Collective Commitment*

Christian Roessler Sandro Shelegia Bruno Strulovici

January 11, 2016

Abstract

Consider collective decisions made by agents with evolving preferences and political power. Faced with an inefficient equilibrium and an opportunity to commit to a policy, can the agents reach an agreement on such a policy? The answer is characterized by a consistency condition linking power structures in the dynamic setting and at the commitment stage: When the condition holds, the only agreement which may be reached, if any, coincides with the equilibrium without commitment. When it fails, as with time-inconsistent preferences, commitment may be valuable. We discuss applications and ways to facilitate the obtention of an agreement under power consistency.

JEL: D70, H41, C70

1 Introduction

In dynamic settings where information, preferences, and political influence evolve over time, successive decision-making by electorates, committees, or individuals often leads to suboptimal outcomes, such as the inability to implement needed reforms (Fernandez and Rodrik (1991)), the use of short-sighted monetary or fiscal policies (Kydland and Prescott (1977) and Battaglini and Coate (2008)), the stability of unpopular regimes (Acemoglu and Robinson (2005)), and the invocation of slippery slope arguments (Volokh (2003)). Voters' behavior reflects in part their desire to protect themselves against such developments: for example, proponents of a moderate reform may fear that it will set the stage for further reforms they would no longer endorse, and thus refuse to support any change in the first place. In these situations, given a chance to commit to a policy at the outset, it would seem that the equilibrium outcome could be improved upon by some commitment. In fact, that is the implicit premise of constitutions, laws, and other contracts that facilitate commitment. This paper studies formally when commitment can, and should, be used to address dynamic inefficiency.

*This paper subsumes earlier versions entitled "The Roman Metro Problem" and "Can Commitment Resolve Political Inertia?" We are grateful to Wiola Dziuda, Georgy Egorov, Jeff Ely, Daniel Garcia, Michael Greinecker, Karl Schlag, Stephen Schmidt, Joel Watson, and numerous seminar participants for their comments. Strulovici gratefully acknowledges financial support from an NSF CAREER Award (Grant No. 1151410) and a fellowship form the Alfred P. Sloan Foundation. Roessler: California State University, East Bay (e-mail: christian.roessler@csueastbay.edu); Shelegia: Universitat Pompeu Fabra and Barcelona GSE (e-mail: sandro.shelegia@upf.edu); Strulovici: Northwestern University (e-mail: b-strulovici@northwestern.edu).

To make the issue concrete, consider a legislature having to decide whether to pass a moderate reform, whose adoption may be followed by a more radical expansion. As noted, some voters in favor of the initial reform may oppose it nonetheless, worried that it may create a "slippery slope" leading to the radical reform if it gathers enough support from other voters. The resulting deadlock could seemingly be resolved by a commitment to implement only the initial reform and rule out any further one. As it turns out, however, such a commitment is majority-preferred to the status quo if and only if it is itself majority-dominated by the policy consisting of implementing the initial reform and then expanding it if the expansion turns out to be desired by a majority of voters. This policy is, in turn, dominated by the status quo, thus creating a Condorcet cycle among policies.

The situation is depicted on Figure 1. Voters are equally divided into three groups (A, B, C) with terminal payoffs as indicated. A majority decision to implement the initial reform (Y) reveals with probability q that an expansion is feasible. In this case, a vote takes place on whether to stop at the moderate reform (M) or implement the radical one (R). As shown in the appendix, for $q \in (2/3, 1]$, implementing the radical reform (YR) at the second stage is the unique equilibrium, while keeping the moderate reform (YM) is the unique equilibrium for $q \in [0, 1/3)$. Moreover, YM also beats the status quo (N) in majority voting and yields higher utilitarian welfare.

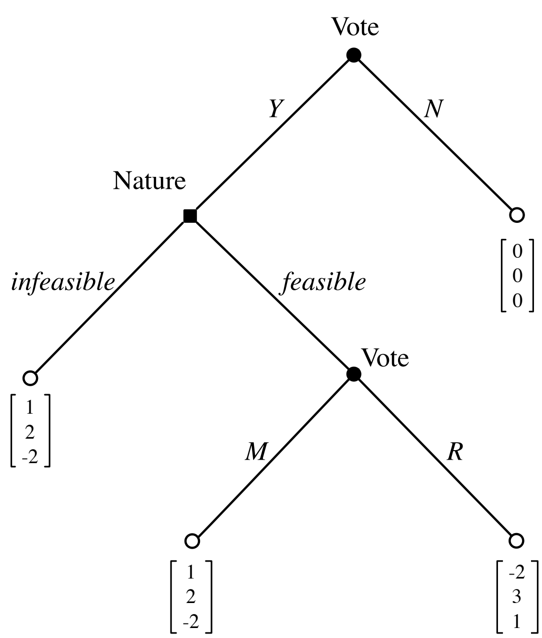

Figure 1: Slippery Slope. Circular nodes indicate majority-based decisions.

For $q \in (1/3, 2/3)$, however, the equilibrium outcome is the status quo, N, because voters from group A deem the risk of ending up with the radical reform too high, while voters from group C find the probability of getting the moderate reform (M), which they do not like, too high. To remedy this situation, suppose that A tries to persuade B to use their joint majority to commit to policy YM (implement and keep the initial reform, regardless of what is learned later). Both favor this proposal over the status quo. However, C may approach B with a counteroffer to instead commit to policy YR (implement the initial reform, but expand it if feasible). Both prefer this proposal to YM. A could then remind C that both of them are better off with the status quo (opposing the initial reform), since YR corresponds exactly to the off-equilibrium path that A and C reject in the dynamic game without commitment. These arguments describe a Condorcet cycle among policies: $YM \prec YR \prec N \prec YM$. Thus, allowing the legislature to commit to a state-contingent plan at the outset is unlikely to resolve the problem. It only leads to disagreements over the plan to follow and, in particular, will not rule out the status quo as a viable option.

When can collective commitment improve dynamic equilibria? As we shall see, the value of commitment depends on the way political power is assigned. Crucially, we must distinguish between influence on dynamic collective decisions in the absence of commitment and power over the ranking of state-contingent policies when commitment is being considered. To develop a preliminary understanding, we first consider the case in which all decisions, in the dynamic game and when voting to commit on a policy, are made according to the simple majority rule. The conclusion of the slippery-slope example then holds in full generality: given any payoffs, either the dynamic equilibrium is undominated or there is a Condorcet cycle with commitment, which involves both the policy corresponding to the dynamic equilibrium and the policy dominating it. As a result, if one applies the simple majority criterion to compare policies, the dominating policy is not ranked higher than the equilibrium itself.

Allowing for commitment under the simple majority rule thus replaces any problem of *inefficiency* with one of *indeterminacy*, even when commitment carries no administrative or other contractual costs and is perfectly credible. Commitment is either unnecessary, when the equilibrium is majority preferred to all other state-contingent policies, or it is impossible to agree on which commitment to choose.

How can this result be reconciled with the apparent value of commitment which is indicated, for example, by the prevalence of contracts? To clarify this issue, we proceed to consider general structures of political power. These may include the use of supermajority rules for some decisions and heterogenous allocations of power across agents. These details, it turns out, do not matter *per se* for the value of commitment.

Instead, what matters is a *power consistency* condition relating political power at the dynamic and commitment stages. When power consistency holds, the introduction of commitment suffers from the same problem as in the case of majority voting: whenever it is potentially valuable, it leads to a cycle among all commitment policies. Furthermore, the power consistency condition is necessary for this result: when it is violated, one may find a preference profile and a policy that dominates not just the equilibrium but also all other policies that are available with commitment.

We focus on dynamic settings where decisions at each period are binary and may be made according to arbitrary—possibly time-varying and state-dependent—voting rules.[1] Power consistency

[1]The focus on binary decisions eliminates "local" Condorcet cycles in each period and thus also an important potential source of indeterminacy which may confuse the main points of the paper. Thanks to this assumption, the cycles which may arise among state-contingent policies have nothing to do with possible cycles in any given period. They also result in equilibrium uniqueness, which simplifies the statements of the paper.

is defined by the following requirement: Consider two policies which are identical except for the decision made in a given period and for a given state (or subset of states) in this period. Then, the social ranking between these two policies must be determined by the same set of winning coalitions as the one arising in the dynamic game when that decision is reached. The condition thus rules out situations in which a subset of persons could impose one policy over another at the commitment stage, but would not be able to choose the action differentiating these policies in the dynamic game.

We explore in detail when one should expect power consistency to hold and when it is likely to be violated. For example, if an important decision requires unanimity in the dynamic setting, the simple majority rule should not be used at the commitment stage to compare policies differing only with respect to this decision. Here, power consistency reflects the notion that the importance of the decision is the same whether it is considered in the dynamic game or at the commitment stage. In other settings, power consistency captures a notion of fairness toward future generations. The condition prohibits current society members from committing to future actions which are contrary to the interest of future society members, who would normally be the ones deciding on these actions. Power consistency may also capture a notion of liberalism similar to the one described by Sen (1970): the social ranking of policies should respect the preferences of individuals who would naturally be making decisions in the dynamic setting.

Yet, violations of power consistency are reasonable in some contexts. In particular, it is well-known that commitment is valuable for a time-inconsistent agent. Time inconsistency creates a particular form of power inconsistency which favors the first-period self (or preference) of the agent. Similarly, current actors or generations may be able to lock in future decisions in various macroeconomic and political economy contexts we discuss in Section 8. In these cases, commitment has value.

Even when power consistency holds, it may be possible to circumvent the indeterminacy result by imposing some restrictions on the type of commitments which may be considered. For example, if some players are ex ante symmetric, it seems reasonable to focus on "anonymous" (i.e., non-discriminatory) policies, which treat these players identically, by giving each of them the same outcome distribution. To explore this idea, we develop a concept of "anonymous" policy which requires that similar agents be treated similarly—in a sense which we formalize—as well as another anonymity criterion based on a veil-of-ignorance argument (Section 7).[2] Both approaches can restore the value of commitment, as we illustrate, either by removing policies which appeared in the cycle or by modifying the individual criteria used to assess policies.

Finally, some forms of commitment, such as a unilateral commitment or a commitment to vote in a particular way on a future collective decision, may be built into the underlying game (one advantage of the generality of the model considered here is precisely to allow for this). The equilibrium of the augmented game may then become efficient thanks to these commitments, removing indeterminacy in accordance to our results.

The observation which motivates our inquiry—that commitments can lead to indeterminacy—

[2]Tabellini and Alesina (1990) use a similar argument to show that commitment to balanced budgets is valuable when agents do not know who will be in the position of power.

has been made before. In particular, as Boylan and McKelvey (1995), Boylan et al. (1996), and Jackson and Yariv (2014) show, when agents have heterogeneous discount factors, no agreement can be reached over consumption streams because no Condorcet winner exists in their settings. The absence of a Condorcet winner weakens the applicability and value of commitment, as in our paper. By contrast, Acemoglu et al. (2012) and Acemoglu et al. (2014) provide single-crossing conditions on agents' preferences under which the equilibrium is undominated and a dynamic median voter theorem applies.[3]

Unlike these earlier works, our result does not affirm or negate the existence of a Condorcet winner among policies. Rather, it provides a necessary and sufficient condition (power consistency) under which the institution of commitment fails to resolve equilibrium inefficiency, either because the equilibrium was a good policy to begin with, or because commitments lead to indeterminacy, absent any external criterion to restrict the class of state-contingent policies which may be considered as commitments.[4] Our result may play a cautionary role in settings where committing to *some* policy could improve upon the equilibrium. Kydland and Prescott (1977) already emphasized the value of commitment in macroeconomic settings, and the observation that a political equilibrium is dominated by some specific commitment has frequently been made (e.g., Strulovici (2010) and Dziuda and Loeper (2014)). These observations characterize equilibrium inefficiency in specific dynamic settings. However, our result shows that the option to commit is not necessarily a cure for inefficiency. In fact, unless there is a reason why the switch from dynamic decisions to comparing policies puts a different group in charge, or certain policies are a priori ruled out, commitment cannot help.

Of particular relevance to our result, the literature on agenda setting has pointed out long ago (e.g., Miller (1977)) that, if the winner of a sequence of binary majority votes over alternatives depends on the order in which alternatives are compared, then there is no Condorcet winner among these alternatives.[5] Our setting is different because i) we allow uncertainty: the state of the world (physical state, information, individual preferences) can evolve stochastically over time, and ii) we consider general decision protocols in which decision rules and individual power may vary over time and depend on past decisions and events. These extensions are relevant in numerous applications—risky reforms, search by committees, theory of clubs, to cite only a few—and require a more sophisticated analysis, as explained next.

Without uncertainty, any policy reduces to a single path in the dynamic game, and can be identified with its unique terminal node. Each policy then corresponds to an "alternative" in the agenda setting literature. With uncertainty, however, this relation breaks down because policies

[3] Appendix C.1 provides similar conditions for our setting.

[4] Imposing such an external criterion already constitutes a form of commitment, which may be acceptable if the criterion is used broadly beyond the particular game under consideration. In Section 7 we provide such a broadly applicable criterion.

[5] In a static choice problem, Zeckhauser (1969) and subsequently Shepsle (1970) study the existence of Condorcet winners in voting over certain alternatives and lotteries over them. Zeckhauser shows that, if all lotteries over certain alternatives are in the choice set, no Condorcet winner can be found, even if there is such a winner among certain alternatives. In a comment on Zeckhauser, Shepsle demonstrates that a lottery can be a Condorcet winner against certain alternatives that cycle.

are state-contingent plans which can no longer be identified with terminal nodes. Choosing among policies at the commitment stage is thus no longer equivalent to making a sequence of binary choices in the dynamic game. Notably, one may construct examples (Appendix B) in which reversing the order of moves in the dynamic game does not affect the equilibrium, and yet the equilibrium is Pareto dominated by some other policy.

While uncertainty makes the analysis more involved, it also makes it more relevant and interesting: many environments (some described in this paper) have the feature that agents learn about their preferences over time, which results in changes in political alliances. These potential changes affect the incentives of winning coalitions early in the game and distort equilibrium away from efficiency, making the commitment issue particularly salient in such environments.

Besides their focus on deterministic settings, earlier works have only considered fixed tournaments as a way to choose between alternatives (an overview is provided by Laslier (1997)). However, when decision rules can evolve over time in response to previous decisions or shocks, the comparison of policies cannot be identified with a static tournament structure. In the theory of clubs, for instance, an early decision to admit new members dilutes the power of preceding members and, hence, affects the subsequent comparisons of alternatives. Our analysis imposes no restriction concerning the dynamics of the power structure.

After describing our main result for the simple majority rule in Section 2 and the general case in Section 3, we discuss interpretations, and violations, of the power consistency condition in Section 4. Section 5 presents three applications. The first one concerns search committees and is illustrated by a job market example. The second application, based on Fernandez and Rodrik (1991), concerns reforms and emphasizes the possible role played by appropriate commitment restrictions to improve equilibrium outcomes. The third one, based on Besley and Coate (1998), concerns the political economy of redistribution and shows that our results are relevant in settings where political power shifts over time, even when there is no uncertainty. Generalizations of our model are considered in Section 6. Section 7 investigates two anonymity criteria which can be used to restore the value of commitment, even when the power consistency condition holds. We conclude with a discussion in Section 8 of the role of commitment in various literatures, demonstrating their relation to our main results and to the power consistency condition introduced in this paper. The appendix contains omitted proofs (Appendix A) and formal details on the difference between our result and agenda setting approaches (Appendix B). Appendix C.2 (online) reviews conditions for the existence of a Condorcet winner and shows how our ideas can be adapted to an infinite horizon.

2 Preliminaries: Simple Majority Rule

There are T periods and N (odd) voters.[6] Each period starts with a publicly observed state $\theta_t \in \Theta_t$, which contains all the relevant information about past decisions and events. At each t, a collective decision must be made from some binary set $A(\theta_t) = \{\underline{a}(\theta_t), \bar{a}(\theta_t)\}$. This choice, along

[6]For this section only we assume for simplicity that the N agents are alive throughout the game, and in particular vote in every state. We relax this assumption in the general setting of section 3.

with the current state, determines the distribution of the state at the next period. Formally, each Θ_t is associated with a sigma algebra Σ_t to form a measurable space, and θ_{t+1} has a distribution $F_{t+1}(\cdot \mid a_t, \theta_t) \in \Delta(\Theta_{t+1})$. If, for instance, the state θ_t represents a belief about some unknown state of the world, θ_{t+1} includes any new information accrued between periods t and $t+1$ about the state, which may depend on the action taken in period t. The state θ_t may also include a physical component, such as the current stage of construction in an infrastructure-investment problem.

Let $\Theta = \bigcup_{t=1}^{T} \Theta_t$ and $A = \bigcup_{\theta \in \Theta} A(\theta)$ denote the sets of all possible states and actions. Each voter i has a terminal payoff $u_i(\theta_{T+1})$, which depends on all past actions and shocks, as captured by the terminal state θ_{T+1}. A *policy* $C : \Theta \to A$ maps at each period t each state θ_t into an action in $A(\theta_t)$.

If a policy C is followed by the group, then given state θ_t, i's expected payoff seen from period t is[7]

$$V_t^i(C \mid \theta_t) = E[u_i(\theta_{T+1}) \mid \theta_t, C].$$

From here onwards, as is standard in the tournaments literature, for simplicity we shall require that no voter is indifferent between the two actions in $A(\theta_t)$ at any state θ_t.[8]

Given a policy C and state θ, let C_θ^a denote the policy equal to C everywhere except possibly at state θ, where it prescribes action $a \in A(\theta)$.

Definition 1 (Voting Equilibrium). *A profile $\{C^i\}_{i=1}^N$ of voting strategies forms a* Voting Equilibrium in Weakly Undominated Strategies *if and only if*

$$C^i(\theta_t) = \arg\max_{a \in A(\theta_t)} V_t^i(Z_{\theta_t}^a \mid \theta_t)$$

for all $\theta_t \in \Theta$, where Z is the policy generated by the voting profile:

$$Z(\theta_t) = a \in A(\theta_t) \text{ if and only if } |C^i(\theta_t) = a| \geq \frac{N}{2}.$$

Z is defined by simple majority voting: at each time, society picks the action that garners the most votes. The definition captures the elimination of weakly dominated strategies: at each t, voter i, taking as given the continuation of the collective decision process from period $t+1$ onwards that will result from state θ_{t+1}, votes for the action that maximizes his expected payoff as if he were pivotal.

Because, by assumption, indifference is ruled out and the horizon is finite, this defines a unique voting equilibrium, by backward induction. The proof of this fact is straightforward and omitted.

Proposition 1. *There exists a unique voting equilibrium.*

[7] Because the terminal state θ_{T+1} includes past states, this formulation includes the time-separable case where $u_i(\theta_{T+1}) = \sum_{t=1}^{T+1} u_{i,t}(\theta_t)$ for some period-utility functions $u_{i,t}$, as well as non-time-separable utility functions.

[8] The literature on tournaments assumes that preference relations across alternatives are asymmetric. See Laslier (1997). Without this strictness assumption, most of Theorem 1 still applies to "weak" Condorcet winner and cycle. See also Remark 1.

Given a pair (Y, Y') of policies, we say that Y dominates Y', written $Y \succ Y'$, if there is a majority of voters for whom $V_1^i(Y \mid \theta_1) > V_1^i(Y' \mid \theta_1)$. A Condorcet cycle is a finite list of policies $Y_0, \ldots, Y_K$ such that $Y_k \prec Y_{k+1}$ for all $k < K$, and $Y_K \prec Y_0$. Finally, X is a Condorcet winner if, for any Y, either $X \succ Y$ or X and Y induce the same distribution over Θ_{T+1}.

Theorem 1. *Let Z denote the equilibrium policy.*

i) If there exists Y such that $Y \succ Z$, then there is a Condorcet cycle including Y and Z.

ii) If there exists a policy X that is a Condorcet winner among all policies, then X and Z induce the same distribution over Θ_{T+1}.

Remark 1. *If voters' preferences allow ties, Part i) still holds with a weak Condorcet cycle: there is a finite list of policies $Y_0, \ldots, Y_K$ such that $Y_k \preceq Y_{k+1}$ for all $k < K$, and $Y_K \prec Y_0$. Furthermore, Z continues to be a Condorcet winner in the sense that there does not exist another policy Y such that $Z \prec Y$.*

The proof, in Appendix A.2, may be sketched as follows. If a policy Y differs from the equilibrium policy Z, then Y must necessarily prescribe, for some states reached with positive probability, actions which the majority opposes. Using this observation, we iteratively construct a sequence of policies by gradually changing Y in these states, in the direction of the majority's will, so that each subsequent policy is majority preferred to the previous one. Because the game is finite, this process eventually ends with the policy Z where all actions follow the majority's preference. More explicitly, we start with the last period, $\hat{T}$, for which Y differs from Z on some subset of states. We then create a new policy, Y_1, identical to Y except in some time-$\hat{T}$ state for which Y differs from Z. On these states, Y takes an action that is not supported by a majority, since Y and Z have the same continuation by definition of $\hat{T}$, and Z was the equilibrium policy. Moreover, Y_1 is now closer to Z as it takes the same actions as Z on the state over which the change took place. We then apply the procedure to another time-$\hat{T}$ state for which Y_1 (and thus Y) prescribes a different action from Z, creating a new policy Y_2, which is identical to Y_1 except for taking the majority preferred action in this state. By construction $Y_1 \prec Y_2$. Once all time-$\hat{T}$ states for which Y differs from Z have been exhausted by the procedure, we move to time $\hat{T} - 1$ and repeat the sequence of changes, constructing a chain of policies which are increasing in the majority ranking and getting gradually more similar to the equilibrium policy, Z. The process ends with a policy Y_K that coincides with Z. Because we know that Y is different from Z, $K \geq 2$, which creates a Condorcet cycle if and only if the initial policy Y dominated Z.[9]

[9] A technical complication, omitted above, is that individual states may have zero probability (e.g., if the state space at each step is a continuum with a continuous distribution). This issue is addressed by partitioning states, in each period, according to the winning coalitions which prefer Z's action over Y's prescription, and having each step of the above procedure simultaneously apply to all the states corresponding to some winning coalition. Because the set of such coalitions is finite, we can reconstruct Z from Y in finitely many steps.

The cycles predicted by Theorem 1, whenever they occur, may be interpreted as follows: If the population were allowed, before the dynamic game, to commit to a policy, it would be unable to reach a clear agreement, as any candidate would be upset by some other proposal. If one were to explicitly model such a commitment stage, the outcome of this stage would be subject to well-known agenda setting and manipulation problems, and the agenda could in fact be chosen so that the last commitment standing in that stage be majority defeated by the equilibrium of the dynamic game.[10]

Theorem 1 distinguishes two cases: when the equilibrium is undominated and when there is no Condorcet winner. These cases can often coexist in the same model, for different parameter values. This was the case in the slippery slope example, where the equilibrium is undominated for $q \in [0, 1/3] \bigcup [2/3, 1]$ and no Condorcet winner existed for $q \in (1/3, 2/3)$.

A more positive interpretation of Theorem 1 is that, even when the equilibrium policy is majority dominated by another policy, it must belong to the top cycle of the social preferences based on majority ranking.[11] In the agenda-setting literature, it is well-known that the equilibrium must belong to the Banks set (Laslier (1997)). This need not be the case here, however, due to the presence of uncertainty, because the dynamic game does not give voters enough choice to compare all policies: the decision set is just not rich enough. In particular, with T periods agents make only T comparisons throughout the dynamic game, but policies, being state-contingent plans, are much more numerous when the state is uncertain. As a result, the equilibrium does not *per se* inherit the Banks-set property.

Another way of understanding the difference between the alternatives compared in the agenda-setting literature and the policies compared in our framework is that a state-contingent policy now corresponds to a *probability distribution* over terminal nodes, and in the dynamic voting game agents do not have rich enough choices to express preferences amongst all these distributions. Put in the more formal language of tournaments, the choice process along the dynamic game may not be summarized by a complete algebraic expression for comparing all policies (Laslier (1997)). These differences are substantial and indeed, the method of proof used for establishing our main theorem is quite different and significantly more involved than the one used in deterministic setting to show that the equilibrium is dominated if and only if there is no Condorcet winner among simple alternatives.

3 General Voting Rules and Power Consistency

Collective decisions often deviate in essential ways from majority voting. In the slippery slope problem, for example, some decisions may be taken by a referendum and others by lawmakers. Another natural example concerns constitutional amendments in the United States, which require a supermajority rule. This section shows that our main result still holds for arbitrary decision rules,

[10] One could also incorporate commitment decisions into the dynamic game, with the state θ_t encoding whether a commitment has been chosen before period t (and if so, which one).

[11] Even then, however, the equilibrium policy may be Pareto dominated by another policy, as in the recruiting application described in Section 5.

under a *power consistency* condition whose meaning and relevance are discussed in detail below.

The formal environment is the same as before except for the structure of political power.[12] Given a period t and state θ_t, the "high" action $\bar{a}(\theta_t)$ might, for instance, require a particular quorum or the approval of specific voters (veto power) to win against $\underline{a}(\theta_t)$. The decision rule may also depend on the current state and, through it, on past decisions. In many realistic applications, some voters may be more influential than others because they are regarded as experts on the current issue, or because they have a greater stake in it, or simply because they have acquired more political power over time.

To each state θ_t corresponds a set $\bar{\mathcal{S}}(\theta_t)$ of coalitions which can impose $\bar{a}(\theta_t)$ in the sense that, if all individuals in $S \in \bar{\mathcal{S}}(\theta_t)$ support $\bar{a}(\theta_t)$, then $\bar{a}(\theta_t)$ wins against $\underline{a}(\theta_t)$ and is implemented in that period. Likewise, there is a set $\underline{\mathcal{S}}(\theta_t)$ of coalitions which may impose $\underline{a}(\theta_t)$. These sets are related as follows: $\underline{\mathcal{S}}(\theta_t)$ contains all coalitions whose complement does not belong to $\bar{\mathcal{S}}(\theta_t)$, and vice versa. We impose the following condition: for any coalitions $S \subset S'$ and state θ, $S \in \bar{\mathcal{S}}(\theta) \Rightarrow S' \in \bar{\mathcal{S}}(\theta)$. This monotonicity condition implies that it is a dominant strategy for each individual to support their preferred action, for any given state: they can never weaken the power of their preferred coalition by joining it.

A *coalitional strategy* C^i for individual i is, as before, a map from each state θ_t to an action in $A(\theta_t)$. It specifies which action i supports in each state. Given any profile $\mathbf{C} = (C^1, \ldots, C^N)$ of coalitional strategies and any state θ, there are two coalitions: those who prefer $\bar{a}(\theta)$ and those who prefer $\underline{a}(\theta)$, and one of them is a winning coalition: it can impose its preferred action.[13] Let $a(\mathbf{C}, \theta)$ denote this action.

Given a policy C and state θ_t, i's expected payoff seen from period t, is given by

$$V_t^i(C \mid \theta_t) = E[u_i(\theta_{T+1}) \mid \theta_t, C].$$

Definition 2 (Coalitional Equilibrium)**.** *A profile* $\{C^i\}_{i=1}^N$ *of coalitional strategies forms a* Coalitional Equilibrium in Weakly Undominated Strategies *if and only if*

$$C^i(\theta_t) = \arg\max_{a \in A(\theta_t)} V_t^i(Z_{\theta_t}^a \mid \theta_t)$$

for all $\theta_t \in \Theta$*, where* Z *is is the policy generated by the profile:* $Z(\theta_t) = a(\mathbf{C}, \theta_t)$*).*

The definition is the same as for majority voting, except that now the action that wins in each period is the one supported by the strongest coalition. We maintain the assumption of the previous section that each voter has, for any policy and state θ_t, a strict preference for one of the two actions in $A(\theta_t)$. Because indifference is ruled out and the horizon is finite, this defines a unique coalitional equilibrium, by backward induction (the proof is omitted).

[12]The number of voters need not be odd any more. We do maintain the assumption that decisions are binary in each period to avoid the complications arising from coalition formation with more choices and equilibrium multiplicity.

[13]That is, the coalition of individuals preferring $\bar{a}(\theta_t)$ belongs to $\bar{\mathcal{S}}(\theta)$ if and only its complement does not belong to $\underline{\mathcal{S}}(\theta)$.

Proposition 2. *There exists a unique coalitional equilibrium.*

Commitment and Indeterminacy

Now suppose that society members are given a chance to collectively commit to a policy instead of going through the sequence of choices in the dynamic game. When can they agree on a policy that dominates the equilibrium? We need to specify the structure of power at the commitment stage. Given a pair (Y, Y') of policies, say that S is a winning coalition for Y over Y' if $Y \succ Y'$ whenever all members of S support Y over Y' when the two policies are pitted against each other. A power structure specifies the set of winning coalitions for every pair of alternatives. Given a power structure and a profile of individual preferences over all policies, one can then construct the social preference relation, which describes the pairwise ranking of every two alternatives: $Y \succ Y'$ if and only if there is a winning coalition S for Y over Y' all of whose members prefer Y to Y'. Our assumptions guarantee that the preference relation is complete.[14]

Given the social preference relation $\succ$, say that a policy Y is a *Condorcet winner* if there is no other policy Y' strictly preferred over Y by a winning coalition. A *Condorcet cycle* is defined as in the previous section with the only difference that $\succ$ is used instead of the simple majority preference relation.[15]

Our main result relies on a consistency condition relating the power structures in the dynamic game and at the commitment stage.

Definition 3 (Power Consistency)**.** *Suppose that Y and Y' differ only on a set $\bar{\Theta}_t$ of states corresponding to some given period t and that S is a winning coalition imposing the action prescribed by Y over the one prescribed by Y' for all states in $\bar{\Theta}_t$. Then, S is also a winning coalition at the commitment stage, imposing Y over Y'.*

Although the power structure at the commitment stage must specify the set of winning coalitions for every pair of policies, the power consistency condition is only concerned with a much smaller subset of those pairs, namely the pairs for which the two policies are identical except on a subset of states in a single period.

Theorem 2. *Assume power consistency, and let Z denote the equilibrium of the coalitional game.*

i) If there exists Y such that $Y \succ Z$, then there is a Condorcet cycle including Y and Z.

ii) If some policy X is a Condorcet winner among all policies, then X and Z must induce the same distribution over Θ_{T+1}.

Proof. Fix any policy Y, let $\bar{\Theta}_T$ denote the set $\{\theta_T \in \Theta_T : \quad Z_T(\theta_T) \neq Y_T(\theta_T)\}$. For $\theta_T \in \bar{\Theta}_T$, let $\mathcal{S}_T(\theta_T)$ denote the coalition of individuals who prefer $Z_T(\theta_T)$ to $Y_T(\theta_T)$. Since Z is the coalition

[14] Although individuals have strict preferences across any two actions in the dynamic game, they will be indifferent between two policies that take exactly the same actions except on a set of states that is reached with zero probability under either policy. We view such policies as identical and say that they "coincide" with each other.

[15] These generalizations of majority-voting concepts to general tournaments are standard. See, e.g., Laslier (1997).

equilibrium policy, $\mathcal{S}_T(\theta_T)$ must be a winning coalition given state θ_T. Let $\mathcal{S}_T = \{\mathcal{S}_T(\theta_T) : \theta_T \in \bar{\Theta}_T\}$ denote the set of all such coalitions and p_T denote the (finite) cardinality of $\mathcal{S}_T$. We index coalitions in $\mathcal{S}_T$ arbitrarily from S_1 to S_{p_T}. For each $p \leq p_T$, let Θ_T^p denote the set of $\theta_T \in \bar{\Theta}_T$ for which the coalition of individuals who prefer $Z_T(\theta_T)$ to $Y_T(\theta_T)$ is equal to S_p and for which S_p is a winning coalition. By construction Θ_T^p is nonempty. Consider the sequence $\{Y_T^p\}_{p=1}^{p_T}$ of policies defined iteratively as follows.

- Y_T^1 is equal to Y for all states except on Θ_T^1, where it is equal to Z.
- For each $p \in \{2, \ldots, p_T\}$, Y_T^p is equal to Y_T^{p-1} for all states except on Θ_T^p, where it is equal to Z.

By construction, $Y_T^1 \succeq Y$ because the policies are the same except on a set of states where a winning coalition prefers Z (and, hence, Y_T^1) to Y, and, by power consistency, they can impose Y_T^1 over Y in the commitment stage. This is because for all states in Θ_T^1, S_1 is a winning coalition, it has to be a winning coalition when comparing Y_T^1 to Y. The winning coalition's preference is strict if and only if Θ_T^1 is reached with positive probability under policy Y.

Therefore, either Y and Y_T^1 coincide (i.e., take identical actions with probability 1), or $Y_T^1 \succ Y$. Similarly, $Y_T^p \succeq Y_T^{p-1}$ for all $p \leq p_T$, and $Y_T^p \succ Y_T^{p-1}$ if and only if $Y_T^p \neq Y_T^{p-1}$ with positive probability. This shows that

$$Y_T^{\bar{p}} \succeq \cdots \succeq Y_T^1 \succeq Y,$$

and at least one inequality is strict if and only if the set of states $\bar{\Theta}_T$ is reached with positive probability under Y. By construction, $Y_T^{p_T}$ coincides with Z on Θ_T: $Y_T^{p_T}(\theta_T) = Z(\theta_T)$ for all $\theta_T \in \Theta_T$.

Proceeding by backward induction, we extend this construction to all periods from $t = T-1$ to $t = 1$. For period t, let $\bar{\Theta}_t = \{\theta_t \in \Theta_t : \quad Z_t(\theta_t) \neq Y_t(\theta_t)\}$. For $\theta_t \in \bar{\Theta}_t$, let $\mathcal{S}_t(\theta_t)$ denote the coalition of individuals who prefer $Z_t(\theta_t)$ to $Y_t(\theta_t)$. Given the continuation policy Z from time $t+1$ onwards, $\mathcal{S}_t(\theta_t)$ is a winning coalition, since Z is the coalitional equilibrium. Also let $\mathcal{S}_t = \{\mathcal{S}_t(\theta_t) : \theta_t \in \bar{\Theta}_t\}$. Letting p_t denote the cardinality of $\mathcal{S}_t$, we index coalitions in $\mathcal{S}_t$ arbitrarily from S_1 to S_{p_t}. Let Θ_t^p denote the set of θ_t's in $\bar{\Theta}_t$ for which the coalition of individuals who prefer $Z_t(\theta_t)$ to $Y_t(\theta_t)$ is equal to S_p and for which this coalition wins. Θ_t^p is nonempty, by construction of S_p. Consider the sequence $\{Y_t^p\}_{p=1}^{p_t}$ of policies defined iteratively as follows, increasing p within each period t, and then decreasing t: for each t,

- For $p = 1$, Y_t^1 is equal to $Y_{t+1}^{p_{t+1}}$ for all states, except on Θ_t^1, where it is equal to Z.
- For each $p \in \{2, \ldots, p_t\}$, Y_T^p is equal to Y_T^{p-1} for all states, except on Θ_t^p, where it is equal to Z.

All the constructed policies have Z as their continuation from period $t+1$ onwards. By construction, $Y_t^{p+1} \succeq Y_t^p$ for all t, and $p < p_t$ and $Y_t^1 \succeq Y_{t+1}^{p_{t+1}}$ for all t. Moreover, the inequality is strict unless the set of states over which they differ is reached with zero probability.

By construction, the last policy $Y_1^{p_1}$ generated by this algorithm is equal to Z. Let $\{Y_k\}_{k=1}^K$, $K \geq 1$, denote the sequence of *distinct* policies obtained, starting from Y, by the previous construction.[16]

If $Y \neq Z$ with positive probability, then $K \geq 2$. Moreover,

$$Y = Y_1 \prec Y_2 \cdots \prec Y_K = Z.$$

Therefore, we get a voting cycle if $Z \prec Y$, which concludes the proof of part i).

Since Z can never be defeated without creating a cycle, we can characterize a Condorcet winner over all policies, if it (they) exist(s), and ii) follows.

□

Theorem 2, implies that, if pairwise comparisons of policies are based on the same power structure as the one used in the binary decisions of the dynamic game, allowing commitment will not lead to an unambiguous improvement of the political equilibrium. While some agenda setter may propose a commitment to resolve political inertia, such a commitment can be defeated by another commitment proposal, and so on, getting us back to political inertia. While one may find some solace in the fact that the equilibrium policy is part of the top cycle among policies, it may of course be Pareto dominated by another policy, and one can choose payoffs to make the domination arbitrarily large.

Remark 2. *As with Theorem 1, a modification of Theorem 2 based on weak Condorcet cycles and weak Condorcet winners holds when agents are allowed to have weak, instead of strict, preferences.*

The model of this section, by allowing history-dependent power structures, extends the agenda-setting and tournament literatures, which have assumed (see Laslier (1997) for an overview) that the pairwise ranking of "alternatives" was prescribed by a single binary complete, asymmetric relation (tournament), regardless of how or when these alternatives were compared. In dynamic settings such as ours, where each decision affects the balance of power for future decisions, this invariance assumption is typically violated. In the theory of clubs, for instance, an early decision to admit new members dilutes the power of preceding members and, hence, affects the subsequent comparisons of alternatives.

The Necessity of Power Consistency

When power consistency fails, one may find some policies which are unambiguously preferred to the equilibrium. More precisely, we will say that the power structures used in the dynamic and commitment stages are *inconsistent* if there exist policies Y and Y' and a coalition S such that i) Y and Y' are identical, except for a subset $\bar{\Theta}_t$ of states of some given period t, reached with positive probability under policy Y (and hence Y'), ii) whenever a state $\theta_t \in \bar{\Theta}_t$ is reached in the dynamic game, S is a winning coalition imposing the action prescribed by Y' over the one prescribed by Y,

[16] We call two policies *distinct* if they induce different distributions over Θ_{T+1}. Policies that differ only at states that are never reached are not distinct.

iii) at the commitment stage, S does not belong to the set of winning coalitions imposing Y' over Y.

Theorem 3. *Suppose that the power structures are inconsistent across stages. Then, there exist utility functions* $\{u_i(\theta_{T+1})\}_{i\in\{1,\ldots,N\},\theta_{t+1}\in\Theta_{T+1}}$ *and a policy* X *such that the equilibrium* Z *is strictly dominated by* X *and* X *is a Condorcet winner.*

4 Interpreting Power Consistency

When does power consistency hold?

The simplest instance of our setting is when the same set of agents is making decisions at the dynamic and commitment stages, and these agents are time consistent. In this case, power consistency may be interpreted and justified in the following ways.

Expertise: Some decisions (choosing an energy policy, addressing international conflicts, setting monetary policy, etc.) require specific expertise. For these decisions, the power should lie with experts, both when these decisions are made in the dynamic game and when comparing policies which differ only with respect to these decisions.

Liberalism: Some decisions primarily concern specific subgroups of the population (e.g., city or statewide decisions, rules governing some associations, etc.). It seems natural to let these groups have a larger say over these decisions both at the dynamic and the commitment stages. This consideration is related to Sen's notion of "liberalism" (Sen (1970)), a link explored further in this section. It may also be applied to minority rights.

Supermajority: Many constituencies require a supermajority rule to make radical changes to their governing statutes. For example, amendments to the United States constitution require two-thirds of votes in Congress, and substantive resolutions by the United Nations Security Council require unanimity. The rules should treat these radical changes consistently, whether they are part of a commitment or arise in the dynamic game.

In several policy applications, such as problems with intergenerational transfers of resources, environmental decisions, and international treaties, commitments involve generations which are unborn when the commitments are made. Whether power consistency holds depends on how one treats unborn generations in practice.

Intergenerational altruism/liberalism: When a decision primarily concerns unborn generations, the social preference concerning policies that differ only with respect to this decision may, normatively, take into account the preferences of these generations – which may depend on the future state –

even though they are absent at the time of commitment. Today's generation is then guided by intergenerational altruism when considering commitments.

Departing generations: Conversely, some agents may die or leave the dynamic game following some actions or exogenous shocks. It is then reasonable to ignore them when comparing policies that differ only with respect to decisions arising after they left the game, which is captured by power consistency.

When is power consistency violated?

At the extreme opposite, another view of future generations is to simply ignore them in the social ranking of policies. This approach violates power consistency, and the current generation will typically find commitment valuable in this case.

Myopic/selfish generation: The current generation ignores the welfare and preferences of future generations. Power consistency is then violated, and this is exposed when the preferences of future generations are in conflict with those of the commitment-making generation.

Time inconsistency: Selfish generations capture a broader time inconsistency problem: the preferences of future decision makers are not reflected in today's preferences. The existence of a relationship between inconsistency and the value of commitment should not be surprising if one considers the case of time-inconsistent agents. Time-inconsistent agents violate power consistency because their initial ranking of social alternatives is not representative of their preferences when they make future decisions. One may think of a time-inconsistent agent as a succession of different selves, or agents, each with their specific preferences. At time t, the t-self of the agent is in power; he is the dictator and the unique winning coalition. When considering commitment at time 0, however, only the initial preferences of the agent are used to rank policies, which violates power consistency.

Commitment is deemed valuable in this case, but only because it is assessed from the perspective of the first-period agent. If one were to take the agent's preferences at various points in time into account, the value of commitment would be subject to the indeterminacy pointed out in Theorem 2.[17]

These observations extend to multiple agents. For example, a set of perfectly identical but time inconsistent agents would obviously face the same issues as a single time-inconsistent agent, regardless of the voting rule adopted in each period. Again, power consistency is violated if future selves have different preferences and their choices are not respected at time zero.

A similar source of time inconsistency concerns institutions whose government changes over time, bringing along different preferences. If an incumbent government can commit to a long-term policy which ties future governments' hands, it will typically find such a commitment valuable, and

[17] The agent's preferences in the first period may incorporate his future preferences, and this very fact may be the source of the agent's time inconsistency, as in Galperti and Strulovici (2014). However, agent's future preferences do not *directly* affect his ranking of policies at time 1.

this commitment may increase overall efficiency. In Tabellini and Alesina (1990) and Alesina and Tabellini (1990), for instance, governments alternate because political power shifts over time (e.g., voting rights are gained by some minorities), changing the identity of the median voter, even though each voter taken individually has a time-consistent preference. The incumbent government borrows too much relative to the social optimum because it disagrees with how future governments will spend the remaining budget. When future governments cannot affect the choice of a commitment policy, the power consistency condition fails. When they can, our theorem has a bite, and a cycle arises among commitment policies. One way out of this cycle is to put all agents behind a veil of ignorance, as suggested by Tabellini and Alesina (1990). We explore this possibility in detail in Section 7.

Law of the current strongest: Another form of power inconsistency arises when some agents become more politically powerful over time. Their influence on future decisions in the dynamic game extends above and beyond their power at the commitment stage. These power changes may be foreseeable or random, depending on the economic or political fortunes of individuals at time zero. Regardless of the cause, commitment may be valuable as a way to insulate future decisions from the excessive power gained by a small minority. Power consistency is violated because the evolution of individual power is not included in the commitment decision.

Choosing future voting rules

In some applications (Barbera et al. (2001), Barbera and Jackson (2004)), earlier decisions determine the voting rule used for ulterior decisions. More generally, early decisions can affect each agent's voting weight for future decisions. This possibility is allowed by our framework because the state θ_t includes any past decision and determines the set of winning coalitions at time t. Settings where the future allocation of political power is determined by current agents appear in the theory of clubs (Roberts (1999)) or in mayoral elections (Glaeser and Shleifer (2005)). Barbera et al. (2001) consider voters deciding on immigration policies that would expand their ranks, while Barbera and Jackson (2004) study the general problem of voters deciding today on voting rules that will be used in the future.

We now discuss in the context of an example whether power consistency should be expected to hold and what Theorem 2 means when power is endogenous. We start with a two-period model. In period 1, a first generation of voters, assumed for now to be homogeneous, chooses the voting rule for period 2, between simple majority and two-thirds majority. In period 2, the next generation votes on whether to implement a reform. It is assumed that a fraction $x \in [1/2, 2/3)$ of period-2 voters favors the reform. In this case, the period-1 generation can obtain whichever outcome it prefers for period 2, by choosing the voting rule appropriately. Whether power consistency holds is irrelevant, because period 2 voters really have no control over the outcome as they are split in their preferences and bound by the voting rule chosen by their elders. In particular, one may assume that the condition holds so that the conclusions of Theorem 2 apply. Here, the equilibrium is efficient for the first generation and dominates any other policy from their perspective, so we are in the case where a Condorcet winner exists and coincides with the equilibrium, as predicted by the theorem.

Suppose next that there is a third period, and that the voting rule chosen by the first generation must also be used for the period-3 decision, with x taking the same value as in period 2. To make the problem interesting, we assume that the first generation wishes to implement the reform in period 2 but not in period 3. In this case, choosing a voting rule in period 1 cannot provide an efficient outcome from the first generation's perspective, and committing to a long-term policy clearly increases that generation's utility. Power consistency is violated because the third generation's power to choose the reform in the third period is not reflected in the social comparisons of policies, which are exclusively based on the first generation's preferences.

Finally, suppose that the three generations are in fact made up of the same individuals at different times. There is a fraction x of people who prefer the reform in the second period and the same fraction x of (partially different) people who support it in the third period. Also suppose that the first-period choice, deciding on which voting rule to use in later periods, is made according to the simple majority rule. If in equilibrium the first-period decision is to use the simple majority rule for future periods, the reform is adopted in both periods. If instead the two-thirds majority rule is chosen in the first period, no reform is adopted in later periods. Suppose that the two-thirds majority rule is chosen in equilibrium. This means that there is a majority of individuals who dislike the reform in at least one period, so much so that they prefer the status quo to having the reform in both periods, even though there is also a majority (x) of people who, in each period, prefer the reform to be implemented in that period.[18] If we use the simple majority rule when comparing any pair of policies other than the pairs differing only at one period, there is a cycle across policies: a majority of people prefer no reform at all (Z) to both reforms (Y), but a majority prefers reform in period 1 only (X) to Z, and a majority prefers reform in both periods (Y) to X, so that $Y \succ X \succ Z \succ Y$. Power consistency seems reasonable in this setting: whatever decision is made in the dynamic game reflects the preferences of the population at the beginning of the game. The theorem applies and, since the equilibrium is dominated by reform in either period, we get a Condorcet cycle.

Power consistency and liberalism

Sen (1970) has demonstrated that a social ranking rule cannot be both Pareto efficient and satisfy what Sen calls "Minimal Liberalism": for at least two individuals there exist two pairs of alternatives, one for each individual, such that the individual dictates the social ranking between the alternatives in his pair. By linking social preferences to individual decisions in a dynamic game, power consistency can capture Sen's notion of liberalism as a particular case.

Sen's setting concerns a static social choice problem, in which an "alternative" entails a complete description of all decisions in society. When these decisions (collective or individual) can be

[18]For example, suppose that $x = 3/5$ and the $2/5$ who oppose the reform in any given period dislike it much more than they value the reform in the other period. By taking the sets of reform opponents to be completely disjoint across periods, we get $4/5$ of agents against simple majority rule in period 1, as it would cause the reform to be implemented in both periods.

represented as a dynamic game, Sen's alternatives correspond to the policies studied here, and there are natural settings in which power consistency corresponds to liberalism.

To illustrate, consider Sen's main example which concerns two individuals, 1 (a 'pervert') and 2 (a 'prude'), and a book, *Lady Chatterley's Lover*. The prude does not want anyone to read the book but, should the book be read by someone (for simplicity, Sen does not allow both individuals to read the book), she prefers to be the one reading it. The pervert, by contrast, would like someone to read the book, and would also prefer the prude to read it rather than himself (the rationale being that he enjoys the idea of the prude having to read this subversive book). Let x, y, and z respectively denote the following alternatives: 2 reads the book; 1 reads the book; no one reads the book. The situation is captured in the game represented in Figure 2: 1 first decides whether to read the book, then 2 makes the same choice if 1 elected not to read the book.[19]

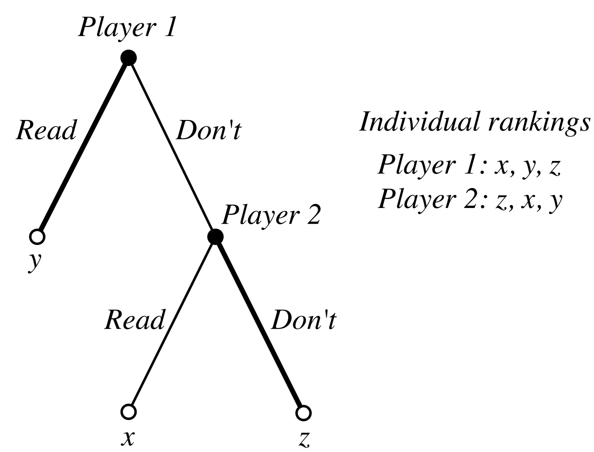

Figure 2: A representation of Sen's game. Individual preferences are indicated from the most preferred to the least preferred alternative.

Power consistency implies that player 2 has the right to choose between reading the book or not. Player 1, too, is entitled to reading the book, regardless of what player 2 does. Thus, power consistency and Sen's version of liberalism are equivalent in this setting. In the coalitional equilibrium of this game, the pervert reads the book and the prude does not (y). Moreover, the Pareto condition of Sen's analysis may be translated into our setting by requiring unanimity for x to win against y. Because x Pareto dominates y given the players' individual preferences, the equilibrium y is defeated by the commitment to a policy in which the pervert does not read the book and the prude does (x). Theorem 2 implies the existence of a Condorcet cycle, which recovers Sen's result on the impossibility of a Paretian liberal.

[19]The reverse sequence of moves yields outcome x (the prude player reads the book) and thus does not capture the tension at the heart of Sen's theorem.

5 Applications

Search Committees

Consider an economics department deciding whether to fly out a particular job candidate. The faculty initially doesn't know whether the candidate's primary interest lies in macroeconomics or labor economics, but this uncertainty will be resolved during the flyout, should it take place. If the faculty choose not to fly out the candidate, they will settle on a previously seen candidate. Similar models appear in the literature on search committees (Compte and Jehiel (2010), Moldovanu and Shi (2013)). The new candidate is risky because she may polarize the committee. Can commitment help improve the committee's decision?

Let N denote the new candidate and S denote the previously seen, "status quo" candidate. If the flyout takes place, N's field becomes known, and the faculty votes between N and S. The department is divided into three equally sized groups with the following preferences. To a third of the faculty (group *I*), it is important to hire a candidate who will exclusively work on macroeconomics; these members would prefer to make an offer to S over a labor economist. Another third (group *II*) is already convinced about N's value and is willing to choose her over S regardless of her field. The remaining third (group *III*) wishes to choose N over S only if she is a labor economist. Figure 3 represents the payoffs of this game. The utility provided by S is normalized to zero. Adding a twist to the game, we assume that the flyout is intrinsically desirable: even if the department ends up hiring S, it receives a higher payoff from having flown N out, perhaps because getting to know this new colleague and learning about her work is valuable regardless of the hiring decision. Furthermore, there is always a majority ex post who would support making the offer to N.

Despite the benefits of N's flyout, groups *I* and *III* block it for opposite reasons: *I* is concerned that N turns out to be a labor economist, in which case *II* and *III* will join forces to make her an offer, while *III* worries that N turns out to be a macroeconomist, in which case *I* and *II* will impose her on the department. Hence, there is no ex ante majority in the department to fly N out. This choice is Pareto dominated by the decision to fly N out and then hiring S. In these circumstances, commitment is a tempting solution. For instance, commitment to hire N only if she is a labor economist is majority preferred to the equilibrium. Theorem 1 then implies that there must be a Condorcet cycle.[20]

While decision makers have no clear way out of an inefficient choice, candidate N could fix the problem by revealing her primary interest. By positioning herself clearly as a macroeconomist or a labor economist, she can remove any uncertainty and guarantee herself a flyout. This observation supports the commonly heard advice that job candidates should avoid mixing fields.

[20] One cycle is this: the commitment to fly N out but hiring S in any case defeats not flying N out, but this commitment is beaten by making an offer to N only if she is a labor economist, which is preferred by *II* and *III*. This policy is itself majority-dominated by hiring N in any case, since this yields better outcomes for *I* and *II*. And this last policy is defeated by not flying N out.

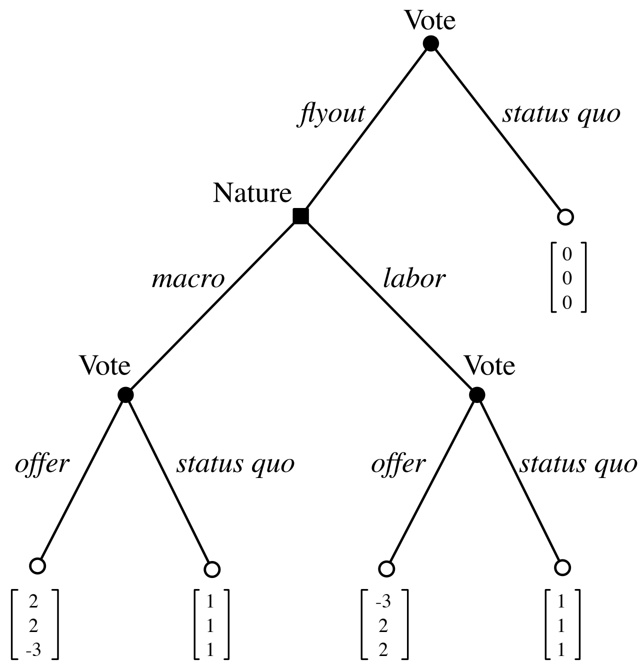

Figure 3: Job Market Game

Reforms

A common source of political inertia concerns the avoidance of socially valuable reforms that carry uncertain outcomes and, hence, political risk (Fernandez and Rodrik (1991)). Theorem 1 suggests that the institution of commitment may fail to resolve political inertia. Nonetheless, restricting feasible commitments to anonymous policies provides a natural way of avoiding Condorcet cycles in this application, when agents are symmetric ex ante. We develop this idea in Section 7.

We build on the two-stage setting of Fernandez and Rodrik (1991). In the first stage, citizens of a country decide whether to institute a trade reform. If the reform is undertaken, each individual learns whether he is a winner or loser of the reform. In the second stage, citizens vote on whether to continue the reform, or to implement it if they hadn't done so in the first period. The game is represented on Figure 4. The reform imposes a (sunk) cost c on each individual that must be borne once, regardless of the duration of the reform. Voters are divided into three groups (*I*, *II*, and *III*) of equal size, and one of the groups is randomly (with uniform probability) chosen as the sole winner from the reform. Individuals in the winning group get a payoff of g per period for the duration of the reform, while remaining individuals lose l per period. If the reform is implemented in the first period and continued in the second, we call it a long-term reform, whereas if it is revoked in the second period, it is a short-term reform.

Provided that g is sufficiently larger than l, the long-term reform is socially valuable. It provides a higher expected payoff to everyone relative to the status quo (not implementing the reform in either period). Commitment to the long-term reform is thus majority preferred to the status quo. However, any initial reform must be revoked in the second period, because two out of three groups

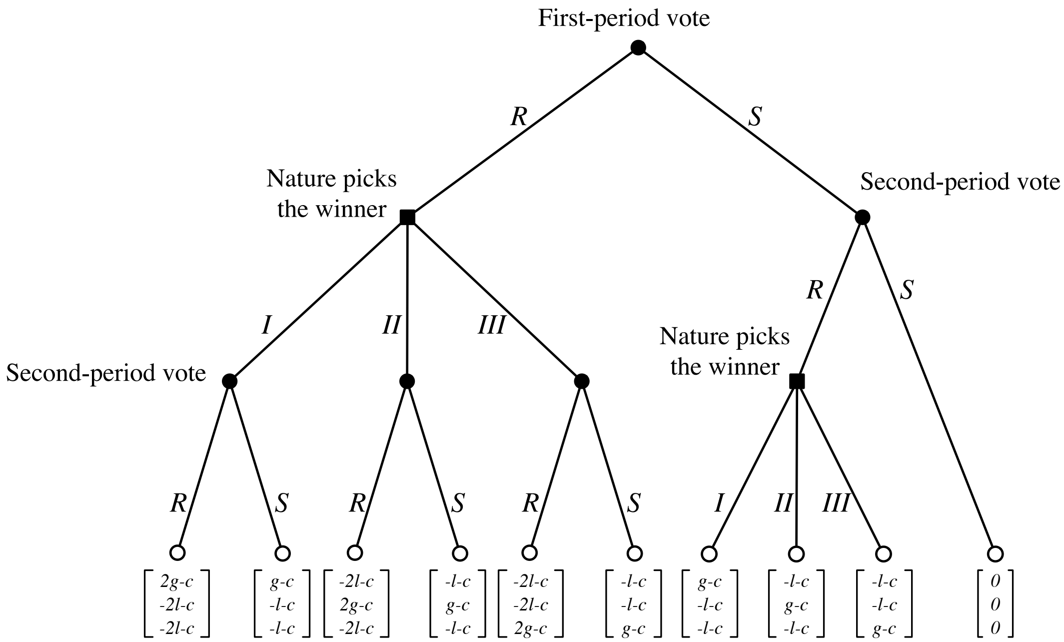

Figure 4: Reform game from Fernandez and Rodrik (1991) for three voters and one winner from the reform.

find out that they are losers of the reform and have an incentive to end it in the second period. A status quo bias arises if the reform is not implemented at all in equilibrium even though committing to it would be socially beneficial.[21]

Theorem 1 implies that, whenever the status quo bias arises, there must exist a Condorcet cycle over policies, and this cycle involves both the status quo and long-term experimentation. The status quo occurs in equilibrium when $g < 2l + 2c$, as in this case the expected payoff from the short-term reform (the reform is ended after the first period, because two groups of losers are identified) is negative for every type. The expected payoff from the long-term reform is positive for each type, provided that $g > 2l + 1.5c$. Other possible commitments, to short-term reform or to delayed reform in the second period also yield negative expected payoffs. Among these policies, the *long-term reform is a Condorcet winner*, which seems to contradict Theorem 1.

This paradox is explained by the consideration of other possible policies. For example, the commitment to a long-term reform is majority-dominated by the policy which consists of implementing the reform in the first period, and then revoking it only if group I is the winner. Groups II and III strictly prefer this policy to unconditional long-term reform. In turn, this policy is majority dominated by the commitment to continue the reform unless I or II is the winner. This differs from the previous policy only when II is the winner, and in that state of the world I and III benefit from ending the reform, so their expected payoffs increase. Now that I and II are at best tempo-

[21] "Status quo bias" is defined as a situation where a socially beneficial reform is not implemented despite having a positive expected payoff relative to the status quo. Strulovici (2010) decomposes the notion of status quo bias into "loser trap" and "winner frustration" effects, which respectively mean being stuck with a reform that turns out to be detrimental to oneself, and being unable to implement in the long run a reform which turned out to be profitable for oneself.

rary winners, both prefer the status quo (recall that the expected payoff from short-term reform is negative), which yields a cycle.

These policies, which single out some groups, seem perhaps unfair. In Theorem 1, the Condorcet cycle over policies is constructed over the full set of feasible policies: any plan of action that conditions on states where voting takes place is under consideration. In this application, however, it makes sense to restrict commitment to anonymous policies, and there does exist a Condorcet winner among anonymous policies. This suggests a way of circumventing the negative results of Theorems 1 and 2 by restricting the policy space.

Redistribution

It is commonly observed in the political economy literature that commitment is valuable for the following reason: welfare-improving decisions are disregarded because the current government expects that such decisions would adversely affect its future political power. For instance, in Besley and Coate (1998), a public investment that increases citizens' ability is not undertaken because it would lead to a change in preferences for redistribution. A commitment restricting future governments' decisions could a priori resolve this inefficiency. As illustrated below, however, the logic underlying our results suggests otherwise: allowing commitment need not yield a clear improvement over the equilibrium policy.

We consider a simplified version of Example 2 in Besley and Coate (1998), where for convenience we have eliminated ties. There are two periods in which citizens inelastically supply a unit of labor to the market and vote before each period on a linear redistribution scheme, (t, T), where t is the tax rate on labor income and T is the lump-sum redistribution obtained from the proceeds of the proportional tax (the budget is balanced in each period). A citizen with productivity a supplies one unit of labor, earns a, and receives utility $u = a(1-t) + T$. There are 3 types of citizens, assumed to be in equal numbers: "Low" types and "High" types have productivity a_L and a_H in both periods, respectively, while "Movers" have a productivity level a_L in the first period which increases to $a_L + \delta (< a_H)$ in the second period if a reform is implemented. In the first period, citizens decide on the tax scheme for this period and on whether to implement the reform. This decision is captured by the tuple (t_1, T_1, I), where $I = 1$ ($I = 0$) means that the reform is (not) implemented. Implementing the reform is costless. In the second period, citizens vote on (t_2, T_2).

Besley and Coate model the political process as follows: Before each period, a citizen of each type decides—at virtually no cost—whether to run for office. Once entry decisions are made, citizens vote for one of the candidates, and the winning candidate implements his optimal policy. Candidates cannot make binding promises, and voters rationally anticipate the consequences of their vote. For example, if a High type wins in the second period, she will set taxes to zero because she gains nothing from any redistribution. By contrast, a Low type fully redistributes earnings by setting $t_2 = 1$.

Besley and Coate show that even though the reform costs nothing and increases productivity, there are parameters for which the equilibrium entails no reform. Intuitively, Movers may prefer to side with the Low types to tax and fully redistribute the proceeds in both periods without the reform, rather than side with the High types who would support the reform but would prevent any redistribution. Besley and Coate suggest that, had commitment been available, the reform would have been undertaken. Our results suggest, instead, that there should be a cycle between commitments.

To see this clearly, we bypass Besley and Coate's running-decision stage and simply assume that exactly one candidate of each type runs for office. In each period, there is a closed primary between the Low and Mover candidates—whose base have initially identical skills; the winner is then pitted against the High-type candidate in a general election. We set parameters to $a_L = 0$, $a_H = 30$, and $\delta = 20$ and assume that voters maximize their expected payoff, eliminating weakly dominated strategies, and that the majority rule is applied to the relevant electorate in each stage. Finally, we assume that in case of a tie in the primary, the Low-type candidate wins.[22] In equilibrium, the Low-type candidate wins the primary and the general election. The Low and Mover types then vote for the Low-type candidate in the general election, who taxes at 100% and does not implement the reform in the first period. This equilibrium policy, X_1, yields a payoff of 20 to everyone, assuming no discounting across periods: each group gets a third of the High type's output. Movers do not vote for the High type because he would not redistribute, giving them a lower utility: the policy which rules out redistribution and implements the reform, X_2, yields payoffs 0,18,60 for the Low, Movers, and High types. In equilibrium, the Low-type candidate does not want to reform because she knows that it would prompt Movers to vote against them and, hence, redistribution, in the second period. This policy, in which redistribution takes place only in the first period and the reform is implemented, denoted X_3, yields payoffs 10,28,40. It achieves the optimal payoff for Movers. X_3 is majority preferred to the status quo (X_1). However, if redistribution is going to take place in a single period, and the reform is implemented, then clearly the High and Low types would prefer it to take place in the second period, in which Movers will contribute their earnings to the redistribution. The resulting policy, X_4, yields payoffs $16, 16, 46$. Notice however that X_4 is itself majority dominated by the status quo, X_1, resulting in the cycle predicted by our theorem.[23]

It seems intuitive, as suggested by Besley and Coate, that citizens would like to commit to a Pareto-improving policy in which the reform is undertaken: this reform increases global output at no cost. In particular, the policy X_5 that implements the same tax rate as the equilibrium (i.e., full redistribution) but implements the reform, clearly Pareto dominates the equilibrium by increasing the pie shared by all citizens. However, this policy, which yields 26 to each citizen, is majority dominated by policy X_3 and is thus part of a cycle.[24]

[22]Equivalently, one could assume that the Low-types slightly outnumber Movers. Keeping an equal number of each type simplifies the computation.

[23]To fit the exact structure of our theorem, notice that each of the four policies above corresponds to an actual policy which would be implemented for some sequences of winning types across the two periods: X_1 is implemented if the Low type wins both elections, X_2 is implemented if the High type wins both, X_3 is implemented if Movers win both elections, and X_4 is implemented if a High type wins the first election and a Low type wins the second election.

[24]Policy X_5 is implemented if Movers win the first election and Low types win the second one.

6 Extensions

Random proposers

In well-known agenda-setting protocols, voters may take turns to make collective proposals, and may be chosen deterministically or stochastically to do so. These protocols are compatible with the setting of this paper. For example, for t odd the state θ_t would include, as well as past information, the identity of a proposer who chooses between two collective proposals. At the next, even, period, the new state θ_{t+1} includes the proposal just made and society decides whether to accept the proposal, given the possibility of future proposals.

Non-binary decisions

Our earlier focus on binary decisions in the dynamic game gets rid of Condorcet cycles at the stage-game level, which avoided confusion between these cycles and those, at the heart of our result, which may arise among commitment policies.[25] Many political problems do, in fact, have this binary structure. For example, choices such as referenda and initiatives take the form of binary decisions. Similarly, lawmakers introduce bills and amendments as "yes" or "no" choices.

With three or more alternatives to choose from at any time, one may attempt to resolve the potential Condorcet cycles by a "binarizing" procedure as in the agenda setting literature. The resulting game then becomes subject to the theorem of this paper: the binary choice sequence leads either to an undominated equilibrium or to an indeterminacy in the ranking of state-contingent policies.

Transfers

Transfers may be explicitly considered in the present setting. For instance, one may include periods at which the binary action corresponds to whether some player i makes a specific transfer to another player j. In this case, i is a 'dictator' over the decision, and the state θ_t keeps track of all past transfers entering players' payoffs at the end of the game.

General utility functions: non separability and past dependence

The utility functions considered in this paper may depend arbitrarily on past states and decisions. For example, they allow decision complementarities across periods and all forms of path-dependence, such as habit formation, addiction, taste for diversity, utility from memories, learning

[25] The approach is also used in the explicit protocol proposed by Acemoglu et al. (2012) (p. 1458). While collective decisions are all binary, that paper allows a player to make proposals among all possible states. This can be easily replicated here by a sequence of at most S periods, where S is the number of states, with each period corresponding to a new state being presented to the proposer, who makes a "no" decision until being presented with the state that he wants to propose to the group, at which point he votes "yes" and the proposal is made to the group.

by experimentation, learning by doing, etc. With such features, one must be careful to correctly interpret power consistency.

7 Commitment and Anonymity

We suggested in Section 5 the use of anonymity as a criterion to restore the value of commitment even when power consistency holds. This section develops two approaches to achieve this. The first approach introduces the concept of an "anonymous" (i.e., non-discriminatory) policy, which is used to eliminate policies violating the definition from consideration at the commitment stage. The approach may be used to destroy some of the Condorcet cycles predicted by our theorems, by removing non-anonymous policies which were entering these cycles. The second approach considers all policies, but compares them behind a "veil of ignorance" (Harsanyi (1955); Atkinson (1970); Rawls (1971)). Identifying which policies are anonymous and evaluating policies behind the veil of ignorance are conceptually demanding tasks in general dynamic games. We thus begin this section with a simple example to illustrate the issues involved.

Consider the following game (see Figure 5) with three agents, 1, 2 and 3. Nature chooses one of the agents with equal probability. If either 2 or 3 is chosen, a majority vote determines whether the chosen agent ends up in "condition" U or D. Here, a condition describes the material contingencies or "physical outcomes" faced by the agent. Identical labels indicate identical situations and, for simplicity, we assume here that no other agent is affected by the vote. If 1 was chosen, a majority vote determines whether she ends up in condition L or R. These conditions are distinct from U and D. For reasons that will become apparent, we do not yet specify agents' utility at terminal nodes, focusing for now on their physical outcomes.

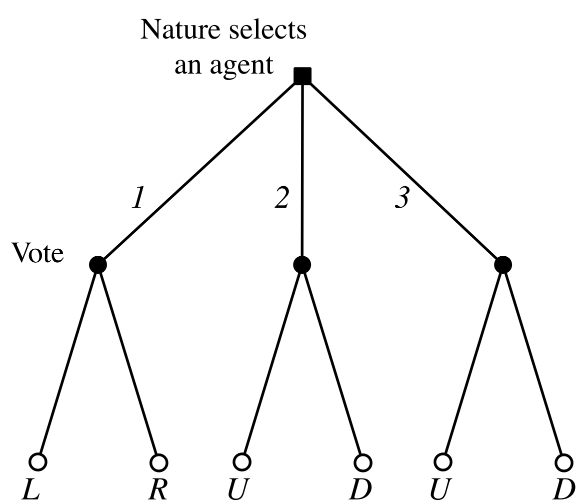

Figure 5

Intuitively, any policy that treats agents 2 and 3 differently must fail anonymity, because these agents are completely symmetric. By contrast, agent 1 is different from the other two because the conditions that she can end up in (L and R) are different from those available for the other two agents (U and D), so even anonymous policies must treat 1 differently from the rest.[26]

We identify anonymous policies by grouping agents into homogeneous categories (in the previous game, there are two such categories: one for agent 1 and the other for agents 2 and 3), and then checking whether a policy treats all agents in the same category equally.

As noted, our definition of anonymity concerns only agents' conditions, not their preferences: it does not aim to provide the same utility to agents in similar physical conditions, merely to provide them with the same material condition. This separation is formalized by distinguishing positions in the game, and preferences over these positions. The separation can be applied to any game considered in earlier sections, by adding to it some structure on the state space, described below. The construction is related to Harsanyi's "extended preferences" (Harsanyi, 1977, Chapter 4). Unlike that concept, however, ours does not require the spectator behind the veil of ignorance to consider other agents' preferences when evaluating positional states, and is thus immune from the criticism which extended preferences have attracted (see, e.g., Adler (2014)) because it requires that the spectator to imagine an "impossible state of affairs" of imagining to be someone she is not.

Formally, we consider a dynamic game with N voters, indexed by i, and N positions, indexed by j (the terminology follows Harsanyi (1955)). Voters are characterized by their preferences, while positions correspond to possible places in society. Because our notion of anonymous policy is purely based on positions, as defined above, we ignore for now voters' preferences, and focus on how the game unfolds for each position. To each time t and role j corresponds a "positional state" θ_t^j, which summarizes the payoff-relevant history of the game up to time t for a voter in position j (e.g., j's employment history, the evolution of her savings, whether she is married to someone in position j' and the circumstance of j', as well as aggregate variables such as income inequality in the society where j lives at time t).[27] Given any position j, a policy C (as defined in the main text) determines a probability distribution over terminal positional states θ_{T+1}^j for j.

Restricting commitment to anonymous policies

A set of positions belongs to the same *category* if switching labels for these positions does not alter the game. For example, if we switch players 2 and 3 in a game and the resulting game is identical to the initial one, 2 and 3 were in the same category. Mathematically, it means that the games before and after relabeling are isomorphic to each other, and that the isomorphism is the identity mapping. In the game described by Figure 5, B and C belong to the same category and A forms a category of its own. If, however, B or C were chosen with different probabilities, then all three players would belong to different categories.

Having partitioned the set of positions into categories, a policy is said to be anonymous if it

[26] For example, if 1 had a higher income than the other two agents, taxing 1 more than these agents would not violate anonymity. Our second approach provides a more comprehensive definition of anonymity which involves all agents in the game.

[27] Positional states contain redundancies among them (for example, the quantity of a public good at time t would be reflected in each θ_t^j). Moreover, the union of positional states at time t across all positions, $\{\theta_t^j : 1 \le j \le N\}$, contains the same information as did state θ_t in the original game of Section 3.

generates the same distribution over terminal positional states for all positions within any given category.[28] Intuitively, this means that the policy does not use any preferential treatment to discriminate between otherwise identical agents. A policy that is ex post asymmetric for agents within the same category will still be anonymous as long as it is ex ante symmetric. For example, in the game in Figure 5 a policy that randomizes between U and D, but does so with the same probabilities regardless of whether B or C is chosen, will be considered anonymous, even though B and C are treated differently ex post.

Finally, one removes from consideration at the commitment stage all non-anonymous policies. The removal may be justified on normative grounds (it is unethical to treat otherwise identical people differently), as well as motivated by practical concerns (it may be difficult to dissociate individuals in any given category). The proofs of our theorems rely on the fact that all state-contingent policies are feasible. Removing discriminatory policies, we may destroy Condorcet cycles and thus identify a socially-preferred commitment.

To illustrate this approach, we revisit our simplified version of Fernandez and Rodrik (Section 5). Each role $j \in \{I, II, III\}$ receives a 0 payoff if no reform takes place, $g - c$ $(l - c)$ if the reform is implemented for one period and j is a winner (loser), $2g - c$ $(2l - c)$ if the reform lasts for two periods and j is a winner (loser). Since the payoffs are symmetric across positions and the probability of being a winner is also the same for each position, all three positions clearly belong to the same category. Therefore, any policy that treats these positions differently is discriminatory. In particular, all "conditional reforms" (e.g., continue the reform only if role 1 is a winner) are discriminatory because they induce different distributions over terminal positions for at least two roles. For example, the policy "start the reform and continue unless I or II are the winners" induces different distributions over positional states between III and the other two agents). The only reforms that remain after we eliminate the discriminatory ones, are the short-term reform in the first period, the short-term reform in the second period, the long-term reform, and the status quo. Among these, the long-term reform is the Condorcet winner. It is worth noting that by allowing voters to have different preferences, we leave enough richness in the game even after discriminatory policies are eliminated. In particular, even if *all* positions belong to the same category, the elimination of discriminatory policies does not preclude Condorcet cycles among the remaining policies.

Veil of ignorance

Many models, in political economy, macroeconomics and other fields of economics, feature interchangeable agents (sometimes a continuum thereof), which confers much strength to the above concept of anonymity.[29] In general, however, agents with distinct situations at the beginning of the game cannot be compared with the first approach, and the corresponding concept of anonymity is unable to make any comparison between such agents. Our second approach circumvents this problem by having voters compare policies behind a veil of ignorance. Given any policy, a voter

[28] For expositional simplicity, we exclude zero-probability events from this presentation.

[29] In general, one may also coarsen categories so that almost interchangeable agents are lumped together. Doing so increases the strength of the concept.

now faces a compound lottery: First, each voter is assigned, with equal probability, one of the N positions in the game (so, for instance, agent '1' in the game of Figure 5 would now be as likely to face the outcome distribution of either "position" 2 and 3 as he would to face his actual condition). Second, each voter evaluates the policy from a particular position against her preferences.

To implement this approach, we need to specify a voter's preference for each position and policy in the game: i's utility function, $\tilde{u}_i(\theta^j_{T+1})$, is defined over the set of terminal "positional" states.[30] Given a policy C, as seen from time t, voter i's expected utility when she occupies a position j in the game is given by $\widetilde{V}^{ij}_t = E(\tilde{u}_i(\theta^j_{T+1}) \mid \theta_t, C)$. For a given bijection $J(i)$ between the set of voters and positions we can then easily define the Coalitional Equilibrium as before by recovering our original utility function as $u_i(\theta_{T+1}) = \tilde{u}_i(\theta^{J(i)}_{T+1})$.

Assuming that voters maximize their expected utility, voter i's value from a policy C is, behind the veil of ignorance, given by

$$U_i(C) = \frac{1}{N} \sum_{j=1}^{N} E\tilde{u}_i(\theta^j_{T+1} \mid \theta_0, C). \tag{1}$$

Equipped with individual preference orderings, we can aggregate individual preferences behind the veil of ignorance, for example using majority voting. If all voters had identical preferences (i.e., identical $\tilde{u}_i$'s), the social ranking would select the utilitarian policy. In general, voters have heterogeneous preferences, and majority voting behind the veil of ignorance may lead to Condorcet cycles of a different nature from those predicted by Theorem 2.

Applying this second approach to our simplified version of Fernandez and Rodrik's model gives the unconditional reform as the Condorcet winner, because voters have identical preferences in this example. The veil of ignorance approach destroys the appeal of position-specific conditional reforms (such as continuing trade liberalization only if a specific industry sector benefits from it) by forcing voters to evaluate these policies without knowing their position in the game. Likewise, in the example of Figure 5, if position 1 were discriminated against in the sense that all 3 voters would suffer had they occupy this position, all voters would take this into account behind the veil when evaluating policies.

8 Discussion

Can political inertia and inefficient equilibria be resolved through the use of commitment? Introducing the option to commit to any state-contingent policy results in an unambiguous social gain only if the power structures used to compare policies under commitment and in the dynamic game are inconsistent. Under power consistency, attempts to improve on an inefficient equilibrium through commitment run into the problem of indeterminacy. This finding holds for general state processes

[30] An agent's terminal state contains his state at all earlier periods of the game. His terminal utility can therefore be an arbitrary function of his instantaneous utilities at all periods, as explained in Section 6. This structure is quite general, but it does not allow for voters to care about other voters' preferences in the game.

and utility functions, allowing social learning, experimentation, and arbitrarily heterogeneous payoffs.

Examples abound in the literature of dynamic games where equilibria are inefficient in the absence of commitment. If we preserve the power structure of the game, then our theorem shows that commitment in general cannot solve the problem, unless there is a rationale for taking certain policies a priori off the table. Policies might be ruled out based on ethical considerations, such as anonymity, or because an agenda setter controls the options (for example, Compte and Jehiel (2010) and Diermeier and Fong (2011).

Without restrictions on the policy space, commitment can be valuable—in the sense that it leads to a policy which is preferred to all others—only if we rank policies in ways that violate power consistency. That is, the social preference for certain future choices that is implied by the social preference over policies must conflict with the will of the decision makers assigned by the underlying dynamic game to make those choices. There is a shift of power toward time-zero decision makers.

In some cases, this is desirable. Commitment devices, such as contracts (backed by law enforcement), are created in order to facilitate power transfers. A contract requires mutual agreement at the outset from those with future decision rights. Why such an agreement is feasible is not fully explored in the contracting literature: it only has to improve on the status quo that each party can unilaterally defect to, but not on the set of all policies that could be considered. This is precisely why renegotiation is an issue: once the parties have locked in a series of choices, scenarios can arise where those with decision power at that time (and potentially everyone else) would like to deviate from the agreed path.[31]

Closely related are models of time inconsistency where inefficiencies arise because voters or governments anticipate that evolving states will create new policy biases in the future. (The general problem was highlighted by Kydland and Prescott (1977).) This includes the literature on redistribution (see, e.g., Campante and Ferreira (2007), Azzimonti et al. (2008), Klein et al. (2008)), where the ex-post allocation of productivity gains cannot be guaranteed, and therefore policies and actions are taken that fail to maximize overall welfare. In the representative democracy models of Persson and Svensson (1989), Besley and Coate (1998), Krusell (2002), or Saint-Paul et al. (2014), governments manipulate the state (debt, wealth), creating inefficiencies, in order to influence future political decisions. Ex ante commitment to long-run policies may restore efficiency, depending on the context: When the power consistency condition holds and some ex ante commitment is clearly beneficial to all agents, an anonymity criterion may be used to select it, as discussed in Section 7. In other cases, such as the variation on Besley and Coate (1998) studied in Section 5, the power consistency condition holds and allowing all commitments is likely to yield an indeterminacy. In the remaining cases, power consistency fails and commitment appears to be valuable. However, when the condition fails, such commitment need not be credible as commitment devices are unlikely to be available.[32] For example, (Tabellini and Alesina, 1990, p. 46) observe, regarding the efficiency

[31] In dynamic contracting models, information is typically learned about the quality of the agent, which causes the principal to want to revise the incentives (see, for example, Laffont and Tirole (1990)).

[32] "Natural" commitment devices that are "hardwired" into the game, such as a first-mover's option to invest in a cost-saving technology, do not violate power consistency and will lead to equilibria that commitment which respects power consistency cannot improve on, so they satisfy our theorem. In general, any "partial" commitment to a specific move, rather than a complete policy, is of this type. Consider, for example, a hold-up scenario: Supplier A doesn't

of a balanced-budget rule, that

> More generally, each current majority does not want to be bound by the rule, even though it wants the rule for all future majorities. However, a budget rule taking effect at some pre-specified future date would be irrelevant: if the rule can be abrogated by a simple majority, then any future majority would (...) abrogate the rule. Using again the terminology of Bengt Holmstrom and Robert Myerson (1983), we conclude that in our model a balanced budget rule, though ex ante efficient,is not "durable" under simple majority.

In overlapping or successive generations models allowing commitment to long-run macroeconomic policy, future decision makers are not involved in selecting among policies simply because they are not around at the outset. While it is then possible to agree on a policy, power consistency is violated, and one should expect attempts by future generations to renege. The assumption that commitment is feasible (can be implemented over time, rather than merely intended by the first generation) has a weaker foundation than if power consistency were satisfied.

In extreme cases, the power shifts associated with commitment that is not power-consistent can lead to outcomes that are disastrous for the group that is deprived of its decision rights. Then, violating power consistency seems ethically hard to defend. Models with equilibria leading to immiseration have this character. In the original example given by Bhagwati (1958), immiserizing growth occurs when increases in the output of an export good reduce its cost so much that the country overall sustains a welfare loss from the rise in the relative cost of imports. Subsequent examples included ex ante optimal policies (such as tariffs in the trade literature, see Johnson (1967) and Bhagwati (1968)), rendered suboptimal by growth, or policies favored by the initial generation (Matsuyama (1991), Farhi and Werning (2007)) that harm future generations. A third strand of this literature are repeated principal-agent models where efficient policies impose costs on a subset of the population in order to achieve incentive compatibility (for example Thomas and Worrall (1990), Atkeson and Lucas (1992), Zhang (2009)).

In all these cases, power consistency is implicitly violated by a long-run policy that ignores the interests of a group that should arguably have control at a future time. Immiseration results reflect that commitments are being allowed, or exclusively considered, that are only optimal with respect to initial preferences and seem infeasible to maintain because they become so objectionable over time.

want to make the upfront investment unless client B commits to a high price. B has an incentive to make that commitment, but the game does not provide a mechanism (if it did, B would simply commit in equilibrium). With complete policies and power consistency, we get a cycle as follows: A invests and B commits to high price is beaten by A invests and B commits to a low price (since B gets to rank these two, by power consistency).

Appendices

A Proofs

A.1 Computations and for the slippery slope example

We solve the game by backward induction, using the elimination of weakly dominated strategies as a refinement. If modest reform is launched, and the expansion turns out to be feasible, then a majority consisting of types B and C votes to institute a radical reform. Anticipating this, type-A voters initially vote for the project if:

$$E(u_A) = -2q + 1 - q \geq 0 \iff q \leq \frac{1}{3}.$$

B votes for the initial reform regardless of q, because B benefits whether or not an expansion is feasible. Type-C voters support the reform if:

$$E(u_C) = q - 2(1-q) \geq 0 \iff q \geq \frac{2}{3}.$$

Overall, there is a majority in favor of the initial reform at the outset if $q \leq 1/3$ (in which case, it is supported by types A and B) or $q \geq 2/3$ (then, the project is supported by types B and C). But in case $1/3 < q < 2/3$, types A and C join forces, so that a majority opposes the reform at the outset.

A.2 Proof of Theorem 1

We fix any policy Y and let $\mathcal{S}$ denote the set of coalitions with at least $N/2$ voters. For each θ_t, $a \in A(\theta_t)$ and policy X, let $S(a \mid \theta_t, X)$ denote the set of voters who strictly prefer a to the other action in $A(\theta_t)$, given the current state θ_t and given that the continuation policy from $t+1$ onwards is X. The set Θ_T can be partitioned into $A_T \bigcup \left(\bigcup_{S\in\mathcal{S}} B_T(S)\right)$, where $B_T(S) = \{\theta_T \in \Theta_T : Z_T(\theta_T) \neq Y_T(\theta_T) \text{ and } S(Z_T(\theta_T) \mid \theta_T, Z)) = S\}$ and A_T consists of all remaining states in Θ_T. In words, $B_T(S)$ consists of all the states at the beginning of period T for which the set of voters who strictly prefer the action prescribed by Z over the one prescribed by Y is equal to S.[33] A_T consists of all the states for which Y_T and Z_T coincide.[34] We will index the coalitions in $\mathcal{S}$ from S_1 to $S_{\bar{p}}$, where $\bar{p}$ is the cardinal of $\mathcal{S}$. Consider the sequence of policies $\{Y_T^p\}_{p=1}^{\bar{p}}$, defined iteratively as follows:

- Y_T^1 is equal to Y for all states except on $B_T(S_1)$, where it is equal to Z.
- For each $p \in \{2, \ldots, \bar{p}\}$, Y_T^p is equal to Y_T^{p-1} for all states except on $B_T(S_p)$, where it is equal to Z.

By construction, $Y_T^1 \succeq Y$ because the policies are the same except on a set of states where a majority of voters prefer Z (and, hence, Y_T^1) to Y. Moreover, because voters are assumed to have strict preferences over actions, the social preference is strict if and only if $B_T(S_1)$ is reached with positive probability under policy Y: $Y_T^1 \succ Y \Leftrightarrow Pr(B_T(S_1) \mid Y) > 0$. If $Pr(B_T(S_1) \mid Y) = 0$, $Y_T^1 = Y$ with probability 1.

Therefore, either Y and Y_T^1 coincide, or $Y_T^1 \succ Y$. Similarly, $Y_T^p \succeq Y_T^{p-1}$ for all $p \leq \bar{p}$, and $Y_T^p \succ Y_T^{p-1}$ if and only if $Y_T^p \neq Y_T^{p-1}$ with positive probability. This shows that

$$Y_T^{\bar{p}} \succeq \cdots \succeq Y_T^1 \succeq Y,$$

[33] In particular, $Y_T(\theta_T) \neq Z_T(\theta_T)$ for all those states.

[34] Because Z is the equilibrium policy, the set of voters who prefer Y over Z at time T must always form a minority, so A_T and $\bigcup_{S\in\mathcal{S}} B_T(S)$ exhaust all states in Θ_T.

and at least one inequality is strict if and only if the set of states in Θ_T over which Z_T and Y_T are different is reached with positive probability under Y. By construction, $Y_T^{\bar{p}}$ coincides with Z on Θ_T: $Y_T^{\bar{p}}(\theta_T) = Z(\theta_T)$ for all $\theta_T \in \Theta_T$.

We now extend the construction by backward induction to all periods from $t = T - 1$ to $t = 1$. For period t, partition Θ_t into $A_t \bigcup \left(\bigcup_{S \in \mathcal{S}} B_t(S)\right)$, where A_t consists of all θ_t's over which Y_t and Z_t coincide, and $B_t(S) = \{\theta_t : Z_t(\theta_t) \neq Y_t(\theta_t) \text{ and } S(Z_t(\theta_t) \mid \theta_t, Z)) = S\}$. That is, $B_t(S)$ consists of all states in Θ_t for which the set of voters who strictly prefer the action prescribed by Z over the one prescribed by Y, *given that Z is used for all subsequent periods*, is equal to S.[35] Y_t^p is defined inductively as follows, increasing p within each period t, and then decreasing t: for each t,

- For $p = 1$, Y_t^1 is equal to $Y_{t+1}^{\bar{p}}$ for all states, except on $B_t(S_1)$, where it is equal to Z.
- For $p > 1$ Y_t^p is equal to Y_t^{p-1} for all states, except on $B_t(S_p)$ where it is equal to Z.

By construction, $Y_t^{p+1} \succeq Y_t^p$ for all t and $p < \bar{p}$ and $Y_t^1 \succeq Y_{t+1}^{\bar{p}}$ for all t. Moreover, the inequality is strict if and only if the policies being compared are not equal with probability 1 on the set of states reached by either of them.

Finally, observe that $Y_1^{\bar{p}} = Z$. Let $\{Y_k\}_{k=1}^K$, $K \geq 1$, denote the sequence of *distinct* policies obtained, starting from Y, by the previous construction, iterating from $t = T$ and $p = 1$ down to $t = 1$ and $p = \bar{p}$.[36]

If $Y \neq Z$ with positive probability, then $K \geq 2$. Moreover,

$$Y = Y_1 \prec Y_2 \cdots \prec Y_K = Z. \tag{2}$$

Therefore, we get a voting cycle if $Z \prec Y$, which concludes the proof of part i).

Since Z can never be defeated without creating a cycle, we can characterize a Condorcet winner out of all policies, if it (they) exists, and ii) follows. As mentioned in Remark 1, the entire proof goes through if one drops the assumption that voters have strict preferences. The only difference now is that (2) only holds with weak inequalities.

A.3 Proof of Theorem 3

We set the terminal payoffs equal to 0 for all policies, except when i) the action sequence until time t has followed policy Y (and hence Y') *and* ii) the state θ_t reached at time t belong to $\bar{\Theta}_t$. In that case, members of coalition S (its complement $\bar{S}$) get 100 (10) if the time t action prescribed by Y' is played and followed by the continuation of policy Y (and hence Y'), with the reverse payoffs if instead the action prescribed by Y is played at time t and followed by the continuation of policy Y. If $\theta_t \notin \bar{\Theta}_t$, everyone's payoff is set to some small $\varepsilon > 0$ for the common continuation of Y and Y', and to zero for all other continuations.[37]

Let Z denote the equilibrium policy. Z must coincide with Y and Y' until period t since it is the only way, for any player, to get a nontrivial payoff. If the state $\theta_t \in \bar{\Theta}_t$, coalition S imposes at time t the action prescribed by Y', so as to achieve its highest possible payoff of 100, and it is in everyone's interest to implement the continuation corresponding to Y' from time $t + 1$ onwards, so that even members of $\bar{S}$ get their second highest payoff of 10. Even if $\theta_t \notin \bar{\Theta}_t$, it is also in everyone's interest to follow the common continuation of Y and Y' so as to get ε. This shows that Z coincides to Y'.

We now consider the social comparisons of policies. Clearly, $\bar{S}$ imposes $Y \succ Y'$ since it has the power to do so and this achieves maximal expected payoff. Moreover, there cannot be any cycle among policies, since Y and Y' Pareto dominate all other policies. Thus, Y is the Condorcet winner[38] among all policies.

[35] Again, by definition of Z, there cannot be a majority who prefer Y_t over Z_t, given the continuation policy $\{Z_t'\}_{t' \geq t}$, so $A_t \bigcup \left(\bigcup_{S \in \mathcal{S}} B_t(S)\right) = \Theta_t$.

[36] We call two policies *distinct* if they induce different distributions over Θ_T. Policies that differ only at states that are never reached are not distinct.

[37] Although this is largely irrelevant to the gist of the present argument (see Remark 2), one may add arbitrarily small, action-dependent payoffs to break ties at all points of the game for all histories as required by the non-indifference assumption of Theorem 2.

[38] Any other Condorcet winner X must be identical to Y except on a set of histories which has probability 0 when Y (or, equivalently, X) is followed.

B Limits of agenda setting: ordering-invariant, Pareto-dominated equilibrium

There are two decisions i) U or D and ii) L or R, and three voters. The decisions can be made in any order and are made according to the simple majority rule. The state of the world is revealed between the decisions. The state of the world can take six possible values with probability 1/6 each, and be broken down into two components: the role played by each decision (2 possibilities) and, independently, which player is the "sucker" (3 possibilities). We index the states as $\theta_{\omega i}$ where i indicates the sucker and $\omega \in \{A, B\}$ indicates the role of decisions. The payoff structure is as follows:

In state θ_{Ai}, all players get a payoff of 1 if L is chosen regardless of the other decision (U or D). If R is chosen, then all players get a payoff of 2 if D is chosen, while player i gets -100 and other players get 3 if U is chosen. Importantly, these payoffs do *not* depend on the order in which the actions (U/D and L/R) are chosen.

In state θ_{Bi}, the role of actions is reversed as follows: all players get a payoff of 1 if U is chosen, regardless of the other decision (L or R). If D is chosen, then all players get a payoff of 2 if R is chosen, while player i gets -100 and other players get 3 if L is chosen. Again, these payoffs *do not depend* on the order of decisions (U/D and L/R).

The game is constructed in such a way that there is a commitment problem: players would ex ante all prefer to get the payoff of 2, which is always achievable by committing to the action profile (D, R), regardless of the order of these decisions. Without commitment however, there is always a probability 1/2 that, in period 2, two players gain by imposing the action that gives them 3 and gives the other player -100. For example, if L/R is the first decision, then if players choose action R and the action state ω is A, two players impose U in the second period and the other one gets -100. If instead they choose action L and the action state ω is B, then two players impose action D in the second period. Because of this, players' expected payoff in equilibrium is of the order $-100/6 \sim -15$.

By symmetry, players get exactly the same expected payoffs if instead the action U/D is taken in the first period.

In contrast, committing to (D, R) yields a payoff of 2, regardless of the state of the world. Thus, reordering actions *per se* does not reveal the value of commitment. Of course, as implied by Theorem 1, there is a cycle among commitment policies, as is easily checked.

References

Acemoglu, D., Egorov, G., and Sonin, K. (2012). Dynamics and stability of constitutions, coalitions, and clubs. *American Economic Review*, 102(4):1446–1476.

Acemoglu, D., Egorov, G., and Sonin, K. (2014). Political economy in a changing world.

Acemoglu, D. and Robinson, J. A. (2005). *Economic origins of dictatorship and democracy.* Cambridge University Press.

Adler, M. D. (2014). Extended preferences and interpersonal comparisons: A new account. *Economics and Philosophy*, 30(02):123–162.

Alesina, A. and Tabellini, G. (1990). A positive theory of fiscal deficits and government debt. *The Review of Economic Studies*, 57(3):403–414.

Atkeson, A. and Lucas, R. E. (1992). On efficient distribution with private information. *The Review of Economic Studies*, 59(3):427–453.

Atkinson, A. B. (1970). On the measurement of inequality. *Journal of economic theory*, 2(3):244–263.

Azzimonti, M., de Francisco, E., and Krusell, P. (2008). Production subsidies and redistribution. *Journal of Economic Theory*, 142(1):73–99.

Barbera, S. and Jackson, M. (2004). Choosing how to choose: Self-stable majority tules and constitutions. *The Quarterly Journal of Economics*, 119(3):1011–1048.

Barbera, S., Maschler, M., and Shalev, J. (2001). Voting for voters: A model of electoral evolution. *Games and Economic Behavior*, 37(1):40–78.

Battaglini, M. and Coate, S. (2008). A dynamic theory of public spending, taxation, and debt. *American Economic Review*, 98(1):201–36.

Besley, T. and Coate, S. (1998). Sources of inefficiency in a representative democracy: a dynamic analysis. *American Economic Review*, pages 139–156.

Bhagwati, J. (1958). Immiserizing growth: A geometrical note. *The Review of Economic Studies*, 25(3):201–205.

Bhagwati, J. N. (1968). Distortions and immiserizing growth: a generalization. *The Review of Economic Studies*, 35(4):481–485.

Boylan, R. T., Ledyard, J., and McKelvey, R. D. (1996). Political competition in a model of economic growth: Some theoretical results. *Economic Theory*, 7(2):191–205.

Boylan, R. T. and McKelvey, R. D. (1995). Voting over economic plans. *American Economic Review*, pages 860–871.

Campante, F. R. and Ferreira, F. H. (2007). Inefficient lobbying, populism and oligarchy. *Journal of Public Economics*, 91(5):993–1021.

Compte, O. and Jehiel, P. (2010). Bargaining and majority rules: A collective search perspective. *Journal of Political Economy*, 118(2):189–221.

Diermeier, D. and Fong, P. (2011). Legislative bargaining with reconsideration. *The Quarterly Journal of Economics*, 126(2):947–985.

Dziuda, W. and Loeper, A. (2014). Dynamic collective choice with endogenous status quo.

Farhi, E. and Werning, I. (2007). Inequality and social discounting. *Journal of political economy*, 115(3):365–402.

Fernandez, R. and Rodrik, D. (1991). Resistance to reform: Status quo bias in the presence of individual-specific uncertainty. *American Economic Review*, 81(5):1146–55.

Galperti, S. and Strulovici, B. (2014). Forward-looking behavior revisited: A foundation of time inconsistency.

Glaeser, E. L. and Shleifer, A. (2005). The curley effect: The economics of shaping the electorate. *The Journal of Law, Economics, & Organization*, 21(1):N1.

Harsanyi, J. (1977). Rational behavior and bargaining equilibrium in games and social situations.

Harsanyi, J. C. (1955). Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *The Journal of Political Economy*, 63(4):309–321.

Jackson, M. O. and Yariv, L. (2014). Collective dynamic choice: The necessity of time inconsistency,. *American Economic Journal: Microeconomics*. Forthcoming.

Johnson, H. G. (1967). The possibility of income losses from increased efficiency or factor accumulation in the presence of tariffs. *The Economic Journal*, 77(305):151–154.

Klein, P., Krusell, P., and Rios-Rull, J.-V. (2008). Time-consistent public policy. *The Review of Economic Studies*, 75(3):789–808.

Krusell, P. (2002). Time-consistent redistribution. *European Economic Review*, 46(4):755–769.

Kydland, F. and Prescott, E. (1977). Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, 85(3):473–491.

Laffont, J.-J. and Tirole, J. (1990). Adverse selection and renegotiation in procurement. *Review of Economic Studies*, 57(4):597–625.

Laslier, J.-F. (1997). *Tournament Solutions and Majority Voting*. Springer, Berlin.

Matsuyama, K. (1991). Immiserizing growth in diamond's overlapping generations model: A geometrical exposition. *International Economic Review*, 32(1):251–62.

Miller, N. R. (1977). Graph-theoretical approaches to the theory of voting. *American Journal of Political Science*, pages 769–803.

Moldovanu, B. and Shi, X. (2013). Specialization and partisanship in committee search. *Theoretical Economics*, 8(3):751–774.

Persson, T. and Svensson, L. E. (1989). Why a stubborn conservative would run a deficit: Policy with time-inconsistent preferences. *The Quarterly Journal of Economics*, 104(2):325–45.

Rawls, J. (1971). A theory of justice.

Roberts, K. (1999). Dynamic voting in clubs. http://ssrn.com/abstract=1160966.

Saint-Paul, G., Ticchi, D., and Vindigni, A. (2014). A theory of political entrenchment. *The Economic Journal*. Forthcoming.

Sen, A. (1970). The Impossibility of a Paretian Liberal. *Journal of Political Economy*, 78(1):152–57.

Shepsle, K. A. (1970). A note on zeckhauser's "majority rule with lotteries on alternatives": The case of the paradox of voting. *Quarterly Journal of Economics*, 84(4):705–09.

Strulovici, B. (2007). Voting and experimentation. Working paper, Oxford University.

Strulovici, B. (2010). Learning while voting: Determinants of collective experimentation. *Econometrica*, 78(3):933–971.

Tabellini, G. and Alesina, A. (1990). Voting on the budget deficit. *American Economic Review*, 80(1):37.

Thomas, J. and Worrall, T. (1990). Income fluctuation and asymmetric information: An example of a repeated principal-agent problem. *Journal of Economic Theory*, 51(2):367–390.

Volokh, E. (2003). The mechanisms of the slippery slope. *Harvard Law Review*, pages 1026–1137.

Zeckhauser, R. (1969). Majority rule with lotteries on alternatives. *Quarterly Journal of Economics*, 83(4):696–703.

Zhang, Y. (2009). Dynamic contracting with persistent shocks. *Journal of Economic Theory*, 144(2):635–675.

C Online Appendix

C.1 Existence of a Condorcet Winner

Terminal States

The setting is specialized as follows: there is finite set S of alternatives. The dynamic game starts with some status quo alternative s_0. Thereafter, in each period, a new alternative in S is proposed against the last accepted alternative. If the challenger is accepted, it becomes the new 'status quo.' Otherwise, the previous status quo is pitted against a new challenger. The state θ_t corresponds to the current status quo. So, without loss, $\theta_t \in S$. It is assumed that the proposal protocol is flexible enough so as to guarantee that each possible sequence of S states may arise under some policy as the succession of status quo. This is achieved for instance if the states are ordered from s_1 to s_S and, the number T of periods is such that $T > S^2$, and at each time $t \leq T$ the challenger is the state s_k where k is the remainder in the Euclidean division of t by S.

The payoff of each player i is determined by the state s at the end of the last period. Formally, i's payoff u_i is a map from S to $\mathbb{R}$. Therefore, any policy reduces, in terms of payoffs, to the last alternative. In particular, there is a Condorcet winner among policies if and only if there is a Condorcet winner among states. As it turns out, this exact condition has been studied by Acemoglu et al. (2012), hereafter (AES). In particular, part i) of their Assumption 2 is identical to ruling out Condorcet cycle.[39] Theorem 4 in AES provide sufficient conditions for acyclicity. By a direct translation of their result, we obtain a sufficient condition for the existence of a Condorcet winner. For completeness, we recall the assumptions of AES:

Say that a set $\mathcal{C}$ of coalitions is *regular* if it satisfies the following properties: i) For any $W \in \mathcal{C}$ and $W \subset W'$, we have $W' \in \mathcal{C}$, ii) For any $W, W' \in \mathcal{C}$, $W \cap W' \neq \emptyset$.

For the sufficient condition, we order the set of players from 1 to N. The preference profile $\{u_i(\cdot)\}_{i \in N}$ obeys the *single-crossing property* if for any ordered states $s < s'$ and ordered players $i < j$, $u_i(s') \geq (>) u_i(s)$ implies $u_j(s') \geq (>) u_j(s)$.

Proposition 3. *Suppose the players satisfy the single-crossing property and that for each state s, the set of winning coalitions is regular. Then, the equilibrium of the coalitional game selects the unique Condorcet winner among the S alternatives.*

Proof. The assumptions of Theorem 4 in AES are satisfied.[40] Theorem 4 then implies acyclicity of coalitional preferences which in turn guarantees the existence of a Condorcet winner s^* among alternatives and, by equivalence, among plans. Theorem 2 then implies that the equilibrium selects that policy and hence s^* as the terminal outcome. $\square$

AES also provide conditions – single-peaked preferences and a nonempty intersection condition concerning winning coalitions across different states – under which the stable set is Pareto efficient (Theorem 3). The conditions may readily be adapted to our setting to guarantee that the equilibrium yields a Pareto efficient outcome (regardless of whether it is a Condorcet winner).

Single-crossing property for policies

Interpret the succession of binary decisions of the dynamic game as society's gradual positioning on a single issue, captured on a one-dimensional spectrum, with $\underline{a}_t$ (resp., $\bar{a}_t$) being a move to the left (right). A state-contingent policy thus corresponds to the gradual refinement of the collective position. The policy $\underline{a}_1, \underline{a}_2$ in a two-period problem is clearly more "left" than the policy $\bar{a}_1, \bar{a}_2$. The comparison of $\underline{a}_1, \bar{a}_2$ with $\bar{a}_1, \underline{a}_2$ depends on the relative magnitude of first and second move. In some settings there is a clear ranking of these magnitudes. For example, the first decision may be of first-order importance, while the second decision concerns a finer point. More generally, policies may be ranked lexicographically: we have $Y \prec Y'$ if the first period t for which there exists a state θ_t such that $Y_t \neq Y'_t$, we have $Y_t = \underline{a}_t$ and $Y'_t = \bar{a}_t$ for all the states of period t on which Y and Y' are distinct. This ordering captures the idea

[39]In both settings, the set of winning coalitions only depends on the current state.

[40]The assumption, maintained throughout the present paper, that players are never indifferent across actions implies that Assumption 6 in AES is satisfied.

that earlier decisions matter more for the policy than later refinements. Policies may then be ranked unequivocally from left to right, and the single-crossing property may be imposed over preferences. With the simple majority rule, the median voter theorem guarantees the existence of a Condorcet winner.

C.2 Voting cycles over infinite-horizon policies: An illustration with collective experimentation

The argument used to prove Theorem 1 may be adapted to Markovian settings, applying backward induction to the underlying state rather than to time. To illustrate this, we revisit the model of Strulovici (2010), focusing on three ex ante symmetric voters. Voters choose at each instant between a risky action (the "reform") and a safe action (the "status quo"). As long as the reform is implemented, each voter may receive some good news, which reveals that the reform benefits him; he becomes a "winner" of the reform. Voters who receive no news, called "unsure voters", become more pessimistic about the reform and wish to abandon it in favor of the status quo. Reform is then abandoned if reform winners did not gain the majority. If the safe action is chosen for some belief of unsure voters, no new learning occurs and experimentation is forever abandoned.

The utilitarian policy, denoted by Y, is ex ante strictly preferred to the equilibrium policy, Z (Theorem 6 in Strulovici (2010)).[41] The equilibrium policy is characterized by cutoffs $p(0) > p(1)$ such that experimentation stops when unsure voters' belief p drops to $p(k)$ and no more than k winners have occurred by then (Theorem 1). The utilitarian policy is determined by similar thresholds $q(k)$ (Theorem 2). From Theorem 3, $q(k) < p(k)$ for $k = 0, 1$. Intuitively, experimentation is more valuable to the social planner than to unsure voters, because he includes the utility of winners in his welfare function and has a higher option value of experimentation than individual voters who have to share power. Suppose that $q(1) = 0$ – it is efficient to play the risky action forever from the moment that one winner has been observed. This condition is equivalent (Theorem 2) to $g \geq 3s$, where $g > 0$ is a winner's flow payoff from the reform and $s \in (0, g)$ is everyone's flow payoff with the status quo. This parametric condition is imposed hereafter.

The thresholds $p(1)$ and $q(0)$ solve the following equations:[42]

$$p(1) = \frac{\mu s}{\mu g + (g - s) + p(1)g - s} \tag{3}$$

$$q(0) = \frac{\mu s}{\mu g + (g - s) + 2(q(0)g - s)}. \tag{4}$$

Because $p(1)$ and $q(0)$ are strictly below the myopic cut-off s/g (Theorem 1), this implies that $q(0) > p(1)$.[43] This implies $p(0) > q(0) > p(1)$.

We now construct a cycle, based on the following modifications of the utilitarian policy. Let Y_1 denote the policy that coincides with the social optimum, Y, except that experimentation stops at the threshold $p(1)$ if the only winner observed by that time is Voter 1. Voters 2 and 3 prefer this policy to the utilitarian policy, conditional on Voter 1 being the only winner by the time $q(0)$ is reached. In fact, this policy maximizes Voters 2 and 3's common expected utility. To see this, notice that the utilitarian cutoff $\tilde{q}$ for Voters 2 and 3, starting from $q(0)$, and given that any winner will cause L to be played forever, is characterized by the indifference equation[44]

$$\tilde{q}g + 2\tilde{q}\lambda\left(\frac{1}{2}\left(\frac{g}{r} + \tilde{q}\frac{g}{r}\right) - \frac{s}{r}\right) = s.$$

The first term on the LHS is the average flow payoff for Voters 2 and 3 when the risky action is played. The second term is their average utility jump if one of them becomes a winner, multiplied by the flow probability of that event.

[41] We refer the reader to the original paper for the results and notation used in this section. Theorem numbers also refer to the original paper.

[42] See equation (17), p. 964, for $N = 3$ and $k_N = 1$ and equation (7), p. 947, for $N = 3$, $k = 0$, and $W(1, p) = \frac{1}{3}(\frac{g}{r} + 2p\frac{g}{r})$.

[43] If $p(1) \geq q(0)$, then we would have $p(1)g - s \geq q(0)g - s > 2(q(0)g - s)$, where the strict inequality comes from $q(0) < s/g$. This would imply that the RHS of (3) is strictly smaller than the RHS of (4) and, hence, that $p(1) < q(0)$.

[44] See Equation (7) on p. 947 with $N = 2$ and $k = 0$, and $W(1, p) = 1/2(g/r + pg/r)$.

The RHS is their flow payoff with the safe action. The last equation is identical to (3) and thus yields the same cut-off.[45] Therefore, $Y_1 \succ Y$.

Consider now the policy Y_2 that is identical to Y_1 except that experimentation stops at $p(1)$ if only Voter 2 has become a winner by that time. This policy is preferred by Voters 1 and 3 over Y_1, by the same reasoning, so $Y_2 \succ Y_1$. Finally, Y_3 is obtained by modifying Y_2 so that experimentation stops at $p(1)$ if only Voter 3 is a winner by that time. Again, $Y_3 \succ Y_2$.

By construction, Y_3 stops experimentation if a single winner has been observed by the time $p(1)$ is reached, as does the equilibrium policy Z. The only difference between Y_3 and Z is that Z stops experimentation earlier, at the threshold $p(0) > q(0)$, if no winner has been observed, whereas Y_3 stops experimentation at $q(0)$. However, subject to the constraint that experimentation stops when p reaches $p(1)$ and at most one winner occurred, and continues forever otherwise, the socially optimal policy is actually to stop at $p(0)$, not $q(0)$. To see this notice that, due to the constraint, the continuation value of a voter who becomes a winner is precisely the continuation value w under the equilibrium policy, and the continuation value for unsure voters when another voter becomes a winner is the equilibrium continuation value, u.

Formally, the socially optimal cut-off $\hat{q}$ for stopping experimentation when no winner has been observed, given the constraint on continuation play, solves the indifference equation

$$3\hat{q} + 3\lambda\hat{q}\left(w(1,\hat{q}) + 2u(1,\hat{q}) - 3\frac{s}{r}\right) = 3s.$$

The first term of the LHS is the aggregate expected flow payoff for the three voters if L is played, and the second term is their aggregate jump in utility if one of them becomes a winner, multiplied by the flow probability of this happening. The right-hand side is the aggregate flow payoff if the safe action is played instead. Dividing by 3, one obtains exactly the characterization of the cut-off $p(0)$.[46] This shows that $Z \succ Y_3$ and yields the voting cycle $Z \prec Y \prec Y_1 \prec Y_2 \prec Y_3 \prec Z$.

The logic of the slippery slope example can also be extended to an infinite-horizon model. In particular, Strulovici (2007) studies experimentation between several restaurants, including a Singaporean restaurant which, in addition to serving Indian cuisine, can teach "voters" about Chinese cuisine.

To reinterpret this example in light of the slippery slope problem, consider the following setting, with three alternatives. S: do not undertake any reform; M: implement only the moderate reform, even if an expansion if feasible; P: implement the moderate reform, and expand if this is feasible. There are also three voters with the following preferences: Voter 1 only values the moderate reform; Voter 2 wishes to implement both reforms; Voter 3 does not want to start the moderate reform a priori, but will support an expansion if the reform is implemented and information is revealed that makes the expansion feasible.

Flow payoffs are as follows. The payoff for the status quo is normalized to 1 for all voters. If the first reform is implemented, Voters 1 and 2 get a flow payoff of 2, while Voter 3 gets 0. If the reform is expanded, Voter 1 gets a flow payoff of -9, Voter 2 gets 2.1, and Voter 3 gets 1.1. Let p denote the ex ante probability that the expansion becomes feasible. The analysis of Strulovici (2007) implies the following result.

Proposition When option P is unavailable, the only Majority Voting Equilibrium (MVE) consists in implementing M. If P is added to the set of feasible alternatives and p lies in $(0.1, 0.8)$, the only MVE consists in implementing S.

Proof. In the example analyzed by Strulovici, Voter 3 gets a payoff $\tilde{g} = 0.1$ under M if only the moderate reform is implemented (action M). The analysis of that example is identical for all $\tilde{g} \in [0, 0.1]$, and in particular for the value $\tilde{g} = 0$ used in the slippery slope example. □

[45]Equation (3) has a unique positive solution (the left-hand side is increasing in p while the right-hand side is decreasing p).

[46]See Equation (20), p. 965, with $N = 3$ and $k_N = 1$, and replacing the factor $N - k_N - 2$ by $N - k_N$).