reason why such a group would refrain from increasing the adoption probability of its preferred policy. As the example in Subsection 3.3 illustrates, however, the logic becomes more subtle in environments with policy experimentation. Shifting the committee's decision toward a policy x_k in period t may affect the evolution of beliefs in such a way that the committee becomes significantly *less* likely to select it in subsequent periods.[10] In that case, the group may well be deterred from manipulating the committee's decision in period t.

To ensure our results are as strong as possible, we adopt a deliberately stringent criterion: we will say that the outcome of a policy-making environment is "prone to special-interest capture" if this is impossible; that is, if increasing the likelihood that the committee selects x_k in any given period increases the frequency with which it is selected in the future. We stress that in adopting this criterion, we are not claiming that a policy-making environment failing to meet it cannot reasonably be viewed as vulnerable to capture. As noted in the introduction, far less may suffice to induce a group to undertake capture efforts. A very shortsighted group, for instance, would gladly exploit any chance to influence today's decision, regardless of the future consequences. Our aim, however, is to establish the *inevitability* of manipulation. Biasing our notion of proneness to capture against obtaining capture only strengthens our point: whenever the environment is agenda-independent, a group always has as strong an incentive as possible to manipulate the committee's decisions, no matter how much weight it places on future policy choices.

To formalize the ideas above, we begin by establishing some notation. Define an *outcome function* as a mapping $\sigma\colon \mathbb{N} \times \mathcal{P} \to \Delta(X)$ that assigns to every period $t \in \mathbb{N}$ and every assessment $\mathbf{P} \in \mathcal{P}$ held by committee members, a probability $\sigma(t, \mathbf{P})(x_k)$ that each policy x_k is selected in t. In a given policy-making environment with agenda $\overline{\Gamma}$, the *default*

[10] A remarkable result by Gossner et al. (2021) shows that such adverse impacts of choice manipulation are in fact impossible in *individual* decision-making settings under a mild independence-of-irrelevant-alternatives condition. However, it is easy to verify that this result does not generally extend to our collective decision-making settings: even if each voter's preferences satisfy this condition, that does *not* imply that the committee's sophisticated voting outcome, $s(\Gamma, \cdot)$, satisfies it as well.

outcome function $\bar{\sigma}$ is the one characterizing the committee's decisions in the absence of special-interest capture; formally, for all $x_\ell \in X$, $\bar{\sigma}(t, \mathbf{P})(x_\ell) = 1$ if and only if $s(\overline{\Gamma}, \mathbf{P}) = x_\ell$.

We conceptualize the capture of the policy-making environment by a special-interest group advocating for a given policy x_k as a change from the default outcome function $\bar{\sigma}$ to an alternative outcome function $\hat{\sigma}$ such that: (i) in every period, the committee is (weakly) more likely to choose policy x_k for any given committee members' assessments; while conditional on the policy choice not being x_k, the outcome is the same under both functions; and (ii) there exists at least one period in which the committee is, in expectation, *strictly* more likely to select x_k.[11]

Formally, any outcome function σ, together with the transitions ϕ_{ik}, induces a stochastic sequence of assessment matrices $\{\mathbf{P}^t(\sigma)\}$ held by committee members over time. For each $k = 1, \ldots, L$, we then say that $\hat{\sigma}\colon \mathbb{N} \times \mathcal{P} \to \Delta(X)$ is a *k-captured outcome function* if the following holds:

(i) for all $t \in \mathbb{N}$ and $\mathbf{P} \in \mathcal{P}$: (i.a) $\bar{\sigma}(t, \mathbf{P})(x_k) \leq \hat{\sigma}(t, \mathbf{P})(x_k)$; and (i.b) $s(\overline{\Gamma}, \mathbf{P}) \neq x_k$ implies $\hat{\sigma}(t, \mathbf{P})(x_\ell)/[1 - \hat{\sigma}(t, \mathbf{P})(x_k)] = \bar{\sigma}(t, \mathbf{P})(x_\ell)$, for each $\ell \neq k$;

(ii) for every $\omega \in \Omega$, there exists $t \in \mathbb{N}$ such that $\mathbb{E}\big[\bar{\sigma}\big(t, \mathbf{P}^{t-1}(\bar{\sigma})\big)(x_k)\big] < \mathbb{E}\big[\hat{\sigma}\big(t, \mathbf{P}^{t-1}(\hat{\sigma})\big)(x_k)\big]$.

Finally, to compare the frequencies with which the committee adopts a given policy x_k under the default and captured outcome functions, we denote by $\mathfrak{n}_k^t(\sigma \mid \omega)$ the (random) number of periods, up to period $t \in \mathbb{N}$, in which x_k is implemented in state ω under any outcome function σ. (The randomness of $\mathfrak{n}_k^t(\sigma \mid \omega)$ arises from the possible realizations of the sequence $\{\mathbf{P}^t(\sigma)\}$ and implemented policies x^t, both of which are induced by the transition functions $\phi_{ik}(\cdot, \omega_k)$ and the outcome function σ.)

It turns out that a manipulation result in the same vein as Proposition 1 extends to S-capture. For every nonempty set $S \subset X$ and every outcome function σ, let $\mathfrak{n}_S^t(\sigma \mid \omega)$ denote the (random) number of periods up to period $t \in \mathbb{N}$ in which an element of S is implemented in state ω under σ.

[11] A more demanding definition would require that there exists a period in which the committee is strictly more likely to select x_k, irrespective of the assessments. While this would not alter our results, it would in fact preclude the existence of a captured outcome in many cases of interest, thus limiting the scope of our findings. Adopting a weaker notion of capture strengthens our result.

Definition 2. *The outcome of a policy-making environment is* prone to special-interest capture *if the following holds for all* $\omega \in \Omega$, $x_k \in X$, *and* k*-captured outcome function* $\hat{\sigma}$*:* $\mathfrak{n}_k^t(\hat{\sigma} \mid \omega)$ *first-order stochastically dominates* $\mathfrak{n}_k^t(\bar{\sigma} \mid \omega)$ *for every period* $t \in \mathbb{N}$*, with strict dominance in at least one period.*

In words, the outcome of a policy-making environment is prone to special-interest capture if for *any* policy x_k, a special-interest group supporting x_k can (stochastically) increase the number of times it is implemented up to *any* given period by manipulating the committee's decision in only *some* arbitrary periods — possibly just one. Moreover, the definition requires that a *strict* increase in the implementation frequency of policy x_k occur after finitely many periods. These intentionally strong requirements imply that any benchmark situations (such as those induced by the default outcome functions $\bar{\sigma}$), in which committee decisions are immune to manipulation by agenda setters and all the existing special-interest groups, are hardly sustainable. The following remark elaborates on this implication.

Remark: Capture-proneness from a noncooperative perspective. The notion of *proneness to capture* introduced above is formulated axiomatically, allowing us to avoid relying on specific behavioural assumptions. While this level of generality is deliberate, the concept also has clear implications for game-theoretic environments. Indeed, at the cost of some generality, it is straightforward to construct explicit noncooperative foundations. A detailed illustration is provided in the appendix; here, we sketch the underlying logic through a heuristic argument.

As illustrated in Example 1 above, the sophisticated outcome $s(\overline{\Gamma}, \cdot)$—and thus the default outcome function $\bar{\sigma}$—can be interpreted as the voting outcome that would arise in a Markovian equilibrium of a dynamic game of collective experimentation among committee members *in the absence of capture.* Any such a game can be extended by introducing vote-buying special-interest groups, in the spirit of Groseclose and Snyder (1996). To fix ideas, suppose first that there is a single special-interest group that strictly prefers some policy, say x_k, and derives a large, positive payoff in any period in which that policy is implemented, and zero otherwise. Assume further that the group seeks to maximize its expected discounted average payoff. In this extended environment, $\bar{\sigma}$ can now be viewed as the outcome induced by a Markov strategy profile ς in which the special-interest group never attempts to buy votes and committee members best respond accordingly. The natural question is whether this "capture-proof" profile ς can actually constitute a (Markov perfect) equilibrium.

To answer this, one must ask whether the special-interest group can profitably deviate from ς. As the group's objective is to maximize its average discounted per-period payoff, the one-shot deviation principle applies. Hence, ς fails to be an equilibrium whenever the group can increase its discounted payoff by acquiring enough votes to enforce x_k in some period (taking into account all future consequences implied by the committee members' strategies under ς). The outcome induced by any such one-shot deviation must be a k-captured outcome function $\hat{\sigma}$.[12] When proneness to capture (Definition 2) holds, *every* such deviation is profitable, as it increases the frequency with which the group's ideal policy is implemented in the future.[13] This implies that the capture-proof strategy profiles like ς can never constitute an equilibrium, no matter which policy x_k the group favors, nor how farsighted it may be. Assuming equilibrium existence, we conclude that *all equilibria must entail capture by the group.*

The same logic extends to settings with multiple groups, each advocating a different policy. Regardless of the technology through which vote-buying efforts translate into success probabilities for each group in the game under consideration, the sustainability of a capture-proof equilibrium requires that no group be able to profitably deviate from a strategy profile in which none of them ever captures the committee's decisions. (Recall that the distribution of initial assessments has full support.) Given the restriction to Markovian equilibria, typical in the dynamic political-economy literature, any such deviation must

[12] Of course, not all k-captured outcome functions can be generated in this way; only those consistent with the committee members' strategies. In this respect, the axiomatic approach makes proneness to capture a more demanding requirement than in specific noncooperative settings, which in turn reinforces our manipulation result below.

[13] This argument implicitly assumes that in periods where the group's ideal policy is not implemented, the distribution over all other policies coincides with that under ς. The proof of Proposition 1 establishes that this is indeed the case whenever the policy environment is agenda-independent.

result in a k-captured outcome for some k. By the same reasoning as above, proneness to capture therefore implies that *all* groups have profitable one-shot deviations from any such profile, and hence that any Markovian equilibrium must exhibit some form of capture. $\square$

Combining the two notions of manipulability yields the following definition.

Definition 3. *A policy-making environment is* manipulable *if it is either agenda-manipulable or prone to special-interest capture.*

3 Inevitable Manipulation

3.1 Manipulation Theorem

Are there policy-making environments that are not susceptible to manipulation by either agenda-setters or special interests? Despite the broad family of environments considered here, the answer is no. As noted in the introduction, the political-economy literature has devoted considerable attention to the manipulation of agenda institutions, and has identified various conditions under which the outcomes of policy-making environments may be agenda-independent. Our main result, however, establishes that any environment that escapes agenda manipulation must exhibit the very strong form of proneness to special-interest capture, as specified in Definition 2.

Proposition 1. *If the outcome of a policy-making environment is agenda-independent, then it is prone to special-interest capture.*

The contrapositive of Proposition 1 is that any policy-making environment that is not prone to special-interest capture must necessarily be agenda-manipulable. We therefore conclude that, despite the wide variety of policy-making environments encompassed by our framework, none is immune to both channels of manipulation. We record this observation in the following corollary.

Corollary 1. *Every policy-making environment is manipulable.*

3.2 The Logic of Dual Manipulation: Heuristics

While the generality of Proposition 1 requires formal machinery and delicate arguments, its core logic may be illustrated by a simple thought experiment.

Consider a committee that, in each of three periods $t = 1, 2, 3$, must select one of $L > 2$ policies, $x_1, \ldots, x_L$. Members are uncertain about the potential impacts of these policies, but at the beginning of each period, they hold individual assessments of the associated benefits and costs. As above, these assessments determine their experimentation preferences, i.e., an ordering of the policies with respect to which they wish to experiment in the current period: implementing a policy yields additional information about its consequences, enabling members to revise their assessments of that policy, and hence to update their ordering in subsequent periods.

At this stage, we impose no restrictions on how individual members' assessments translate into their orderings of policies, except for the independence-of-irrelevant-alternatives condition. In each of the three periods, individual experimentation preferences are aggregated into a collective policy choice via a sequential majoritarian voting procedure (e.g., a standard amendment agenda).

Consider now the position of a special-interest group that seeks to maximize the implementation frequency with which its preferred policy, x_1, is implemented, and has the ability to capture the policy-making process in the first period. The group anticipates the following sequence of experimentation outcomes: (i) before any policy is tried, committee members' assessments are such that agenda voting would yield policy x_2; (ii) after one trial of x_2, individual assessments would be revised in a manner that leads to the selection of x_1; and (iii) after one trial of both x_1 and x_2, further updates in assessments would again result in the implementation of x_1. Accordingly, the group expects the policy sequence over the three periods to be (x_2, x_1, x_1).

Can the group be confident that enforcing its preferred policy x_1 in period 1 will not backfire in later periods? In general, the answer is no. Changing the period-1 policy from x_2 to x_1 means that it is then x_1, rather than x_2, whose assessment is revised at the beginning of period 2, as illustrated in Figure 3. Consequently, committee members enter period 2 with assessments different from those they would have held in the absence of capture. This, in turn, may alter their preferences in such a way that agenda voting yields a policy other than x_1. The resulting assessments at the start of period 3 would likewise differ, possibly leading again to a policy choice other than x_1. Special-interest capture can therefore have adverse consequences, reducing the frequency with which policy x_1 is implemented from two periods to only one.

		$t = 1$	$t = 2$	$t = 3$
Without capture	Assessments	Initial	After one trial of x_2 only	After one trial of each x_1 and x_2 only
	Policy choice	x_2	x_1	x_1
	Frequency of x_1 after t periods	0	1	2
Capture at $t = 1$	Assessments	Initial	After one trial of x_1 only	After one trial of each x_1 and x_k only
	Policy choice	x_1 (capture)	x_k ($\neq x_1$)	x_ℓ ($\neq x_1$)
	Frequency of x_1 after t periods	1	1	1

Figure 3: Illustration of how capture can alter the committee's learning trajectory in a way that ultimately harms the special-interest group. While the policy x_1 advocated by the group would be implemented twice after three periods in the absence of capture, it is implemented only once when the group intervenes in the first period.

Such adverse consequences of capture would however be precluded if after one trial of x_1 in period 1, the policy period 2 were necessarily either x_1 or x_2. Indeed, if x_1 were chosen again in period 2, the group would be guaranteed at least as many occurrences of x_1 as in the absence of capture, and possibly more if x_1 were also selected in period 3. If x_2 were chosen in period 2 then after two periods, each of x_1 and x_2 would have been tried once, exactly as in the absence of capture. Consequently, at the start of period 3, the assessments inherited from experimentation in the preceding periods would be identical with or without capture. This in turn would yield the same individual orderings of policies and, therefore, the same policy choice x_1.

How can the special-interest group be assured that, after one trial of x_1, the period-2 choice will be restricted to x_1 or x_2 after one trial of x_1? This is precisely where the notion of *agenda-independence* is critical. Suppose that voting is immune to agenda manipulation, in the sense that whatever the committee's policy assessments, it is impossible for an agenda-setter to change the outcome of the vote by modifying the agenda. A classical result in formal political theory establishes that such immunity obtains if and only if, for every realization of members' assessments during experimentation, their orderings are such that there exists a unique policy that defeats all others in pairwise majority comparisons. This policy, the *Condorcet winner,* is then the committee's collective choice.

Suppose that agenda-independence holds. Recall that in the absence of capture, it is policy x_2 that is chosen in period 1. Thus, before any policy is tried, the committee's initial assessments are such that x_2 is the Condorcet winner; by definition, it defeats every untried policy in a pairwise majority vote. Hence, even if the special-interest group enforces x_1 in the first period, the independence-of-irrelevant-alternatives condition ensures that x_2 would remain majority-preferred to any other untried policy at the start of period 2. As the voting outcome must coincide with the Condorcet winner under agenda-independence, the committee's choice in that period must then be either x_1 or x_2. This illustrates how the very condition preventing agenda-setters from manipulating voting outcomes ensures for the special-interest group that capturing the committee's policy choice in period 1, while immediately profitable, will also never prove disadvantageous in subsequent periods.

This simple heuristic assumes that the special-interest group can predict with certainty how committee members will learn from each policy trial, and thus the resulting policy trajectory. Of course, this is highly unrealistic, as policy experimentation typically involves substantial uncertainty about the possible outcomes of different trials. Proposition 1 establishes that the conclusions drawn from this illustrative example do not depend on this simplification or any of the other assumptions we have made here.

3.3 Agenda-Manipulability and Adverse Capture

To better understand the role of agenda-independence in fostering capture, we now illustrate why proneness to capture may fail to obtain if agenda-independence does *not* hold.

We do so by presenting a simple policy-making environment whose outcome is neither agenda-independent nor prone to special-interest capture. In particular, it illustrates that *without agenda-independence, attempts at capture can produce long-term adverse consequences for the special interests that seek to manipulate policy outcomes.*

A three-armed bandit environment. We now consider a committee of three members, $N = \{1, 2, 3\}$, who must choose a policy from the set $X = \{x_1, x_2, x_3\}$ in each period. Policy x_1 is "safe" ($\Omega_1 = \{\text{Safe}\}$) in the sense that it yields a certain positive payoff $u_{i1} > 0$ to each committee member i in every period where it is implemented. Policy x_2 constitutes a policy experiment that may turn out to be either good or bad for each member. If it is good for i, then she receives a period payoff $u_{i2} > 0$ with probability $\lambda > 0$, and zero with probability $1-\lambda$; if it is bad for her, she receives zero. Policy x_3 is also risky and, similarly, may be good or bad for each member. Unlike x_2, however, it entails the stochastic arrival of *bad* news. Specifically, whenever x_3 is selected by the committee, each i receives a certain payoff of $u_{i3} > 0$, from which a cost $c_i > 0$ is subtracted if bad news occurs. She receives bad news with probability $\mu > 0$ if x_3 is bad for her, and with probability zero otherwise.

We assume that if policy x_k, $k = 2, 3$, is good for members 1 and 2, then it is necessarily bad for member 3, and vice versa. Accordingly, we can write $\Omega_2 = \Omega_3 = \{\text{Good}, \text{Bad}\}$, with the common understanding that $\omega_k = \text{Good}$ indicates that x_k is good *only* for members 1 and 2. We further assume that members' assessments of each policy x_k, $k = 2, 3$, are represented by their beliefs that the policy is good, i.e., that $\omega_k = \text{Good}$. They hold common initial beliefs that x_2 and x_3 are good, drawn from the sets

$$P_2 \equiv \left\{ \frac{\bar{p}_2(1-\lambda)^l}{\bar{p}_2(1-\lambda)^l + 1 - \bar{p}_2} : l = 0, 1, \dots \right\} \cup \{1\} \ ,$$
$$P_3 \equiv \left\{ \frac{\bar{p}_3}{\bar{p}_3 + (1-\bar{p}_3)(1-\mu)^l} : l = 0, 1, \dots \right\} \cup \{0\} \ ,$$

respectively, for some $\bar{p}_2, \bar{p}_3 \in (0, 1)$. The set of assessment matrices $\mathcal{P}$ comprises all the (3×3) matrices of the form

$$\mathbf{P} = \begin{bmatrix} 1 & p_2 & p_3 \\ 1 & p_2 & p_3 \\ 1 & p_2 & p_3 \end{bmatrix} ,$$

where $p_2 \in P_2$ and $p_3 \in P_3$.

Committee members' beliefs are updated according to Bayes' rule, so that the transitions ϕ_{i1} and ϕ_{i2} are as in Example 1, and ϕ_{i3} is defined by:

$$\phi_{i3}(p_3, \text{Good}) \left(\frac{p_3}{p_3 + (1-p_3)(1-\mu)} \right) = 1 \ ,$$
$$\phi_{i3}(p_3, \text{Bad}) \left(\frac{p_3}{p_3 + (1-p_3)(1-\mu)} \right) = 1-\mu \ ,$$

and $\phi_{i3}(p, \text{Bad})\,(0) = \mu$, for all $p_3 \in P_3$.

Each committee member i is assumed to myopically maximize her expected payoff in each period. Thus, for any assessment matrix $\mathbf{P}$, we have

$$x_1 \, R_i(\mathbf{P}) \, x_2 \text{ if and only if } u_{i1} \geq p_2 \lambda u_{i2},$$
$$x_1 \, R_i(\mathbf{P}) \, x_3 \text{ if and only if } u_{i1} \geq u_{i3} - (1-p_3)\mu c_i,$$
$$x_2 \, R_i(\mathbf{P}) \, x_3 \text{ if and only if } p_2 \lambda u_{i2} \geq u_{i3} - (1-p_3)\mu c_i \ ,$$

for each member $i = 1, 2$. Member 3's experimentation preferences are defined in like manner, simply by replacing the probabilities p_2 and p_3 with $1-p_2$ and $1-p_3$, respectively, in each of the inequalities above.

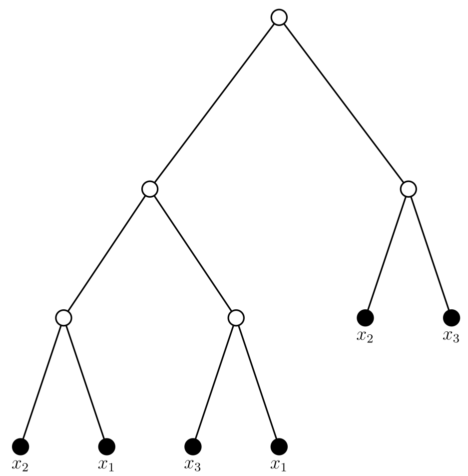

Figure 4: Agenda for the three-armed bandit environment.

In each period, collective decision-making is governed by the agenda $\overline{\Gamma}$, depicted in Figure 4. In the first round, the committee must choose between an amendment agenda and a pairwise vote between the two risky policies. If it opts for the amendment agenda, policy x_2 is first compared to x_3 in a pairwise vote, and the winner is then pitted against x_1 in a second pairwise vote.

Agenda-manipulability and adverse capture. There exist parametric configurations of the policy-making environment described above under which it is agenda-manipulable, but capture activities may have adverse consequences. To see this, consider the following parametrization: $\lambda = 0.8$, $\mu = 0.5$, $(c_1, c_2, c_3) = (1, 4, 3)$, $\bar{p}_2 = 0.7$, $\bar{p}_3 = 0.6$, and

$$[u_{ik}] = \begin{bmatrix} 8 & 4 & 2 \\ 2 & 16 & 8 \\ 3 & 8 & 4 \end{bmatrix}.$$

Let $\overline{\mathbf{P}}$ denote the assessment matrix in which $p_k = \bar{p}_k$ for each $k = 2, 3$. It is readily checked that these parameter values yield experimentation preferences for which the sophisticated outcome of the agenda $\overline{\Gamma}$ is $s(\overline{\Gamma}, \overline{\mathbf{P}}) = x_2$. Moreover, an agenda setter whose ideal policy is x_1 could change $\overline{\Gamma}$ to the successive-elimination agenda Γ' with voting sequence (x_1, x_2, x_3), thus obtaining her most-preferred policy as the sophisticated outcome. If, instead, x_3 is her most-preferred policy, she could enforce it by changing $\overline{\Gamma}$ to the amendment agenda Γ'' with voting sequence (x_2, x_3, x_1), for example. Thus, $s(\Gamma, \overline{\mathbf{P}}) \neq s(\overline{\Gamma}, \overline{\mathbf{P}})$, for $\Gamma = \Gamma', \Gamma''$. We conclude that whenever $\overline{\mathbf{P}}$ is the initial assessment matrix, an agenda setter can manipulate the agenda to her advantage from the very first period.

Consider now the position of a special-interest group that has the ability to enforce its favored policy x_3 in the first period, inducing the 3-captured outcome function

$$\hat{\sigma}(t, \mathbf{P}) = \begin{cases} x_3 & \text{if } t = 1, \\ \bar{\sigma}(t, \mathbf{P}) & \text{otherwise.} \end{cases}$$

It turns out that *such capture behavior can significantly reduce the likelihood that the group's preferred policy is selected in the long run.* Indeed, suppose the state is $\omega = (\text{Safe}, \text{Bad}, \text{Good})$. Under the default outcome function $\bar{\sigma}$, the committee experiments with x_2 in the first period and subsequently engages in repeated experimentation with x_3 from period $t = 2$ onward. Hence, $\mathfrak{n}_3^t(\bar{\sigma} \mid \omega) = t - 1$ for all $t \in \mathbb{N}$. By contrast, under the captured outcome function $\hat{\sigma}$, the committee has to experiment with x_3 in the first period but then switches to the safe policy x_1 from period $t = 2$ onward, so that $\mathfrak{n}_3^t(\hat{\sigma} \mid \omega) = 1$ for all $t \in \mathbb{N}$. It follows immediately that $\mathfrak{n}_3^t(\bar{\sigma} \mid \omega)$ first-order stochastically dominates $\mathfrak{n}_3^t(\hat{\sigma} \mid \omega)$ for all $t \geq 2$, with strict dominance for all $t > 2$.

4 Discussion and Extensions

A central issue in voting theory since its very beginnings has been the manipulation of policy-making institutions by those who control them. Most of the literature has focused on identifying environments in which such manipulation can be avoided. But are these environments also immune to another form of manipulation: influence by special interest groups? This paper has examined the relationship between these two forms of manipulation and reached a stark conclusion: whenever the conditions are met for a policy-making environment to be immune to institutional manipulation, it is highly vulnerable to capture by special interests. The generality of the environments to which this conclusion applies makes it an impossibility result, with significant implications for attempts to insulate legislative policy-making from manipulation by self-interested actors.

Evidently, if the set $\mathcal{A}$ of possible voting procedures available to the committee is enlarged, then the ability of those (e.g., the committee chair) responsible for selecting these procedures to manipulate the committee's decisions increases, thus reinforcing our result. It is however natural to ask whether our result still holds when the set of collective decision-making procedures available to the committee is instead restricted. After all, not every agenda in the family $\mathcal{A}$ that we have considered is commonly used in real-world legislative bodies. In fact, the literature on agenda manipulation has often focused on a specific subset of $\mathcal{A}$, i.e., the empirically-relevant class of amendment agendas. In this case, agenda control is limited to selecting a voting sequence for amendments. It is readily checked that all the arguments in the proof of Proposition 1 still hold when $\mathcal{A}$ is restricted to amendment agendas. This also applies when $\mathcal{A}$ is restricted to successive-elimination agendas.

For expositional convenience, we have assumed an (odd-sized) majoritarian committee. However, in many situations of interest — particularly outside legislative contexts — quota rules other than majority voting are used by committees. A well-known difficulty in modeling such situations is that one must make assumptions about how ties in the agendas' pairwise votes are broken; moreover, these assumptions must be applicable across all possible agendas in $\mathcal{A}$. In the case where $\mathcal{A}$ consists of amendment agendas, two natural assumptions commonly made in the literature (e.g., Banks and Bordes, 1988; Duggan, 2006; or Barberà and Gerber, 2017) are that, in the event of a procedural tie between two policies at a decision node, the winner is either the policy appearing earlier or the one appearing later in the agenda's voting sequence. As explained in the Appendix, our manipulation result also applies to cases where voting is by quotas other than majority rule (and/or the committee is even-sized) under either of these two assumptions.

Finally, we have considered situations in which one or more special-interest groups advocate for a single policy at a time. However, empirically relevant scenarios may arise in which (i) a group or coalition of groups simultaneously advocates for multiple policies, and/or (ii) groups aim not to capture the committee's decision-making process *in favor of* a given policy, but rather *against* certain policies. In Appendix D, we develop a framework that accommodates these scenarios and establish a corresponding manipulation result analogous to Proposition 1.

Appendix

A Binary Agendas: Formal Definitions

A *(binary) agenda* (or *voting tree*) is a tuple $\Gamma = \langle \Lambda, \succ, \theta \rangle$, whose three components are described as follows. First, Λ is a set of nodes. Second, $\succ$ is a transitive and asymmetric relation on Λ that satisfies the following two conditions: (i) there exists a unique node $\lambda_0 \in \Lambda$ such that $\lambda_0 \not\succ \lambda$ for all $\lambda \in \Lambda$; and (ii) for any pair $\lambda, \lambda' \in \Lambda$ such that $\lambda' \succ \lambda$, there exists a unique sequence of nodes $\lambda_1, \ldots, \lambda_m$ in Λ such that $\lambda' \succ \lambda_1 \succ \cdots \succ \lambda_m \succ \lambda$. Intuitively, the relation $\succ$ serves as a "precedence relation." Condition (i) ensures thus that there is a unique initial node that is not preceded by any other node, while condition (ii) ensures that there is exactly one path in the tree from λ to λ'. This relation partitions the set of nodes Λ into the subset of *terminal nodes* $\Lambda^T \equiv \{\lambda \in \Lambda \colon \nexists \lambda' \in \Lambda \text{ such that } \lambda' \succ \lambda\}$ — i.e., those that do not precede any other node — and the subset of *decision nodes* $\Lambda^D \equiv \Lambda \setminus \Lambda^T$. In a binary agenda, every decision node $\lambda \in \Lambda^D$ has exactly two nodes immediately following it. Finally, the third component of the agenda Γ is a surjective function $\theta \colon \Lambda^T \to X$, which assigns each terminal node of the tree to exactly one alternative in X.

The committee's final choice resulting from any agenda $\Gamma \in \mathcal{A}$ is characterized by the *sophisticated outcome* (Shepsle and Weingast, 1984) of that agenda. More specifically, each assessment matrix $\mathbf{P}$ induces a preference profile $\big(R_1(\mathbf{P}), \ldots, R_n(\mathbf{P})\big)$ on X. This profile in turn gives rise to a majority preference relation $R^m(\mathbf{P})$ on X: for all $x_k, x_\ell \in X$, $x_k R^m(\mathbf{P}) x_\ell$ if and only if $\big|\{i \in N \colon x_k R_i(\mathbf{P}) x_\ell\}\big| > n/2$. Equipped with $R^m(\mathbf{P})$, we can then apply the following *sophisticated voting* procedure to Γ. By construction, each non-terminal node λ in the tree is followed by exactly two terminal nodes, say λ' and λ''. The *sophisticated equivalent* of λ, denoted $\bar{s}(\lambda)$, is defined as the majority winner between $\theta(\lambda')$ and $\theta(\lambda'')$:

$$\bar{s}(\lambda) \equiv \begin{cases} \theta(\lambda') & \text{if } \theta(\lambda') R^m(\mathbf{P}) \theta(\lambda'') \ , \\ \theta(\lambda'') & \text{otherwise.} \end{cases}$$

Applying the same reasoning to every final decision node, we can delete the set of terminal nodes of Γ by associating each final decision node λ with its sophisticated equivalent $\bar{s}(\lambda)$, i.e., by converting each final decision node into a terminal node whose associated outcome is given by $\bar{s}(\lambda)$. This yields a new binary agenda. The same logic can then be applied to the final decision nodes of this new agenda, and so on, recursively up the tree, until we are left with a unique policy $\bar{s}(\lambda_0)$, which is the sophisticated equivalent of the initial node of Γ. This policy is referred to as the *sophisticated outcome* of Γ, and is denoted by $s(\Gamma, \mathbf{P})$.

B Proof of Proposition 1

Take any policy-making environment whose outcome is agenda-independent; so that we must have $s(\Gamma, \mathbf{P}) = s(\overline{\Gamma}, \mathbf{P})$, for all $\Gamma \in \mathcal{A}$ and all $\mathbf{P} \in \mathcal{P}$. A theorem, proved by McKelvey and Niemi (1978), establishes that the sophisticated outcome of any binary agenda Γ must be an element of the top cycle set, whose unique element is the maximum of the majority preference relation when there is one. It follows that for every $\mathbf{P} \in \mathcal{P}$, the majority preference relation induced by $\big(R_1(\mathbf{P}), \ldots, R_n(\mathbf{P})\big)$ has a unique maximum. (We reach the same conclusion in the variant of our model discussed in Section 4, where voting is by any quota rule in amendment agendas, using Barberà and Gerber's (2017) Theorem 5.1 instead of McKelvey and Niemi's.)

Now take an arbitrary policy $x_k \in X$, and let $\mathbf{P}$ and $\mathbf{Q}$ be any two assessment matrices in $\mathcal{P}$ such that $s(\overline{\Gamma}, \mathbf{P}), s(\overline{\Gamma}, \mathbf{Q}) \neq x_k$, and $p_{\cdot\ell} = q_{\cdot\ell}$, for all $\ell \neq k$. Let x_ℓ and x_m denote the maxima of the majority preference relations induced by $\mathbf{P}$ and $\mathbf{Q}$, respectively. Applying McKelvey and Niemi's theorem again, we must have $x_\ell = s(\overline{\Gamma}, \mathbf{P}) \neq x_k$ and $x_m = s(\overline{\Gamma}, \mathbf{Q}) \neq x_k$. This in turn implies that $p_{\cdot\ell} = q_{\cdot\ell}$ and $p_{\cdot m} = q_{\cdot m}$. Hence, by the independence-of-irrelevant-alternatives property, we have $x_\ell R_i(\mathbf{P}) x_m$ if and only if $x_\ell R_i(\mathbf{Q}) x_m$, for each $i \in N$. It follows that the majority preference relation between x_m and x_ℓ is the same under $\mathbf{P}$ and $\mathbf{Q}$; and therefore, $s(\overline{\Gamma}, \mathbf{P}) = x_\ell = x_m = s(\overline{\Gamma}, \mathbf{Q})$. (Otherwise, the sophisticated outcome of $\overline{\Gamma}$ under either $\mathbf{P}$ or $\mathbf{Q}$ would not be the maximum of the corresponding majority-preference relation; a contradiction.) As $x_k \in X$ and $\mathbf{P}, \mathbf{Q} \in \mathcal{P}$ were taken arbitrarily, we conclude that: for every policy x_k and assessment matrices $\mathbf{P}$ and

$\mathbf{Q}$ satisfying $s(\overline{\Gamma}, \mathbf{P}), s(\overline{\Gamma}, \mathbf{Q}) \neq x_k$, and $p_{\cdot\ell} = q_{\cdot\ell}$ for all $\ell \neq k$, we have $s(\overline{\Gamma}, \mathbf{P}) = s(\overline{\Gamma}, \mathbf{Q})$. Henceforth, we refer to the latter property more succinctly as "condition (C)."

What remains to establish in order to obtain Proposition 1, therefore, is that condition (C) is sufficient for the outcome of the policy-making environment under consideration to be prone to self-interest capture. In fact, this condition is a natural extension of Gossner et al.'s (2021) condition of independence of irrelevant alternative x_k for attention strategies in their individual decision-making framework with limited attention, for each $x_k \in X$. The rest of the proof consists in verifying that condition (C) suffices to extend the coupling argument underlying their attention theorem to our agenda-voting framework with multiple individuals, and in obtaining strict first-order stochastic dominance in at least one period. To this end, we must first introduce some terminology and mathematical machinery. For every $\omega = (\omega_1, \ldots, \omega_L) \in \Omega$, define a "learning process" π as a tuple of independent n-dimensional processes $\pi_\ell = \{\pi_\ell^m\}$, $\ell = 1, \ldots, L$, where each π_ℓ is a Markov process with the initial state $p_{\cdot\ell}^0$, and whose transitions are governed by the $\phi_{i\ell}(\cdot, \omega_\ell)$'s. Thus, the mth component of π_ℓ, π_ℓ^m, is equal to the assessment profile of policy x_ℓ after it has been implemented in m periods. Observe that for any state ω and any outcome function $\tilde{\sigma}$, we can define the joint stochastic process $\{(\mathbf{P}^t, \chi^t)\}$ of assessment matrices and chosen policies, with the set of possible draws of π as its sample space, as follows: $p_{\cdot\ell}^t \equiv \pi_\ell^{\mathfrak{n}_\ell^t(\tilde{\sigma}|\omega)}$, and the probability distribution of $\chi^t(\tilde{\sigma})$ is $\tilde{\sigma}(t, \mathbf{P}^{t-1})$, for every $t \in \mathbb{N}$. The law of this process is denoted by $\mathfrak{L}_{\tilde{\sigma}}^\omega$.

Next, fix some $\omega = (\omega_1, \ldots, \omega_L) \in \Omega$, $x_k \in X$, and k-captured outcome function $\hat{\sigma}$. For every assessment matrix $\mathbf{P}$, define an "implementation process" as a triple of stochastic processes $\big(A(\cdot \mid \mathbf{P}), B(\cdot \mid \mathbf{P}_{-k}), C(\cdot \mid \mathbf{P})\big)$, where $\mathbf{P}_{-k}$ denotes the $n \times (L-1)$ matrix obtained from $\mathbf{P}$ by deleting its kth column, such that for every $t \in \mathbb{N}$: (i) $A(t \mid \mathbf{P})$ is a $\{0,1\}$-valued random variable, which is equal to one with probability $\hat{\sigma}(t, \mathbf{P})(x_k)$; $B(t \mid \mathbf{P}_{-k}) \in X \setminus \{x_k\}$ is defined as policy $x_\ell \neq x_k$ if there exists $p_{\cdot k}$ such that $s(\overline{\Gamma}, p_{\cdot k}\mathbf{P}_{-k}) = x_\ell$ (where $p_{\cdot k}\mathbf{P}_{-k}$ is the assessment matrix obtained by adding the assessment profile $p_{\cdot k}$ to the matrix $\mathbf{P}_{-k}$, as its kth column), and as an arbitrary element of $X \setminus \{x_k\}$ otherwise; and $C(t \mid \mathbf{P})$ is a $\{0,1\}$-valued random variable which is equal to one with probability

$\mathbf{1}_{\{s(\Gamma,\mathbf{P})=x_k\}}/\hat{\sigma}(t,\mathbf{P})(x_k)$ if $\hat{\sigma}(t,\mathbf{P})(x_k) > 0$, and with probability zero otherwise. Observe that the definition of the $B(t \mid \mathbf{P}_{-k})$'s is where the condition (C), established above, kicks in: it guarantees that each $B(t \mid \mathbf{P}_{-k})$ is well-defined, since the choice of the profile $p_{\cdot k}$ such that $s(\overline{\Gamma}, p_{\cdot k}\mathbf{P}_{-k}) \neq x_k$ (if there is any) is irrelevant under this condition. Observe further that the probability distribution of each $C(t \mid \mathbf{P})$ is also well-defined, since $\hat{\sigma}$ is a k-captured outcome function and, consequently, $\hat{\sigma}(t,\mathbf{P})(x_k) = 1$ whenever $s(\overline{\Gamma},\mathbf{P}) = x_k$.

Equipped with the implementation processes, we now turn to the construction of joint stochastic processes $\{(\mathbf{P}^t(\tilde{\sigma}), \chi^t(\tilde{\sigma}))\}$, for each $\tilde{\sigma} \in \{\bar{\sigma}, \hat{\sigma}\}$. As above, the ℓth column of each $\mathbf{P}^t(\tilde{\sigma})$, $p^t_{\cdot\ell}(\tilde{\sigma})$, is defined as the assessment profile of the n committee members after policy x_ℓ has been implemented in exactly $\mathfrak{n}^t_\ell(\tilde{\sigma} \mid \omega)$ periods. Having defined the $\mathbf{P}^t(\tilde{\sigma})$'s, we can now construct the $\chi^t(\tilde{\sigma})$'s as follows. Let $\nu^t(\omega, \mathbf{P}_{-k} \mid \tilde{\sigma}) \equiv \big|\{\tau < t \colon \mathbf{P}^\tau_{-k} = \mathbf{P}_{-k} \ \&\ \chi^\tau(\tilde{\sigma}) \neq x_k\}\big|$. Then, let

$$\chi^t(\bar{\sigma}) = \begin{cases} x_k & \text{if } A\big(t \mid \mathbf{P}^{t-1}(\bar{\sigma})\big)C\big(t \mid \mathbf{P}^{t-1}(\bar{\sigma})\big) = 1 \\ B\big(\nu^t(\omega, \mathbf{P}^{t-1}_{-k} \mid \bar{\sigma}) \mid \mathbf{P}^{t-1}_{-k}(\bar{\sigma})\big) & \text{otherwise,} \end{cases}$$

and

$$\chi^t(\hat{\sigma}) = \begin{cases} x_k & \text{if } A\big(t \mid \mathbf{P}^{t-1}(\hat{\sigma})\big) = 1 \\ B\big(\nu^t(\omega, \mathbf{P}^{t-1}_{-k} \mid \hat{\sigma}) \mid \mathbf{P}^{t-1}_{-k}(\hat{\sigma})\big) & \text{otherwise,} \end{cases}$$

for all $t \in \mathbb{N}$. As desired for the coupling argument, the law of the joint stochastic process $\{(\mathbf{P}^t(\tilde{\sigma}), \chi^t(\tilde{\sigma}))\}$ is by construction $\mathfrak{L}^\omega_{\tilde{\sigma}}$, for each $\tilde{\sigma} \in \{\bar{\sigma}, \hat{\sigma}\}$.

To complete the proof of the proposition, therefore, it suffices to show that for every $\omega \in \Omega$: (i) for every period $t \in \mathbb{N}$, $\mathfrak{n}^t_k(\hat{\sigma} \mid \omega) \equiv \{\tau \leq t \colon \chi^\tau(\hat{\sigma}) = x_k\}$ first-order stochastically dominates $\mathfrak{n}^t_k(\bar{\sigma} \mid \omega)$; and (ii) this first-order stochastic dominance is strict for at least one $t \in \mathbb{N}$.

We begin with part (i). Fix $\omega \in \Omega$. Because $\bar{\sigma}$ is deterministic, then by definition of a k-captured outcome function, we have that for every draw, policy x_k is implemented under $\hat{\sigma}$ whenever it is implemented under $\bar{\sigma}$ in period 1. Proceeding by induction, suppose that $\mathfrak{n}^t_k(\bar{\sigma} \mid \omega) \leq \mathfrak{n}^t_k(\hat{\sigma} \mid \omega)$ in all periods up to some fixed t. For every arbitrary draw where $\mathfrak{n}^t_k(\bar{\sigma} \mid \omega) < \mathfrak{n}^t_k(\hat{\sigma} \mid \omega)$, the condition trivially holds at $t+1$. Now, take a draw where $\mathfrak{n}^t_k(\bar{\sigma} \mid \omega) = \mathfrak{n}^t_k(\hat{\sigma} \mid \omega)$, so that the kth columns of $\mathbf{P}^t(\bar{\sigma})$ and $\mathbf{P}^t(\hat{\sigma})$ coincide. It follows from a variant on Gossner et al.'s (2021) Lemma 2 that the other columns coincide as well, i.e., the joint processes $\{(\mathbf{P}^t_{-k}(\bar{\sigma}), \chi^t(\bar{\sigma}))\}$ and $\{(\mathbf{P}^t_{-k}(\hat{\sigma}), \chi^t(\hat{\sigma}))\}$ are identical when restricted to periods in which policy x_k is not implemented. Formally, let $\overline{T}(m)$ denote the mth period t such that $\chi^t \neq x_k$ under the outcome function $\bar{\sigma}$; and let $\widehat{T}(m)$ denote the mth period t such that $\chi^t \neq x_k$ under the outcome function $\hat{\sigma}$. We then have $(\mathbf{P}^{\overline{T}(m)}_{-k}(\bar{\sigma}), \chi^{\overline{T}(m)}(\bar{\sigma})) = (\mathbf{P}^{\widehat{T}(m)}_{-k}(\hat{\sigma}), \chi^{\widehat{T}(m)}(\hat{\sigma}))$, for all m. (The proof parallels their lemma, just substituting assessment (sub)matrices in our multiple-agent framework to assessment (sub)vectors in their single-agent framework.) Coupled with the observation that the kth columns of $\mathbf{P}^t(\bar{\sigma})$ and $\mathbf{P}^t(\hat{\sigma})$ coincide, we obtain that $\mathbf{P}^t(\bar{\sigma}) = \mathbf{P}^t(\hat{\sigma})$ in the draw under consideration, it follows from the construction of χ^t above that $\chi^{t+1}(\bar{\sigma}) = x_k$ only if $\chi^{t+1}(\hat{\sigma}) = x_k$. Hence, $\mathfrak{n}^{t+1}_k(\bar{\sigma} \mid \omega) \leq \mathfrak{n}^{t+1}_k(\hat{\sigma} \mid \omega)$, as desired. As this inequality holds for every draw, we obtain part (i).

We now move to part (ii). Here, we use the observation that a random variable X strictly first-order stochastically dominates another random variable Y if and only if the following two conditions hold: X first-order stochastically dominates Y, and their cumulative distribution functions are not equal (e.g., Chapter 1 in Shaked and Shanthikumar, 2007). It follows from part (i) above that, for all t, $\mathfrak{n}^t_k(\hat{\sigma} \mid \omega)$ first-order stochastically dominates $\mathfrak{n}^t_k(\bar{\sigma} \mid \omega)$. What remains to be established, therefore, is that there is some period t for which the cumulative distribution functions of $\mathfrak{n}^t_k(\bar{\sigma} \mid \omega)$ and $\mathfrak{n}^t_k(\hat{\sigma} \mid \omega)$ are different.

To this end, we first need to establish some additional notation. Fix any $\omega \in \Omega$. For every learning draw π, let $t^*(\pi) \equiv \inf\{t \in \mathbb{N} \colon \bar{\sigma}(t, \mathbf{P}^{t-1})(x_k) < \hat{\sigma}(t, \mathbf{P}^{t-1})(x_k)\} \in \mathbb{N} \cup \{\infty\}$, where $\mathbf{P}^{t-1}$ is the realization of the period-$(t-1)$ assessment matrix in π. By construction, therefore, $\{(\mathbf{P}^t(\bar{\sigma}), \chi^t(\bar{\sigma}))\}$ and $\{(\mathbf{P}^t(\hat{\sigma}), \chi^t(\hat{\sigma}))\}$ are identical up to period $t = t^*(\pi) - 1$. This in turn implies that $\mathfrak{n}^{t^*(\pi)-1}_k(\bar{\sigma} \mid \omega) = \mathfrak{n}^{t^*(\pi)-1}_k(\hat{\sigma} \mid \omega)$. Next, for every $t \in \mathbb{N}$, let $\Pi^t \equiv \{\pi \in \Pi : t^*(\pi) = t\}$; and let $t^{**} \equiv \inf\{t \in \mathbb{N} \colon \Pi^t \text{ is not negligible}\}$. Condition (ii) in the definition of a k-captured outcome function guarantees that $t^{**} \in \mathbb{N}$, because we must have $\sum_{t=1}^{\tau} \Pr\{\pi \in \Pi^t\} = 1$ for any τ satisfying condition (ii) and the fact that $t^{**} \leq \tau$. Observe that by construction, $\mathfrak{n}^{t^{**}}_k(\tilde{\sigma} \mid \omega) = \mathfrak{n}^{t^{**}-1}_k(\tilde{\sigma} \mid \omega) + \mathbf{1}_{\chi^{t^{**}}(\tilde{\sigma}) = x_k}$ for each $\tilde{\sigma} = \bar{\sigma}, \hat{\sigma}$. It follows that for every $\mathfrak{n} \in \{0, 1, \cdots, t^{**}-1\}$ such that $\Pr\{\mathfrak{n}^{t^{**}-1}_k(\tilde{\sigma} \mid \omega) = \mathfrak{n} \mid \pi \in \Pi^{t^{**}}\} > 0$, we have

$$
\begin{aligned}
&\Pr\left\{\mathfrak{n}_k^{t^{**}}(\tilde{\sigma}\mid\omega)\le\mathfrak{n}\right\}\\
&=\Pr\left\{\mathfrak{n}_k^{t^{**}}(\tilde{\sigma}\mid\omega)\le\mathfrak{n}\mid\pi\in\Pi^{t^{**}}\right\}\Pr\{\pi\in\Pi^{t^{**}}\}\\
&\qquad\qquad+\Pr\left\{\mathfrak{n}_k^{t^{**}}(\tilde{\sigma}\mid\omega)\le\mathfrak{n}\mid\pi\notin\Pi^{t^{**}}\right\}\Pr\{\pi\notin\Pi^{t^{**}}\}\\
&=\Big[\Pr\left\{\mathfrak{n}_k^{t^{**}-1}(\tilde{\sigma}\mid\omega)<\mathfrak{n}\mid\pi\in\Pi^{t^{**}}\right\}+\Pr\left\{\mathfrak{n}_k^{t^{**}-1}(\tilde{\sigma}\mid\omega)=\mathfrak{n}\mid\pi\in\Pi^{t^{**}}\right\}\\
&\qquad\times\Pr\left\{\mathbf{1}_{\chi^{t^{**}}(\tilde{\sigma})=x_k}=0\mid\mathfrak{n}_k^{t^{**}-1}(\tilde{\sigma}\mid\omega)=\mathfrak{n},\pi\in\Pi^{t^{**}}\right\}\Big]\Pr\{\pi\in\Pi^{t^{**}}\}\\
&\qquad\qquad+\Pr\left\{\mathfrak{n}_k^{t^{**}}(\tilde{\sigma}\mid\omega)\le\mathfrak{n}\mid\pi\notin\Pi^{t^{**}}\right\}\Pr\{\pi\notin\Pi^{t^{**}}\}\ .
\end{aligned}
$$

Observe first that, since $\mathfrak{n}_k^{t^{**}-1}(\hat{\sigma}\mid\omega)=\mathfrak{n}_k^{t^{**}-1}(\bar{\sigma}\mid\omega)$ for every draw in $\Pi^{t^{**}}$, we have $\Pr\left\{\mathfrak{n}_k^{t^{**}-1}(\hat{\sigma}\mid\omega)<\mathfrak{n}\mid\pi\in\Pi^{t^{**}}\right\}=\Pr\left\{\mathfrak{n}_k^{t^{**}-1}(\bar{\sigma}\mid\omega)<\mathfrak{n}\mid\pi\in\Pi^{t^{**}}\right\}$ and $\Pr\left\{\mathfrak{n}_k^{t^{**}-1}(\hat{\sigma}\mid\omega)=\mathfrak{n}\mid\pi\in\Pi^{t^{**}}\right\}=\Pr\left\{\mathfrak{n}_k^{t^{**}-1}(\bar{\sigma}\mid\omega)=\mathfrak{n}\mid\pi\in\Pi^{t^{**}}\right\}$. Moreover, by construction, we also have that $\mathfrak{n}_k^{t^{**}}(\hat{\sigma}\mid\omega)=\mathfrak{n}_k^{t^{**}}(\bar{\sigma}\mid\omega)$ for all draws outside $\Pi^{t^{**}}$; so that $\Pr\left\{\mathfrak{n}_k^{t^{**}}(\hat{\sigma}\mid\omega)\le\mathfrak{n}\mid\pi\notin\Pi^{t^{**}}\right\}=\Pr\left\{\mathfrak{n}_k^{t^{**}}(\bar{\sigma}\mid\omega)\le\mathfrak{n}\mid\pi\notin\Pi^{t^{**}}\right\}$. By the same logic as in part (i), $\mathbf{P}^{t^{**}-1}(\hat{\sigma})=\mathbf{P}^{t^{**}-1}(\bar{\sigma})$, so that $\chi^{t^{**}}(\bar{\sigma})=x_k$ only if $\chi^{t^{**}}(\hat{\sigma})=x_k$, in every draw in $\Pi^{t^{**}}$. We then have that in every draw of $\Pi^{t^{**}}$, $\bar{\sigma}\big(t^{**},\mathbf{P}^{t^{**}-1}(\bar{\sigma})\big)(x_k)<\hat{\sigma}\big(t^{**},\mathbf{P}^{t^{**}-1}(\hat{\sigma})\big)(x_k)$. This implies that $\Pr\left\{\mathbf{1}_{\chi^{t^{**}}(\hat{\sigma})=x_k}=0\mid\mathfrak{n}_k^{t^{**}-1}(\hat{\sigma}\mid\omega)=\mathfrak{n},\pi\in\Pi^{t^{**}}\right\}<\Pr\left\{\mathbf{1}_{\chi^{t^{**}}(\bar{\sigma})=x_k}=0\mid\mathfrak{n}_k^{t^{**}-1}(\bar{\sigma}\mid\omega)=\mathfrak{n},\pi\in\Pi^{t^{**}}\right\}$. As $\Pr\{\pi\in\Pi^{t^{**}}\}>0$, we conclude that $\Pr\left\{\mathfrak{n}_k^{t^{**}}(\hat{\sigma}\mid\omega)\le\mathfrak{n}\right\}<\Pr\left\{\mathfrak{n}_k^{t^{**}}(\bar{\sigma}\mid\omega)\le\mathfrak{n}\right\}$ and, consequently, $\mathfrak{n}_k^{t^{**}}(\bar{\sigma}\mid\omega)$ and $\mathfrak{n}_k^{t^{**}}(\hat{\sigma}\mid\omega)$ have different cumulative distribution functions. This completes the proof of the proposition.

C Example 1 (Continued): Collective Experimentation and Vote Buying

We consider a variant on the collective-experimentation model of Anesi and Bowen (2021), described in Example 1. A committee of $n\ge 3$ members (with n odd) meets in discrete time $t=0,1,2,\ldots$. At the start of each period, members vote simultaneously by simple majority on whether to implement a safe policy x_1 or pursue a risky reform x_2.

The reform produces uncertain returns that depend on whether it is "good" or "bad." If it is good, it succeeds in any given period with probability $\lambda > 0$, in which case all members enjoy heterogeneous benefits. If it is bad, it never succeeds. By contrast, the safe policy yields a sure payoff to each member every period. Member i receives stage payoff $b_i > 0$ in the event of a successful reform, and $s > 0$ from the safe policy. We assume that conditional on the reform being good, its expected payoff exceeds that of the safe alternative. For expositional ease, we further asume that $b_1 < \cdots < b_n$.

Committee members share a common prior belief $\bar{p} \in (0, 1)$ that the reform is good. Beliefs are updated in a Bayesian manner after each unsuccessful attempt, implying that the collective belief evolves endogenously over time. A first success fully reveals the true quality of the reform, and beliefs jump to one thereafter.

We enrich this baseline environment by introducing an additional player: a vote-buying special-interest group as in Groseclose and Snyder (1996), which lobbies in favor of the risky reform. More precisely, this player earns a policy benefit of $b_0 > 0$ in every period where the reform x_2 is implemented, and receives zero whenever the safe policy x_1 is chosen.

At the beginning of each period t, before the committee votes, the group may seek to capture the committee's vote. It selects a vector of take-it-or-leave-it offers $(m_1^t, \ldots, m_n^t) \in \mathbb{R}_+^n$, where m_i^t denotes the monetary transfer promised to member i in exchange for a favorable vote. Committee members then simultaneously decide whether to accept or refuse the offer made to them. The subsequent voting stage unfolds as follows: all members who accepted the group's offer are committed to voting for x_2, while those who declined remain free to cast their vote for either alternative. We assume that the interest group faces no budget constraint, so that the sole limit on the payments it offers is its own (endogenous) willingness to pay for changing the committee's decision.

In each period, the stage-payoff of committee member i is given by the sum of her expected policy payoff and any payment offered by the special-interest group (if accepted). Similarly, the stage payoff of the special-interest group is equal to its expected policy benefit minus the total of the payments accepted by committee members. All players maximize the expected discounted value of their stage-payoffs, using a common discount factor $\delta \in (0, 1)$. The solution concept is Markov perfect equilibrium. In order to limit the number of possible cases (without affecting our conclusions), we assume that in case of a tie: (i) committee members who are indifferent vote in favor of the risky reform; and (ii) the special-interest group, when indifferent between offering a positive amount or zero, chooses to offer the positive amount. We refer to any pure-strategy Markov perfect equilibrium with these tie-breaking rules simply as an "equilibrium."

A particularly intuitive result, formally established in Anesi and Bowen (2021), is that the equilibrium continuation value of implementing the risky reform is strictly increasing in the benefit b_i of each committee member. Consequently, the median committee member $m \equiv (n + 1)/2$ is pivotal, regardless of the discount factor δ, in evaluating (dynamically) the risky reform x_2 relative to the safe policy x_1, when potential transfers from the special-interest group are ignored.

It follows that in any period in which it is optimal for the group to purchase votes, it is optimal to buy the minimal winning coalition that is least costly, i.e., $C \equiv \{m, \ldots, n\}$. The optimal offer to each member of this coalition is the minimum amount that they are willing to accept, i.e., the difference between the equilibrium continuation values (of her policy benefits) under x_1 and x_2.

In the trivial case where $b_0 \geq ms$, buying off every member of C always benefits the group, since the net surplus is strictly positive. Consider instead the case $b_0 < ms$. As the belief p tends to zero, the expected payoff from implementing the reform also converges to zero. The cost of purchasing the cheapest majority coalition C therefore exceeds the group's willingness to pay, implying that abstaining from intervention is optimal. We thus conclude that there exists a sufficiently small threshold $\underline{p} > 0$ such that whenever beliefs fall below $\underline{p}$, the group refrains from vote buying and the committee implements x_1 forever in any equilibrium.

We can then use the fact that the belief that the reform is good decreases discretely with each unsuccessful attempt at implementing it to characterize equilibrium behavior, using standard recursive arguments. More specifically, let $\underline{p}'$ denote the largest belief at which if the committee were to implement x_2 and the attempt failed, the belief at the beginning of the next period would fall under $\underline{p}$. In any period where the belief equals $\underline{p}'$, the arguments above imply that we know the (unique) equilibrum continuation values associated with implementing each policy for all committee members, and thus the minimum transfer each would require to vote in favor of x_2. This in turn determines the group's optimal decision (whether to buy the coalition C or not) and therefore the equilibrium continuation values for all players when beliefs are $\underline{p}'$. In particular, the group chooses to buy off C in any such period if and only if the monetrary cost of doing so does not exceed b_0. It is also worth noting that, when it is optimal for it to buy off coalition C, the group' transfers are chosen so that coalition members are just indifferent between implementing x_1 and x_2. It follows that their equilibrium continuation payoffs at $\underline{p}'$ must be s — exactly as in the case without the group.

Proceeding in this way, we recursively work backward to the initial belief $\bar{p}$, fully characterizing equilibrium behavior. Note in particular that the equilibrium continuation values obtained in this way are unique and, therefore, the equilibrium itself is unique.

Combined with the observation that the group always leaves members of C indifferent between accepting and rejecting its offers, the fact that the continuation value of implementing x_2 increases with the belief implies that it is less costly for the group to buy off committee members for larger values of p. In fact, the group secures x_2 without paying anything as long as $p > q_m$, since the committee then implements x_2 without any outside intervention. It is only when the belief p falls below q_m that intervention becomes necessary. At the first period in which this occurs (if any), the group buys off C if and only if the net benefit of doing so exceeds the monetary cost; that is, if its policy benefit b_0 exceeds a threshold $\underline{b} > 0$.

We conclude that whenever $b_0 \geq \underline{b}$, there exist two time cutoffs $t_1, t_2 \in \mathbb{N}$, with $1 \leq t_1 \leq t_2$, such that the above game has a unique equilibrium in which the group captures the committee's decision making between periods t_1 and t_2 (unless the committee discovers that x_2 is good in the meantime). During this interval, the expected frequency with which the risky reform is implemented is strictly greater than it would be in the benchmark game without a special-interest group. This can be seen as noncooperative version of Proposition

1: as there are only two feasible policies, the environment is trivially agenda-independent and, clearly, the group's capture efforts between t_1 and t_2 never backfire in equilibrium. Any strategy profile in which the group refrains from interfering in the committee's decision-making process is unsustainable in equilibrium, as any deviation toward capture — yielding a k-captured outcome path in the continuation equilibrium — would be profitable.

D Multi-Policy Capture

In the framework analyzed in the main text, we primarily consider situations in which one or more groups advocate for a single policy at a time. However, empirically relevant scenarios may arise in which (i) a group or coalition of groups simultaneously advocates for multiple policies; and/or (ii) groups seek not to capture the committee's decision-making process *in favor of* a given policy, but rather *against* certain policies. In this section, we propose a variant of Definition 2 that incorporates these extensions and establish a corresponding version of the manipulation result for this setting.

For every nonempty set $S \subset X$ and every outcome function σ, we indulge in a slight abuse of notation by letting $\sigma(t,\mathbf{P})(S) \equiv \sum_{x\in S} \sigma(t,\mathbf{P})(x)$, for all $t \in \mathbb{N}$ and $\mathbf{P} \in \mathcal{P}$. Recall that the default outcome function is defined by $\bar{\sigma}(t,\mathbf{P}) \equiv s(\overline{\Gamma},\mathbf{P})$. We say that $\hat{\sigma}\colon \mathbb{N}\times\mathcal{P} \to \Delta(X)$ is an *S-captured outcome function* if there exists $t_0 \in \mathbb{N}$ such that the following conditions hold:

(i) for all $t \in \mathbb{N}$ and $\mathbf{P} \in \mathcal{P}$: (i.a) $\bar{\sigma}(t,\mathbf{P}) = \hat{\sigma}(t,\mathbf{P})$ if $t \neq t_0$; (i.b) $\bar{\sigma}(t_0,\mathbf{P})(S) \leq \hat{\sigma}(t_0,\mathbf{P})(S)$; and (i.c) if $\hat{\sigma}(t_0,\mathbf{P})(S) < 1$, then $\frac{\hat{\sigma}(t_0,\mathbf{P})(x)}{1-\hat{\sigma}(t_0,\mathbf{P})(S)} = \bar{\sigma}(t_0,\mathbf{P})(x)$ for all $x \notin S$;

(ii) for every $\omega \in \Omega$, there exists $t \geq t_0$ such that $\mathbb{E}\big[\bar{\sigma}\big(t,\mathbf{P}^{t-1}(\bar{\sigma})\big)(S)\big] < \mathbb{E}\big[\hat{\sigma}\big(t,\mathbf{P}^{t-1}(\hat{\sigma})\big)(S)\big]$.

Condition (i) first requires that capture occur at a single period t_0. It further requires that in that period, capture affect only the total probability assigned to policies in S. That is, conditional on the selected policy being outside S, the relative probabilities coincide with those implied by the default outcome function. Condition (ii) ensures that capture strictly increases the expected probability that an alternative in S is selected.

Two remarks concerning this definition are in order. First, the notion of an S-captured outcome is not a generalization of that of a k-captured outcome. In particular, a k-captured outcome function need not be $\{x_k\}$-captured (while the reverse is true). The reason is that S-capture requires deviations from the default outcome function $\bar{\sigma}$ to occur in a fixed period t_0 (that may depend on ω). This requirement is nonetheless consistent with the noncooperative interpretation in terms of one-shot deviations, developed in Subsection 2.2. Second, S-capture may equivalently be interpreted as capture *against* a policy $x_k \notin S$. Indeed, opposing a given policy necessarily entails increasing the probability of implementation of a set of alternative policies, potentially the entire set $X \setminus \{x_k\}$.

It turns out that a manipulation result in the same vein as Proposition 1 extends to S-capture. For every nonempty set $S \subset X$ and every outcome function σ, let $\mathfrak{n}_S^t(\sigma \mid \omega)$ denote the random number of periods up to period $t \in \mathbb{N}$ in which an element of S is implemented in state ω under σ.

Proposition D1. *Suppose the policy-making environment is agenda independent. Then, the following holds for every nonempty $S \subset X$ and every S-captured outcome function $\hat{\sigma}$: for all $\omega \in \Omega$ and $t \in \mathbb{N}$, $\mathfrak{n}_S^t(\hat{\sigma} \mid \omega)$ first-order stochastically dominates $\mathfrak{n}_S^t(\bar{\sigma} \mid \omega)$, with strict dominance for at least one period.*

Proof. Take any $S \subset X$ and any S-captured outcome function $\hat{\sigma}$. Let t_0 be the period at which $\hat{\sigma}$ "departs" from $\bar{\sigma}$. Then fix any $\omega \in \Omega$, and any draw of the resulting learning process, as in the proof of Proposition 1. We consider in turn two possible types of events. First, consider the event in which capture "succeeds" in favor of S in t_0, i.e., the realization of the period-t_0 policy choice is some $x_k \in S$ under $\hat{\sigma}$, and any other policy outside S under $\bar{\sigma}$. In any such event, conditional on the realized assessment path up to $t_0 - 1$, the outcome function $\hat{\sigma}$ coincides with the default $\bar{\sigma}$ for all $t \neq t_0$ and $\mathbf{P} \in \mathcal{P}$. Hence, conditional on this realization, the deviation is a k-captured outcome function. By the same argument as in the proof of Proposition 1, the corresponding counting process $\mathfrak{n}_k^t(\hat{\sigma} \mid \omega)$ first-order stochastically dominates $\mathfrak{n}_k^t(\bar{\sigma} \mid \omega)$, with strict dominance in at least one period. It follows that, for this realization,

$$\begin{aligned}\mathfrak{n}_S^t(\hat{\sigma} \mid \omega) = \mathfrak{n}_k^t(\hat{\sigma} \mid \omega) + \sum_{\ell \in S \setminus \{k\}} \mathfrak{n}_\ell^t(\hat{\sigma} \mid \omega) = \mathfrak{n}_k^t(\hat{\sigma} \mid \omega) + \sum_{\ell \in S \setminus \{k\}} \mathfrak{n}_\ell^t(\bar{\sigma} \mid \omega) \\ \succeq_{\text{SD}} \mathfrak{n}_k^t(\bar{\sigma} \mid \omega) + \sum_{\ell \in S \setminus \{k\}} \mathfrak{n}_\ell^t(\bar{\sigma} \mid \omega) = \mathfrak{n}_S^t(\bar{\sigma} \mid \omega) \ ,\end{aligned}$$

where $\succeq_{\text{SD}}$ denotes the first-order stochastic dominance relation.

Now consider any event in which capture in favor of S "fails," so that the two processes $\left\{(\mathbf{P}^t(\tilde{\sigma}),\ \chi^t(\tilde{\sigma}))\right\}_{t \geq 1}$, $\tilde{\sigma} \in \{\hat{\sigma}, \bar{\sigma}\}$ coincide pathwise. In particular, $\mathfrak{n}_S^t(\hat{\sigma} \mid \omega) = \mathfrak{n}_S^t(\bar{\sigma} \mid \omega)$, for all $t \in \mathbb{N}$. We conclude that for all $t \in \mathbb{N}$ and $\omega \in \Omega$, $\mathfrak{n}_S^t(\hat{\sigma} \mid \omega)$ first-order stochastically dominates $\mathfrak{n}_S^t(\bar{\sigma} \mid \omega)$, with strict dominance in at least one period. □

References

Anesi, V., and T.R. Bowen (2021) "Policy Experimentation in Committees: A Case against Veto Rights under Redistributive Constraints," *American Economic Journal: Microeconomics* 13, 124-162.

Austen-Smith, D., and J.S. Banks (2005) *Positive Political Theory II: Strategy and Structure.* University of Michigan Press.

Baker, S.R., S.J. Davis, and J.A. Levy (2022) "State-Level Economic Policy Uncertainty," *Journal of Monetary Economics* 132, 81-99.

Banks, J.S. (1985) "Sophisticated Voting Outcomes and Agenda Control," *Social Choice and Welfare* 1, 295-306.

Banks, J.S., and G.A. Bordes (1988) "Voting Games, Indifference, and Consistent Sequential Choice Rules," *Social Choice and Welfare* 5, 31-44.

Barberà, S., and A. Gerber (2017) "Sequential Voting and Agenda Manipulation," *Theoretical Economics* 12, 211-247.

Barberà, S., and A. Gerber (2022) "Deciding on What to Decide," *International Economic Review* 63, 37-61.

Bils, P., and F. Izzo (2024) “Policy Gambles and Valence in Elections.” Available at SSRN 4847072.

Black, D. (1958) *The Theory of Committees and Elections.* London: Cambridge University Press.

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

Chen, Y., and H. Eraslan (2017) “Dynamic Agenda Setting,” *American Economic Journal: Microeconomics* 9, 1-32.

Denzau, A.T., and M.C. Munger (1986) “Legislators and Interest Groups: How Unorganized Interests Get Represented,” *American Political Science Review* 80, 89-106.

Duggan, J. (2006) “Endogenous Voting Agendas,” *Social Choice and Welfare* 27, 495-530.

Farquharson, R. (1969) *Theory of Voting.* New Haven: Yale University Press.

Ferejohn, J., M. Fiorina, and R.D. McKelvey (1987) “Sophisticated Voting and Agenda Independence in the Distributive Politics Setting,” *American Journal of Political Science* 31, 169-193.

Ferejohn, J.A., and D. Grether (1974) “On a Class of Rational Social Decision Procedures,” *Journal of Economic Theory* 8, 471-482.

Freer, M., C. Martinelli, and S. Wang (2020) “Collective Experimentation: A Laboratory Study,” *Journal of Economic Behavior & Organization* 175, 365-379.

Gans, J.S., and M. Smart (1996) “Majority Voting with Single-Crossing Preferences,” *Journal of Public Economics* 59, 219-237.

Gieczewski, G., and S. Kosterina (2024) “Experimentation in Endogenous Organizations,” *Review of Economic Studies* 91, 1711-1745.

GILLIGAN, T., AND K. KREHBIEL (1987) "Collective Decision Making and Standing Committees: An Informational Rationale for Restrictive Amendment Procedures," *Journal of Law, Economics, and Organization* 3, 287-335.

GRANDMONT, J.-M. (1978) "Intermediate Preferences and the Majority Rule," *Econometrica* 46, 317-330.

GOSSNER, O., J. STEINER, AND C. STEWART (2021) "Attention Please!" *Econometrica* 89, 1717-1751.

GROSECLOSE, T., AND J.M. SNYDER JR. (1996) "Buying Supermajorities in Legislatures," *American Political Science Review* 90, 303-315.

GROSSMAN, G.M., AND E. HELPMAN (2001) *Special Interest Politics.* MIT Press.

HWANG, I. (2023) "Policy Experimentation with Repeated Elections," *Games and Economic Behavior* 142, 623-644.

IZZO, F. (2023) "Ideology for the Future," *American Political Science Review* 117, 1089-1104.

MCKELVEY, R.D., AND R.G. NIEMI (1978) "A Multistage Game Representation of Sophisticated Voting for Binary Procedures," *Journal of Economic Theory* 18, 1-22.

MILLER, N.R. (1980) "A New Solution Set for Tournaments and Majority Voting," *American Journal of Political Science* 24, 68-96.

MILLER, N.R. (1995) *Committees, Agendas and Voting.* Chur: Harwood Academic Publishers.

MOULIN, H. (1979) "Dominance-Solvable Voting Schemes," *Econometrica* 47, 1337-1351.

ORDESHOOK, P.C., AND T. SCHWARTZ (1987) "Agendas and the Control of Political Outcomes," *American Political Science Review* 81, 179-199.

PLOTT, C.R., AND M.E. LEVINE (1978) "A Model of Agenda Influence on Committee Decisions," *American Economic Review* 68, 146-160.

RIKER, W.H. (1982) *Liberalism Against Populism.* San Francisco: W.H. Freeman.

RIKER, W.H. (1986) *The Art of Political Manipulation.* New Haven, CT: Yale University Press.

SHAKED, M., AND J.G. SHANTHIKUMAR (2007) *Stochastic Orders.* Springer, New York.

SHEPSLE, K.A. AND B.R. WEINGAST (1984) "Uncovered Sets and Sophisticated Voting Outcomes with Implications for Agenda Institutions," *American Journal of Political Science* 28, 49-74.

SHEPSLE, K.A., AND B.R. WEINGAST (1987) "The Institutional Foundations of Committee Power," *American Political Science Review* 81, 85-104.

SNYDER, J.M. JR. (1991) "On Buying Legislatures," *Economics and Politics* 3, 93-109.

STIGLER, G.J. (1971) "The Theory of Economic Regulation," *Bell Journal of Economics and Management Science* 2, 3-21.

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