A Model of Repeated Collective Decisions*

Antonin Macé[1] and Rafael Treibich[2]

[1]Paris School of Economics, CNRS and École Normale Supérieure-PSL
[2]University of Southern Denmark

February 28, 2023

Abstract

The theory of repeated games offers a compelling rationale for cooperation in a variety of environments. Yet, its consequences for collective decision-making have been largely unexplored. In this paper, we propose a general model of repeated voting in committees and study equilibrium behavior under alternative majority rules. Our main characterization reveals a complex, relationship between the majority threshold, the preference distribution, and the set of equilibrium payoffs. In contrast with the stage game, equilibria in the repeated game may involve a form of implicit logroll, individuals sometimes voting against their preference to achieve the efficient decision. In turn, this affects the optimal voting rule, which may significantly differ from the optimal rule under sincere voting. The model provides a rationale for the use of unanimity rule, while accounting for the prevalence of consensus in committees which use a lower majority threshold.

*We are very grateful to Vincent Anesi, Geir Asheim and Alessandra Casella for useful comments, and to participants of the following seminars: University of Strasbourg, Paris School of Economics, Online Social Choice and Welfare, COMSOC, and of the following conferences: *Aggregation Across Disciplines*, Social Choice and Welfare, ACM Conference on Economics and Computation (EC'22).

1 Introduction

Collective decisions in many committees display two key features: (i) decisions are taken repeatedly over time, and (ii) committee members often differ in how much they care about each decision. Examples include international organizations, standard-setting organizations, city councils, hiring committees, etc. For instance, the Council of the European Union, one of the three main decision bodies of the EU, issues more than 300 legislative acts every year on a wide range of topics. Countries' stakes on these proposed reforms may vary significantly depending on their economy, citizens' preferences, historical background, cultural norms, current national legislation, etc. In this paper, we propose a model of collective decisions that captures these two defining characteristics and study how they affect equilibrium voting behavior, and in turn the comparison of alternative voting rules.

In the classical setting where the committee takes a single binary decision, sincere voting is a dominant strategy under any qualified majority rule. In a repeated setting, individuals may choose to condition their voting behavior on past votes. This history dependence may lead to strategic (non-sincere) behavior, thus changing the outcome of collective decisions compared to the static, sincere, benchmark. It is thus essential to account for the repeated nature of collective decisions to understand the implications of different voting rules.

Our paper sheds light on two important facts about collective decision-making that can be hard to reconcile with standard voting theory. The first observation is that, despite strong similarities in their organization and structure, committees vary widely in the procedures and voting rules used for making collective decisions: from simple majority, unanimity rule, and weighted majority rules to the use of veto power. The heterogeneity of voting procedures can sometimes be observed within the same organization, decisions of different nature being taken according to different voting rules. For instance, the Council of the EU uses simple majority for some decisions, but the unanimity rule for others. Why do otherwise similar committees use different voting rules? And why is the unanimity rule so prevalent, when it is usually considered to be inefficient and prone to gridlock? The second observation is that decisions in many committees are often made by consensus, without taking a formal vote, even when the committee is supposed to use a (possibly non-unanimous) voting rule (Urfalino, 2014). The prevalence of consensus may seem puzzling, especially in large committees, where preferences over decisions are rarely consensual. We show that it is possible to rationalize such consensual behavior without assuming any type of social preference. Instead, our model features a social norm of consensus emerging from equilibrium behavior in a repeated interaction.

In our model an ex-ante symmetric committee makes repeated binary decisions about whether to accept or reject proposed reforms. At each stage utilities are drawn and observed publicly. Utilities reflect individuals' cardinal preferences for the reform relative to the status quo. Individuals then vote simultaneously either in favor or against the reform and a collective decision is taken according to a qualified majority rule. We study the equilibrium outcomes of this repeated game.

Our first result characterizes the set of equilibrium payoffs. The width of this set and the discount factor together determine the power of intertemporal incentives in the repeated game. In turn, this power of intertemporal incentives drives the optimal equilibrium, the one achieving the highest level of (ex-ante) utility. This equilibrium takes the form of a cooperation norm, the committee adopting efficient decisions unless it becomes too costly to incentivize the pivotal voter to abide by them. The power of intertemporal incentives thus provides a measure of the degree of cooperation that can be sustained at equilibrium. When the discount factor is large enough, even the first best can be achieved.

We then refine the analysis by focusing on individual voting behavior. We show that the cooperation norm of the optimal equilibrium may allow for a substantial level of consensual decisions, even if preferences are never consensual. When the discount factor is large enough, not only all decisions can be efficient, but any accepted reform can receive unanimous approval.

To explore the implications of the cooperation norm we uncover, we then focus on a more stylized setting where utilities write as a sum of common and private components. We derive explicit formulas for the set of equilibrium payoffs under any possible majority rule and discount factor. Taking stock, we perform two comparative statics exercises. First, we look for the optimal voting rule, the one which allows to reach the highest level of utility. While simple majority appears optimal in the static benchmark, a super-majority (or even the unanimity) can be optimal in the repeated game. Second, we show that, contrary to what we would obtain in a static model, the level of consensus may be higher when committee members have more heterogeneous preferences or when the majority threshold is lower.

1.1 Illustrative example

A group of three individuals decides whether to accept or reject repeated proposals at either *unanimity* or *simple majority*. Assume proposals are of the following two types:

- Proposals of type A benefit one individual strongly ($u_i = 5$), one individual weakly ($u_i = 1$), and hurt one individual mildly ($u_i = -3$).
- Proposals of type B benefit one individual strongly ($u_i = 5$) and hurt two individuals weakly ($u_i = -1$).

Individuals are equally likely to occupy any position for each kind of proposals. We denote by $p \in (0, 1)$ the probability of occurrence of proposals A. Since both proposals generate an average utility of 1, the first-best consists in accepting any proposal and yields ex-ante utility $U^e = 1$.

If the organization only makes one decision, voting sincerely is a weakly dominant strategy. Under *unanimity*, no reform is ever accepted at equilibrium, yielding ex-ante utility $U^u = 0$. Under *simple majority*, only proposals A are accepted at equilibrium, yielding ex-ante utility $U^s = p$. Simple majority dominates unanimity but does not achieve the first best.

If the organization makes (infinitely) repeated decisions, it is possible to incentivize individuals to sometimes vote against their preference so as to achieve a higher level of utility. Consider the following strategy profile under *unanimity*: on the equilibrium path all three individuals always vote in favor of proposals A and B; any deviation is punished by a permanent reversal to the stage-game equilibrium. Such a profile is an equilibrium if and only if the worst-off individual in proposals A has an incentive to vote in favor of the proposal. This is the case if her inter-temporal utility from the reform being adopted, $-3(1-\delta)+\delta U^e$, exceeds her inter-temporal utility from the reform being rejected, δU^u. As a result, the first-best can be sustained under this profile if $\delta \geq \delta_u^F := 3/4$. In fact, this equilibrium is optimal, and the previous condition is thus necessary for the first-best to be achieved at equilibrium.

If the discount factor is too small, full cooperation cannot be achieved at equilibrium but it may still be possible to improve over the stage-game equilibrium. Consider the following strategy profile: on the equilibrium path all three individuals vote in favor of proposals B, but they vote sincerely on proposals A; any deviation is punished by a permanent reversal to the stage-game equilibrium. Only proposals B are now accepted, which yields a lower level of utility than under the first-best, $\hat{U} = 1 - p < U^e$. However, individuals who need to be incentivized only have a stake of 1, as opposed to 3 previously. Such a profile is an equilibrium if and only if $-(1-\delta) + \delta\hat{U} \geq \delta U^u$, or equivalently $\delta \geq \delta_u^P := 1/(2-p)$. When $p < 2/3$, threshold δ_u^P is smaller than the threshold for full cooperation δ_u^F: we get an intermediate range of discount factors (δ_u^P, δ_u^F) such that unanimity cannot achieve the first-best, yet improves over the stage-game equilibrium.

Under *simple majority*, sustaining the first-best no longer requires incentivizing A's worst-off individual to vote in favor since proposals A are accepted under sincere voting. Consider the following strategy profile: on the equilibrium path all three individuals vote in favor of proposals B, but they vote sincerely on proposals A; any deviation is punished by a permanent reversal to the stage-game equilibrium. This profile is an equilibrium if and only if $-(1-\delta)+\delta U^e \geq \delta U^{sm}$, or equivalently $\delta \geq \delta_{sm}^F := 1/(2-p)$. This equilibrium is optimal, and the previous condition is thus necessary for the first-best to be achieved at equilibrium.[1] In contrast to unanimity, achieving efficiency under simple majority only requires incentivizing individuals with a stake of 1, as opposed to 3, and may thus seem easier to achieve. However, since simple majority outperforms unanimity in the stage game, the long term benefit from complying with prescribed behavior is smaller, which can eventually

[1]Under simple majority, the payoff of the stage-game equilibrium is not the lowest feasible payoff ($U^s = p > 0$). However, it can be shown (in this example) that the worst equilibrium payoff in the repeated game does in fact coincide with the payoff of the stage-game equilibrium.

make efficiency harder to sustain. Here, if $p > 2/3$, there exists a range of discount factors $(\delta_u^F, \delta_{sm}^F)$ where only unanimity can sustain the first-best.[2]

This simple example illustrates how the repeated nature of collective decisions may affect equilibrium voting behavior and allow for better outcomes to be sustained relative to the stage game. In turn, this may have consequences for the optimal voting rule and the level of consensus.[3]

1.2 Literature Review

Most of the literature on repeated collective decision-making considers decisions over a single, persistent, issue (e.g. a central bank setting the interest rate). The objective is then to understand how the endogeneity of the status quo (today's decision becomes tomorrow's status quo) affects voting behavior, most often under the assumption of Markovian strategies. For instance, Baron (1996) proposes a model of dynamic redistributive politics where legislators have single peaked preferences over a one-dimensional policy space. He characterizes voting behavior at the stationary equilibrium and show how it relates to the preference of the median voter. More recently, Dzuida and Loeper (2018) study an infinite horizon model where a group of legislators repeatedly make the same binary decision. They show how the endogeneity of the status quo leads to more polarized voting behavior at equilibrium than under sincere voting.[4] By contrast, our model considers a committee making repeated collective decisions about exogenous and independent proposals. The repeated, as opposed to dynamic, nature of our model allows us

[2]The fact that δ_{sm}^F coincides with δ_u^P is specific to this simple example. In general, unanimity can also outperform simple majority in a regime of partial cooperation.

[3]Note for instance that it is possible to achieve full consensus under simple majority if individuals are patient enough. Consider the following strategy profile: on the equilibrium path all three individuals always vote in favor of proposals A and B; any deviation is punished by a permanent reversal to the stage-game equilibrium. This profile is an equilibrium if and only if $\delta \geq 3/(4-p) > \delta_s^F$. In contrast with the original profile used to sustain the first-best, the worst-off individual in proposals A now has to be incentivized to vote in favor as well, which is more demanding and may only be achieved at equilibrium for a larger discount factor.

[4]Other references in this literature include Kalandrakis (2004), who extends Baron (1996)'s model of legislative bargaining to a dynamic setting by considering a three players version of the divide the dollar game with endogenous reversion point, Duggan and Kalandrakis (2012), who prove the existence of stationary Markov perfect equilibria in an infinite-horizon model of legislative policy making with endogenous status quo, under general assumptions about the policy space and preferences, and Penn (2009), who studies the formation of "far-sighted" preferences under endogenous status quo when voters must choose, at each stage, between a randomly drawn proposal and the status quo.

to consider non-Markovian strategies, which are essential to generate the implicit logrolling behavior characteristic of our optimal equilibrium. While the benefit of repetition on cooperation has long been acknowledged in the repeated games literature (Mailath and Samuelson, 2006), very few papers have explored the implications for voting behavior and optimal voting rules. This is the main purpose of our paper.

Going back to Buchanan and Tullock (1962), it has been argued that making multiple, as opposed to a single, collective decisions opens up the possibility of vote trading between committee members. As noted in Buchanan and Tullock (1962), "The existence of a logrolling process is central to our general analysis of simple majority voting". The literature on logrolling usually considers the possibility of vote trading over a finite and explicit set of alternatives under complete information (Park, 1967; Casella and Palfrey, 2019; Casella and Macé, 2021). An important take-away from this literature is that logrolling does not necessarily improve welfare, as a mutually beneficial exchange of votes might end up hurting the voters who are not part of the trade. Our model does not assume explicit vote trading agreements between voters. Instead, logrolling emerges endogenously at equilibrium, as an implicit agreement between voters. Increased cooperation can be sustained at equilibrium because of the threat of reverting to an inefficient equilibrium in case of failure to achieve the efficient decision. In contrast with most of the existing literature on vote trading, logrolling in our model is associated with welfare improvements.

A few papers from the logrolling literature are particularly connected to our study. In the experiments reported in Fischbacher and Schudy (2014, 2020), logrolling also arises from a dynamic sequence of votes and can be explained by social preferences. Voters that have benefited from non-sincere voting from other voters in the past reciprocate when given the opportunity. In turn, trusting voters may vote non-sincerely at the beginning of the sequence, anticipating that their fellow committee members will reciprocate. By contrast, logrolling is obtained in our model as an equilibrium outcome of the repeated voting game. This feature is also present in Carrubba and Volden (2000), who consider a sequential model of logrolling embedded in a repeated model of coalition formation in a legislature. At each period a winning coalition of legislators is picked endogenously and a series of proposals, each benefiting only one of the coalition's members (at the expense of all other legislators), is put to a vote. The voting rule is chosen endogenously by the legislators at the beginning of the super game. The objective of the model is to show why legislators may form oversized coalitions (larger than the majority threshold) to enter into logrolling agreements. The result relies on the legislators having a smaller incentive to deviate once their own bill has been passed if rejecting a proposal requires the defection of more than one legislator in the coalition. In contrast with Carrubba and Volden (2000), logrolling in our model is implicit (no explicit agreement between individuals within a specific coalition), and the intensity of individual preferences varies across proposals. Furthermore, we characterize the optimal equilibrium, as opposed to only considering a specific kind of equilibrium.

To the best of our knowledge, the sole paper that has considered the implications of logrolling for the choice of a voting rule is Charroin and Vanberg (2021). They study a classical logrolling model over a finite set of decisions and restrict attention to the comparison between simple majority and unanimity, when there are only three voters. Relying on simulations and lab experiments, they obtain that logrolling is welfare-improving under unanimity but not necessarily under simple majority. As in our model, they show that unanimity may outperform simple majority. Yet, the results we obtain in the current paper are more general in the following sense : we consider any number of voters and any super-majority requirement, and we explicitly characterize the circumstances under which a given majority rule is optimal on the stylized (yet broad) model. Moreover, by resorting to analytical results, we are able to highlight a precise mechanism (the truncated rule of the optimal equilibrium) underlying the comparison of voting rules.

Finally, the paper closest to ours is Maggi and Morelli (2006), who study self-enforcing voting rules in international organizations. A group of countries takes repeated collective actions under the assumption that an action is only effective if taken by all countries (what they call a *pure collective action*). Before taking action, countries engage in a cheap talk stage where they announce whether they are in favor or against the proposed action. At the optimal symmetric public perfect equilibrium, countries reveal their preference truthfully and take action (unanimously) if the number of countries in favor exceeds a fixed threshold. The decision rule on the equilibrium path thus mimics the outcome of a (super)majority rule under enforceable decisions. Similar to our model, Maggi and Morelli (2006) consider repeated decisions over exogenous and independent proposals, and strategies may depend on the past (public) history of collective actions and votes. However, our model differs from theirs in three important respects. First, we consider a model where decisions can be enforced. The voting rule (majority threshold) matters in an explicit way, and a winning coalition can impose its preference on the losing minority. Second, we consider a model where members' preferences have varying intensities. This aspect is crucial in many environments and greatly exacerbates the benefits of logrolling. Third, we assume complete information.[5] Similar to Maggi and Morelli (2006), we find that it is possible to achieve the first-best if the discount factor is large enough. However, our analysis of the optimal equilibrium differs substantially.[6]

The rest of the paper is organized as follows. Section 2 introduces the model and lays out the main assumptions. Section 3 characterizes the set of equilibrium payoffs and the optimal consensus probability. Section 4 revisits our general results on a more stylized model to perform comparative statics. Section 5 discusses our main assumptions. Section 6 concludes.

2 Setup

2.1 Stage game

Model - A group of individuals $N = \{1, \ldots, n\}$ must collectively decide whether to adopt a proposed *reform* or keep the status quo. If enacted, the reform yields utility $u_i \in \mathbb{R}$ to individual $i \in N$, while the status quo's utility is normalized to 0.

[5] Note that truthful revelation under incomplete information, as obtained in Maggi and Morelli (2006), would not be possible in an environment where preferences have varying intensities.

[6] Under unanimity rule, our model can be interpreted as a model of repeated collective decisions without enforcement under a pure collective action assumption. Our paper can thus be seen as an extension of Maggi and Morelli (2006)'s analysis of self-enforcing voting rules to preferences that exhibit varying intensities under complete information.

An individual $i \in N$ is thus *favorable* to the reform if $u_i \geq 0$, *opposed* if $u_i < 0$, and we refer to $|u_i|$ as her *stake* in the collective decision. The reform is characterized by the utility vector $\mathbf{u} = (u_i)_{i \in N}$, with mean $\overline{u} = (\sum_{i \in N} u_i)/n$. We say that a reform is *good* when $\overline{u} \geq 0$ and *bad* when $\overline{u} < 0$.

The reform $\mathbf{u}$ is drawn from a cumulative distribution function G, whose support is denoted by $S \subseteq \mathbb{R}^N$. We make the following three assumptions on G.

Assumption 1. G is symmetric: for any $\mathbf{u} \in S$ and any permutation π of N, $G(\mathbf{u}) = G(\mathbf{u}_\pi)$.

Assumption 2. G is smooth and its support S is bounded and convex.

Assumption 3. $\mathbb{E}_G[|\overline{u}|] < +\infty$.

Assumption 1 reflects the fact that individuals are ex-ante identical, but ex-post heterogeneous. Assumption 2 is made for ease of exposition, it ensures in particular that the distribution G has no atom. Finally, Assumption 3 guarantees that the game has well-defined expected payoffs.

After observing reform $\mathbf{u}$, each individual i votes either in favor of the reform, $\mathrm{v}_i = 1$, or against, $\mathrm{v}_i = 0$. A *majority rule* with threshold $k \in \{1/2, \dots, 1\}$ then decides the reform's fate $\mathrm{d} \in \{0, 1\}$. If at least kn individuals vote in favor, the reform is adopted, $\mathrm{d} = 1$, and each individual i gets utility u_i. If not, the status quo remains, $\mathrm{d} = 0$, and each individual gets utility 0.

Strategies and Equilibrium - A *voting strategy* v_i for player i associates to any reform $\mathbf{u}$ a vote $v_i(\mathbf{u}) \in \{0, 1\}$. At the unique Nash equilibrium in weakly undominated strategies,[7] every individual votes sincerely, i.e. $v_i(\mathbf{u}) = \mathbb{1}\{u_i \geq 0\}$. As a result, a reform is collectively accepted if and only if the individual with the $[kn]$-th largest utility is favorable, i.e. $u_{[kn]} \geq 0$. This individual plays a central role in our analysis as she acts as a pivot, both in the static and in the repeated game. Henceforth, we refer to her as the *critical individual*, and to her utility $u_c := u_{[kn]}$ as the *critical utility*. We note that the equilibrium is usually inefficient, as bad reforms may be accepted (when $u_c \geq 0$ but $\overline{u} < 0$) and good reform may be rejected (when $u_c \leq 0$

[7] Individuals with utility $u_i = 0$ are indifferent between voting in favor or against, but this instance almost never arises under Assumption 2. The equilibrium in undominated strategies is thus *essentially* unique.

but $\overline{u} > 0$) under sincere voting.

Decision rules - Given majority rule k, a strategy profile $(v_i)_{i \in N}$ induces a (group) decision rule. Formally, a *decision rule* d associates to any reform $\mathbf{u}$ a collective decision $d(\mathbf{u}) \in \{0, 1\}$. The expected utility of individual i under decision rule d is given by:

$$U_i(d) = \int u_i d(\mathbf{u}) \, dG(\mathbf{u}). \tag{1}$$

When the rule d is symmetric, we simply denote by $U(d)$ the common expected utility. We denote by d^0 the *sincere* decision rule induced by the stage-game equilibrium, $d^0(\mathbf{u}) = \mathbb{1}\{u_c \geq 0\}$.

Cost of Implementation - For any reform $\mathbf{u} \in S$ and any collective decision $\mathrm{d} \in \{0, 1\}$, let $c(\mathbf{u}, \mathrm{d})$ denote the smallest transfer that would make the critical individual favor decision d,

$$c(\mathbf{u}, \mathrm{d}) = \begin{cases} 0 & \text{if } \mathrm{d} = d^0(\mathbf{u}) \\ |u_c| & \text{if } \mathrm{d} \neq d^0(\mathbf{u}). \end{cases}$$

We refer to $c(\mathbf{u}, \mathrm{d})$ as the *cost of implementing* decision d on reform $\mathbf{u}$. It is equal to 0 if the critical individual favors d and to the *critical stake* $|u_c|$ if she does not. The cost $c(\mathbf{u}, \mathrm{d})$ reflects the smallest common reward needed to incentivize the required majority of k individuals to favor decision d on reform $\mathbf{u}$.[8] For any decision rule $d(.)$, we define $C(d)$ and $\Delta(d)$ as, respectively, the expected and the largest cost of implementing decision $d(\mathbf{u})$,

$$C(d) = \mathbb{E}\left[c(\mathbf{u}, d(\mathbf{u}))\right] \quad \text{and} \quad \Delta(d) = \sup_{\mathbf{u} \in S} c(\mathbf{u}, d(\mathbf{u})).$$

Note that $C(d^0) = \Delta(d^0) = 0$ since the critical individual always agrees with d^0.

Two benchmark rules - Two decision rules play an important role in our characterization of the repeated game's equilibrium payoffs. First, we define the *efficient* rule d^e

[8]The focus on a common (rather than personalized) reward is a consequence of Assumption 5 below.

as the rule which accepts reforms if and only if they are good, i.e. $d^e(\mathbf{u}) = \mathbb{1}\{\overline{u} \geq 0\}$. In what follows, we focus on the interesting case where the stage-game is inefficient, i.e. $U^e := U(d^e) > U^0 := U(d^0)$.

Second, we define the *cost-adjusted inefficient rule* d^i as the rule which selects the inefficient decision if and only if the inefficiency exceeds the cost of implementing the corresponding decision: $d^i(\mathbf{u}) = 1 - d^e(\mathbf{u})$ if and only if $c(\mathbf{u}, 1 - d^e(\mathbf{u})) \leq |\overline{u}|$. We note that d^e maximizes $U(.)$ over all symmetric rules, while d^i minimizes $U(\cdot) + C(\cdot)$ over all symmetric rules. In the sequel, we will denote the largest costs of implementation of the two benchmark rules by $\Delta^e = \Delta(d^e)$ and $\Delta^i = \Delta(d^i)$.

We illustrate decision rules d^0, d^e and d^i on Figure 1. A reform is identified graphically by the average utility $\overline{u}$ and the critical utility u_c. The colored area represents the (projected) support S of G. Good reforms are represented in green, while bad reforms are represented in red. Collectively accepted proposals are located inside the black polygon. Inefficient decisions, either proposals inefficiently rejected or inefficiently accepted, are represented in darker tones. Reforms in hatched areas have a positive cost of implementation, i.e. $c(\mathbf{u}, d(\mathbf{u})) = |u_c|) > 0$, while reforms outside of hatched areas have a zero cost of implementation.

Figure 1: Three decision rules.

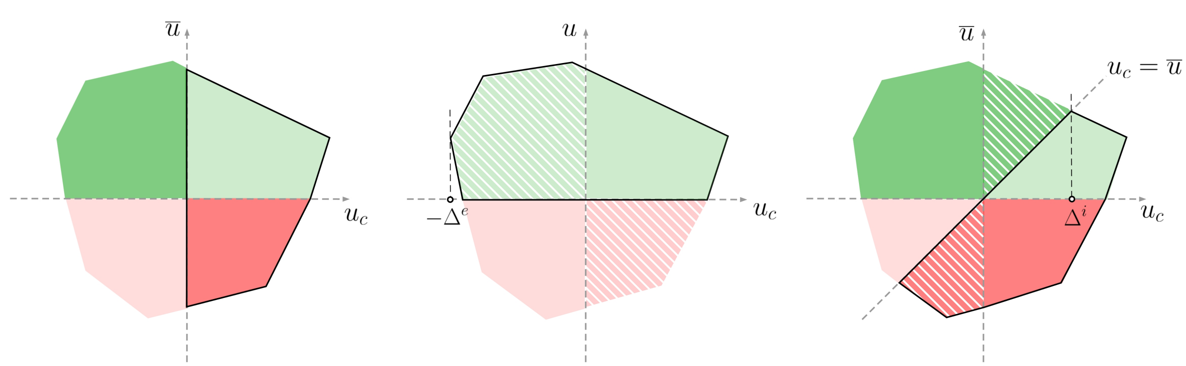

(a) Sincere rule d^0. (b) Efficient rule d^e. (c) Cost-adjusted inefficient rule d^i.

2.2 Repeated Game

We now consider an infinitely repeated version of the stage game.

Timing - At each stage, a reform $\mathbf{u}$ is drawn from G independently of previous stages. The reform $\mathbf{u}$ is publicly observed, then individuals simultaneously vote under majority rule k to decide the reform's fate. The *history* at time t, denoted by h^t, consists of all reforms and votes prior to t. A *strategy* σ_i for individual i associates to every history h and reform $\mathbf{u} \in S$ a vote $\sigma_i(h, \mathbf{u}) \in \{0, 1\}$. Utilities are discounted with a discount factor $\delta \in (0, 1)$. We say that a history h is *on the path* of a strategy profile $\sigma = (\sigma_i)_{i \in N}$ if the votes at each period are the ones specified by σ given the utility realizations.[9]

Assumptions - We restrict our analysis to strategy profiles that satisfy the following three properties.

Assumption 4 (Symmetry)**.** All individuals use the same strategy.

Assumption 5 (Independence of Individual Votes (IIV))**.** Voting behavior only depends on past anonymized realizations of utilities and decisions.

Assumption 6 (As-if-pivotal Voting)**.** Taking as given future play prescribed by the strategy profile, players play as if they were pivotal.

Symmetry is a natural assumption in our context since the model is fully symmetric ex-ante. IIV requires that strategies only depend on past decisions, and not on past individual votes. This implies that deviations from the equilibrium path may only be punished collectively. IIV can be justified by a reluctance to single out or antagonize specific individuals for their votes.[10] Similar requirements have been used in the literature to allow for history-dependent strategies, without losing too much tractability (Bernheim and Slavov, 2009; Anesi and Seidmann, 2015). As-if-pivotal Voting allows us to pin down the voting behavior of individuals who are

[9] Such histories are also referred to in the literature as consistent histories (Mailath and Samuelson (2006)).

[10] In a previous version of the paper, we showed that IIV can be weakened to the requirement of *Anonymity*, i.e. that voting behavior only depends on past anonymized realizations of utilities and votes (voting behavior may then be conditioned on the fraction of individuals voting in favor of the reform). We decided to require the slightly stronger property for ease of exposition.

not pivotal, which is left unconstrained by subgame perfection. Individuals then vote for the alternative that maximizes their continuation utility. The requirement is commonly used in dynamic voting games (see for instance Ali et al. (2022)), it was first proposed in Baron and Kalai (1993) as *stage-undomination*.[11]

Equilibrium - We focus on subgame perfect equilibria that satisfy the above three assumptions, and simply refer to them as *equilibria*. Denoting by $\underline{v}(\delta)$ and $\overline{v}(\delta)$ the lowest and highest equilibrium payoffs,[12] respectively, we define,

$$b_\delta := \frac{\delta}{1-\delta}\Big(\overline{v}(\delta) - \underline{v}(\delta)\Big) \tag{2}$$

as the *power of intertemporal incentives.* Parameter b_δ reflects how costly a decision can be implemented at equilibrium in the first stage, given that continuation promises must themselves be equilibrium payoffs and future is discounted by δ.

3 Equilibrium

We start by introducing truncated decision rules, which play an important role in our characterization of equilibrium payoffs.

3.1 Truncated Rules

We have highlighted in the previous section that decision rules sometimes involve positive implementation costs, which might be prohibitive in the repeated game if intertemporal incentives are not powerful enough. This consideration leads us to the following definition. For any decision rule d and any cutoff $b \in \mathbb{R}_+$, the *truncated rule* d_b selects decision $d(\mathbf{u})$ unless its cost of implementation is greater than b, in

[11] A formal definition of Stage-undomination is given in the appendix. Note that when strategies do not depend on the detail of past individual votes, but only on the resulting collective decision (as under IIV), imposing Stage-Undomination does not affect the set of equilibrium utilities.

[12] The set of equilibrium payoffs is a compact interval, see Lemmas 1 to 3.

which case it selects the sincere decision,

$$d_b(\mathbf{u}) = \begin{cases} d(\mathbf{u}) & \text{if } c(\mathbf{u}, d(\mathbf{u})) \leq b \\ d^0(\mathbf{u}) & \text{if } c(\mathbf{u}, d(\mathbf{u})) > b. \end{cases}$$

For any cutoff $b \in \mathbb{R}_+$, let $\mathcal{D}_b$ denote the set of all symmetric decision rules truncated by b. A symmetric decision rule can be implemented at equilibrium in the first stage if and only if it belongs to $\mathcal{D}_{b_\delta}$. Indeed, as long as the cost of implementation is smaller than b_δ, the critical individual may be incentivized to vote in any desired direction by rewarding compliance with continuation promise $\overline{v}(\delta)$ and punishing deviations with continuation promise $\underline{v}(\delta)$.

Two truncated rules play a central role in our characterization of equilibrium payoffs. Under the *truncated efficient rule* d^e_b, the decision taken is efficient unless its cost of implementation exceeds b. That is, good reforms are accepted unless the critical utility is too small ($u_c \leq -b$), while bad reforms are rejected unless the critical utility is too high ($u_c \geq b$). When $b = 0$, the rule d^e_b coincides with the sincere rule d^0 as no decision with positive implementation cost is allowed. For $b \in (0, \Delta^e)$, the rule d^e_b coincides with the sincere rule whenever it prescribes efficient decisions, while reversing some, but not all, inefficient decisions. For $b \geq \Delta^e$, the truncation becomes moot (since Δ^e is the highest implementation cost of an efficient decision) and d^e_b coincides with the efficient rule. In terms of utility, $U(d^e_b)$ is increasing in b, going from U^0 when $b = 0$ to the first best U^e when $b \geq \Delta^e$.

Under the *truncated cost-adjusted inefficient rule* d^i_b, the decision taken on reform $\mathbf{u}$ is inefficient unless its cost of implementation exceeds $\overline{u}$ or exceeds b. That is, good reforms are rejected unless the critical utility is too high ($u_c \geq \min(\overline{u}, b)$), while bad reforms are accepted unless the critical utility is too low ($u_c \leq \max(\overline{u}, -b)$). As the cutoff b increases, more inefficient decisions are taken. The utility $U(d^i_b)$ thus decreases with b, and the rule d^i_b coincides with d^i whenever $b \geq \Delta^i$.

We illustrate the two truncated rules on Figure 2 below.

Figure 2: Truncated decision rules.

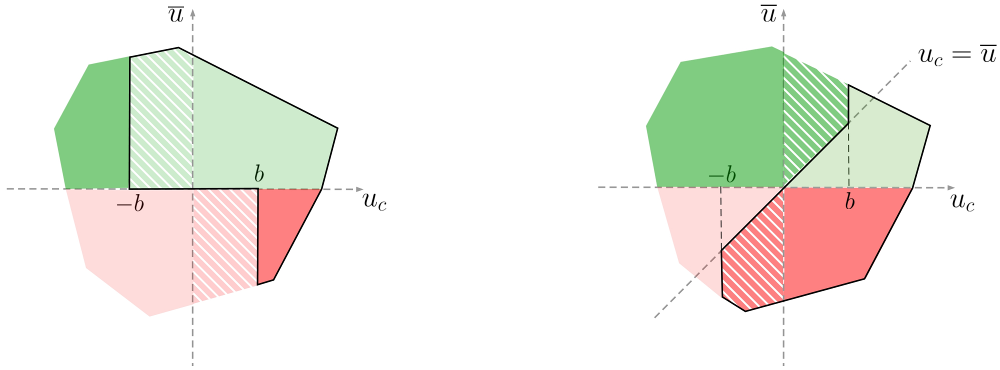

(a) Truncated efficient rule d_b^e. (b) Truncated cost-adjusted inefficient rule d_b^i.

3.2 Characterization of equilibrium payoffs

For any cutoff $b \in \mathbb{R}_+$, let $\overline{U}(b)$ and $\underline{U}(b)$ be such that,

$$\overline{U}(b) := \max_{d \in \mathcal{D}_b} U(d) \quad \text{and} \quad \underline{U}(b) := \min_{d \in \mathcal{D}_b} \Big(U(d) + C(d) \Big).$$

In words, $\overline{U}(b)$ is the highest payoff achievable by a b-truncated rule, while $\underline{U}(b)$ is the lowest cost-augmented payoff achievable by a b-truncated rule. Since both optimizations can be performed reform by reform, it follows that $\overline{U}(b)$ and $\underline{U}(b)$ can be implemented, respectively, by the truncated efficient rule d_b^e and the truncated cost-adjusted inefficient rule d_b^i,

$$\overline{U}(b) = U(d_b^e) \quad \text{and} \quad \underline{U}(b) = U(d_b^i) + C(d_b^i).$$

Furthermore, because set $\mathcal{D}_b$ expands when b increases, we get that $\overline{U}(b)$ increases with b, while $\underline{U}(b)$ decreases with b.

Theorem 1. *The set of equilibrium payoff is equal to $[\underline{U}(b_\delta), \overline{U}(b_\delta)]$, where the power of intertemporal incentives b_δ is given by:*

$$b_\delta = \max \Big\{ b \geq 0 \;\Big|\; (1-\delta)b = \delta \Big[\overline{U}(b) - \underline{U}(b) \Big] \Big\}. \tag{3}$$

Theorem 1 characterizes the set of equilibrium payoffs in the repeated game through the power of intertemporal incentives b_δ, defined as a solution to fixed-point equation:

$$(1-\delta)b_\delta = \delta\left[\overline{U}(b_\delta) - \underline{U}(b_\delta)\right]. \tag{4}$$

The intuition for Theorem 1 follows from the dual relationship between the power of intertemporal incentives b_δ and the extremal equilibrium payoffs $\underline{v}(\delta)$ and $\overline{v}(\delta)$.[13] First, the power b_δ can be simply expressed from the extremal payoffs, as in (2). Second, to see how extremal payoffs depend on power b_δ, recall that any decision rule in $\mathcal{D}_{b_\delta}$ can be implemented at the first stage of the repeated game. Now, the key observation is that there is no on-path cost at incentivizing high payoffs, since high payoffs can be sustained by high rewards (on-path) and low punishments payoffs (off-path). Hence, the highest equilibrium payoff that can then be sustained is simply given by the payoff-maximizing decision rule in $\mathcal{D}_{b_\delta}$, that is $\overline{v}(\delta) = \overline{U}(b_\delta) = U(d^e_{b_\delta})$.

By contrast, there may be on-path costs at incentivizing low payoffs, since low payoffs may require high rewards which are then realized on-path. As a result of this trade-off, the lowest equilibrium payoff is obtained by minimizing the cost-augmented payoff $U(\cdot)+C(\cdot)$ on $\mathcal{D}_{b_\delta}$, and we obtain $\underline{v}(\delta) = \underline{U}(b_\delta) = U(d^i_{b_\delta}) + C(d^i_{b_\delta})$. This intertemporal utility can be achieved by a non-stationary equilibrium which implements $d^i_{b_\delta}$ at the first stage. Intuitively, to achieve the lowest equilibrium payoff, an inefficient decision is taken at the first stage only if its inefficiency exceeds its implementation cost, which then realizes on the equilibrium path.

3.3 Equilibrium outcomes

In this section, we derive a number of implications of Theorem 1 on equilibrium outcomes.

Corollary 1. *The optimal equilibrium payoff can be achieved by a stationary equilibrium in which the truncated efficient rule $d^e_{b_\delta}$ is implemented at each stage.*

Corollary 1 implies that the optimal equilibrium may be understood as a (self-

[13]The proof is provided in the appendix, it relies on the technique of self-generation, see for instance Mailath and Samuelson (2006).

enforcing) *cooperation norm*,[14] whereby voters routinely abide by collectively efficient decisions unless the critical stake exceeds cutoff b_δ. In turn, the parameter b_δ may be interpreted as the optimal *degree of cooperation* that is achievable in the repeated game.

Corollary 2. *The optimal degree of cooperation b_δ is weakly increasing with δ.*

A consequence of Corollary 2 is that the set of equilibrium payoffs expands with δ. Then, with a higher discount factor δ, a higher degree of cooperation can be achieved for two reasons: because voters attach a greater importance to the future (relative to the present) and because the available rewards and punishments become richer.

The next result addresses the following questions: Can even the first best be achieved? If not, can there by any cooperation? We say that *partial cooperation* is achieved when an equilibrium delivers a utility strictly higher than U^0. We say that *full cooperation* is achieved when an equilibrium delivers the first-best utility U^e.

Corollary 3. *There are thresholds δ^P, δ^F with $0 < \delta^P \leq \delta^F < 1$ such that:*

- *if $\delta < \delta^P$, the unique equilibrium payoff is U^0, the stage-game equilibrium payoff.*
- *if $\delta \in [\delta^P, \delta^F)$, only partial cooperation is possible at equilibrium.*
- *if $\delta < \delta^P$, full cooperation is possible at equilibrium.*

Corollary 3 characterizes when either partial or full cooperation can be achieved at equilibrium. For intermediate values of the discount factor, the optimal cooperation norm coincides with a (strictly) truncated efficient rule. For higher values of δ, even the efficient rule can be sustained.

3.4 Consensus

While the previous section focuses on equilibrium outcomes at the group level, we now derive implications of Theorem 1 for individual voting behavior. We focus on optimal equilibria and ask how much *consensus* they can generate.

[14] In fact, this stationary path constitutes the unique optimal equilibrium path.

For a stationary equilibrium σ, we let v^σ denote the voting profile used at each stage. The *consensus probability* under σ is then defined as the probability that a reform is approved unanimously, that is:

$$P_C(\sigma) = \mathbb{P}\left(\prod_{i\in N} v_i^\sigma(\mathbf{u}) = 1\right).$$

We focus on the *optimal consensus probability* which is defined as the highest consensus probability that can be achieved at an optimal equilibrium: if $\Sigma^*(\delta)$ is the set of stationary optimal equilibria, we define:[15]

$$P^*(\delta) = \max_{\sigma\in\Sigma^*(\delta)} P_C(\sigma).$$

Theorem 2. *The optimal consensus probability $P^*(\delta)$ is weakly increasing with δ. There is $\delta^C \in (0,1)$ such that the optimal consensus probability is maximal for $\delta \geq \delta^C$. All good reforms are then accepted with consensus.*

When the discount factor is higher, the higher power of intertemporal incentives allows to construct optimal equilibria with worse punishments. The increase in the optimal consensus probability thus reflects two effects. First, the optimal degree of cooperation is higher and thus more good reforms are approved.[16] Second, any good reform is more likely to receive consensus, as harsher punishments can incentivize even the least favorable voter to vote in favor of the reform.

The first substantive lesson of Theorem 2 is that consensus becomes easier to achieve when voters become more patient. This refines the observation made in Corollary 2 on self-enforcing cooperation norms. When δ increases, it becomes possible to implement cooperation norms such that collective decisions are better on average and also such that reforms are more likely to be approved unanimously. Moreover, if the discount factor is high enough, i.e. $\delta \geq \max(\delta^F, \delta^C)$, a cooperation norm can be sustained whereby only good reforms are approved, and all are

[15]We note that the highest consensus probability at an equilibrium of the repeated game might be attained for non-optimal equilibria. Here, we treat consensus as a secondary objective that comes after payoff maximization. While not exhaustive, we believe that this approach still provides a reasonable and tractable benchmark to highlight how consensus can emerge in the repeated game.

[16]This also means that less bad reforms are approved, but as we note in the proof, these reforms never receive a consensual approval.

approved unanimously.

4 Comparative statics

In this section, we derive rich comparative statics in a simplified model. We focus on a stylized class of preference distributions where (i) individual utilities write as the weighted sum of a common and a private component, and (ii) the distribution of private components is fixed across reforms. The model illustrates how repetition may significantly affect the relative efficiency of alternative majority rules.

4.1 Preference distributions

We assume that the utility of each individual $i \in N$ now writes as,

$$u_i = \theta + \alpha\varepsilon_i$$

where (i) the ordered vector of private components is fixed across reforms: there exist $e_1 \leq \cdots \leq e_n$ such that $\varepsilon_{[j]} = e_j$ for all $j \in \{1, \ldots, n\}$, and (ii) the common component θ is drawn uniformly on $[-1, 1]$. As in the general model, we assume that the distribution of $\mathbf{u} = (u_i)_{i \in N}$ is symmetric, so that any voter has an equal chance to have a private component at any given rank of the distribution $\mathbf{e} = (e_i)_{i \in N}$. Parameter α reflects the (relative) weight attached to the private component, which we interpret as a measure of *diversity* within the committee. For ease of exposition, we assume that $e_1 = -1$ and $e_{n/2} < \overline{e} = 0$.

In this setting, the difference between the average and the critical utility is constant, as $\overline{u} - u_c = -\alpha e_c$, while the average utility, $\overline{u} = \theta$, varies across reforms. This feature greatly simplifies the computation of the optimal equilibrium, allowing us to study the effects of the majority threshold and of the degree of diversity on equilibrium outcomes.

A proposal is good if and only if the common component is positive, $\theta \geq 0$, while it is collectively accepted at the stage-game equilibrium if and only if the critical individual is in favor, $\theta \geq -\alpha e_c$. As $e_{n/2} < 0$, we have $e_c \leq e_{n/2} < 0$ for any majority rule, so that the inefficiency of the stage-game equilibrium comes from good reforms $\theta \in [0, -\alpha e_c]$ being rejected.

Under any voting rule k, the expected payoff under the efficient and sincere decision rules are now equal to:

$$U^e = \frac{1}{4} \quad \text{and} \quad U^0(k) = U^e - \frac{(\Delta(k))^2}{4}, \tag{5}$$

where we abuse notation to write $\Delta(k) := \Delta^e(k) = \alpha|e_c| = \alpha|e_{n-k+1}|$. Here, the parameter $\Delta(k)$ can be interpreted as a measure of the *inefficiency* associated with majority rule k in the stage game. In particular, we get that the equilibrium utility in the stage-game is decreasing with threshold k.

Proposition 1. *In the stage game of the simplified model:*

- *the optimal rule is simple majority (*$k = 1/2$*)*
- *the worst rule is unanimity (*$k = 1$*).*

We note that, although majority rule is optimal, it does not achieve the first best since $\Delta^e(1/2) = |\alpha e_{n/2}| > 0$ and thus $U^0(1/2) < U^e$.

4.2 Optimal voting rule

We turn to the analysis of the repeated game. As we show in the appendix, by applying Theorem 1 for each possible value of the inefficiency Δ, we can obtain closed-form formulas for the optimal degree of cooperation b_δ. In turn, we obtain the partial and full cooperation thresholds δ^F and δ^P, as well as the optimal utility $\overline{v}(\delta)$. Taking stock, our first result characterizes the optimal voting rule, i.e. the rule maximizing the highest equilibrium payoff $\overline{v}(\delta)$.

Theorem 3. *In the simplified model, there are thresholds* $\underline{\delta}, \overline{\delta}$ *with* $0 < \underline{\delta} < \overline{\delta} < 1$ *and continuous functions* $k^*, k^{**} : [0, 1] \to [1/2, 1]$ *such that:*

- *if* $\delta < \underline{\delta}$*, the optimal rule is simple majority (*$k = 1/2$*)*
- *if* $\underline{\delta} \leq \delta < \overline{\delta}$*, the optimal rule is a (super-) majority* $k^*(\delta)$
- *if* $\delta \leq \overline{\delta}$*, any majority rule* k *with* $k^*(\delta) \leq k \leq k^{**}(\delta)$ *is optimal.*

Moreover, $k^*(\cdot)$ is weakly decreasing, from $k^*(\underline{\delta}) > 1/2$ up to $k^*(1) = 1/2$, while $k^{**}(\cdot)$ is weakly increasing, from $k^{**}(\overline{\delta}) = k^*(\overline{\delta})$ up to $k^{**}(1) = 1$.

Theorem 3 shows that accounting for the repetition of collective decisions can drastically affect the assessment of alternative majority rules. In the benchmark case where the discount factor is low ($\delta < \underline{\delta}$), future is so discounted that the stage-game result of Proposition 1 applies, i.e. simple majority ($k = 1/2$) is the optimal rule. In the opposite benchmark where the discount factor is high ($\delta \geq \overline{\delta}$), intertemporal incentives are sufficient to achieve full cooperation for some rules ($k^*(\delta) \leq k \leq k^{**}(\delta)$), and eventually for all rules when δ is large enough. This last remark coincides with the standard folk theorem.

The most interesting case of Theorem 3 arises for intermediate discount factors ($\underline{\delta} \leq \delta < \overline{\delta}$). The optimal rule is then a majority k^* that is a strict supermajority at first ($k^*(\underline{\delta}) > 1/2$) and then weakly decreases for higher discount factors. In particular, this shows that a strict supermajority can dominate simple majority when accounting for the repetition of collective decisions, while the opposite conclusion would be drawn by focusing only on the stage game. Moreover, as we show below, a common case is that unanimity becomes uniquely optimal in the repeated game, thus completely overturning the result of Proposition 1.

Corollary 4. *In the simplified model, when $\alpha \leq 1/2$, there is a range of discount factors for which unanimity ($k = 1$) is uniquely optimal.*

4.3 Consensus

Following Section 3.4, we look for the maximal consensus that can be obtained at the optimal equilibrium of the repeated game. This is achieved by providing to each voter the highest possible incentive (given by the power b_δ) to vote in favor of decisions that must be accepted at equilibrium. The optimal consensus probability is thus given by

$$P^*(\delta) = \mathbb{P}(u_{[1]} + b_\delta \geq 0) = \mathbb{P}(\theta \geq \alpha - b_\delta) = \frac{1 + b_\delta - \alpha}{2}$$

whenever $\alpha - b_\delta \in [-1, 1]$. This probability can be computed explicitly in the simplified model by plugging in the formulas for b_δ.

Proposition 2. *The optimal consensus probability P^* can be:*

- *decreasing with diversity α*
- *decreasing with majority rule k.*

Proposition 2 testifies of the complex mapping from parameters to equilibrium outcomes in the repeated game. First, when preferences in the committee are more diverse (α is higher), voting outcomes may be more often unanimous. Hence, the observation of vote tallies may be a poor indicator of preferences within a committee when a norm of cooperation operates. Second, by increasing the majority k (the required threshold to pass a reform), consensus might in fact decrease within the committee. This prediction contrasts with standard models of strategic voting to aggregate dispersed information under common preferences (Maug and Rydqvist, 2009), where voters respond to an increase in the majority threshold by a higher propensity to vote in favor of the reform (so as to counteract the voting rule's bias), thus making consensus more likely.

5 Discussion

Malevolent Behavior - Some of the equilibrium payoffs we characterize in Theorem 1 rely on the possibility of punishing individuals for voting in favor of *efficient* decisions (either rejecting good proposals or accepting bad proposals). One may consider such behavior to be unrealistic, or at the very least undesirable. How much cooperation could be achieved at equilibrium without having to rely on these kind of malevolent incentives? We say that a strategy profile is *non-malevolent* if after any history, implementing the inefficient decision does not lead to a greater continuation utility than implementing the efficient decision. Under the assumption of non-malevolent strategies, the worst equilibrium payoff necessarily coincides with the payoff of the stage-game equilibrium U^0. The set of equilibrium payoffs is then equal to $E^* = [U^0, \overline{U}(b^*_\delta)]$, where cutoff $b^*_\delta \le b_\delta$ is given by,

$$b^*_\delta = \max \left\{ b \ge 0 \mid (1-\delta)b = \delta \left[\overline{U}(b) - U^0 \right] \right\}.$$

The restriction to non-malevolent strategies weakens the optimal punishment (i.e. the lowest equilibrium payoff), thus reducing the largest equilibrium payoff. The set of equilibrium payoffs under non-malevolence E^* is thus always included in the set of equilibrium payoffs E. The inclusion is strict unless there is no cooperation (i.e. $E = \{U^0\}$). Yet, we can show that the result on the optimal majority rule in the simplified model (Theorem 3) remains valid under non-malevolence.

Dropping IIV - If strategies can be conditioned on the detailed history of past votes (as opposed to only past collective decisions), then deviations can be punished individually. In some cases, this may lead to worst equilibrium outcomes than the ones we characterize under IIV. In that sense, our analysis can be understood as providing a lower bound on the maximal equilibrium payoff.

Dropping Symmetry - Our analysis describes the emergence of a norm of cooperation in a stylized setting where committee members are ex-ante symmetric. This is a substantive assumption that allows the precise characterization of equilibrium outcomes (Theorem 1) and optimal voting rules (Theorem 3). In many committees, voters are often organized, formally or not, in groups (e.g., political parties in assemblies) which share similar preferences. While the precise analysis seems more difficult in this context, we may speculate that similar cooperation norms can emerge in two ways. First, a cooperation norm can emerge within groups, voters sometimes voting against their will to better satisfy the overall preference of the group, if a form of symmetry in preferences exists within the group. Second, a cooperation norm can also be sustained across groups, provided that "groups' preferences" are sufficiently symmetric.

Incomplete information - Our model assumes complete information in the stage game. This assumption is essential for individuals to identify good proposals and punish deviations in order to sustain better outcomes at the equilibrium of the repeated game. Yet, we can show a property of continuity : if the incompleteness of information is low enough, the norm of cooperation that we showcase in our analysis can still emerge in the repeated game.

6 Conclusion

In this paper, we propose a general model of repeated voting in committees and study equilibrium outcomes under alternative majority rules. In contrast with most of the existing literature, which usually focuses on repeated decisions over a single issue under Markovian strategies, we consider a committee making decisions over issues of varying nature, while allowing for history-dependent strategies. Our characterization reveals a complex relationship between the majority rule, the preference distribution, and the set of equilibrium payoffs. We thus turn to a simplified preference domain to perform various comparative statics. Our findings provide theoretical support for two important stylized facts about collective-decision-making in committees: the ubiquity of unanimity as a formal voting rule, and the prevalence of consensus decision-making.

References

Ali, S. N., Bernheim, B. D., Bloedel, A. W., and Battilana, S. C. (2022). Who controls the agenda controls the polity. *arXiv preprint arXiv:2212.01263*.

Anesi, V. and Seidmann, D. J. (2015). Bargaining in standing committees with an endogenous default. *Review of Economic Studies*, 82:825–867.

Baron, D. (1996). A dynamic theory of collective goods programs. *The American Political Science Review*, 90(2):316–330.

Baron, D. and Kalai, E. (1993). The simplest equilibrium of a majority-rule division game. *Journal of Economic Theory*, 61(2):290–301.

Bernheim, B. D. and Slavov, S. N. (2009). A solution concept for majority rule in dynamic settings. *Review of Economic Studies*, 76:33–62.

Buchanan, J. and Tullock, G. (1962). The calculus of consent.

Carrubba, C. J. and Volden, C. (2000). Coalitional politics and logrolling in legislative institutions. *American Journal of Political Science*, 44:261–277.

Casella, A. and Macé, A. (2021). Does vote trading improve welfare? *Annual Review of Economics*, 13:57–86.

Casella, A. and Palfrey, T. (2019). Trading votes for votes. a dynamic theory. *Econometrica*, 87:631–652.

Charroin, L. and Vanberg, C. (2021). Logrolling affects the relative performance of alternative q-majority rules. *Working paper.*

Duggan, J. and Kalandrakis, A. (2012). Dynamic legislative policy making. *Journal of Economic Theory*, 147(5):1653–1688.

Dzuida, W. and Loeper, A. (2018). Dynamic pivotal politics. *American Political Science Review*, 112(3):580–601.

Fischbacher, U. and Schudy, S. (2014). Reciprocity and resistance to comprehensive reform. *Public Choice*, 160(3):411–428.

Fischbacher, U. and Schudy, S. (2020). Agenda control and reciprocity in sequential voting decisions. *Economic Inquiry*, 58(4):1813–1829.

Kalandrakis, A. (2004). Why not proportional? *Journal of Economic Theory*, 116:294–322.

Maggi, G. and Morelli, M. (2006). Self-enforcing voting in international organizations. *The American Economic Review*, 96(4):1137–1158.

Mailath, G. J. and Samuelson, L. (2006). Repeated games and reputation - long run relationships.

Maug, E. and Rydqvist, K. (2009). Do shareholders vote strategically? voting behavior, proposal screening, and majority rules. *Review of Finance*, 13(1):47–79.

Park, R. (1967). The possibility of a social welfare function: Comment. *The American Economic Review*, 57:1300–1304.

Penn, E. M. (2009). A model of farsighted voting. *American journal of Political Science*, 53(1):36–54.

Urfalino, P. (2014). The rule of non-opposition: Opening up decision-making by consensus. *Journal of Political Philosophy*.

A Proofs

Note : the proofs are incomplete in the current version of the file. All proofs will be typed shortly, and the updated file will be made available online.

A.1 Equilibrium characterization

In the sequel, we associate to any utility vector u the *reduced vector* $\tilde{u} := (u_c, \overline{u})$ and we denote by $\tilde{S}$ the image of the support S of utilities by the projection $u \mapsto \tilde{u}$.

Definition 1. *A payoff $w \in \mathbb{R}$ is decomposable on W if there exists a decision rule $d^w : \tilde{S} \mapsto \{0,1\}$ and continuation payoffs $v_+^w : \tilde{S} \mapsto W$ and $v_-^w : \tilde{S} \mapsto W$ such that:*

$$w = \mathbb{E}[(1-\delta)d^w(\tilde{u})\overline{u} + \delta v_+^w(\tilde{u})] \tag{6}$$

and

$$\forall \tilde{u} = (u_c, \overline{u}) \in \tilde{S}, \qquad (1-\delta)d^w(\tilde{u})u_c + \delta v_+^w(\tilde{u}) \geq (1-\delta)(1-d^w(\tilde{u}))u_c + \delta v_-^w(\tilde{u}). \tag{7}$$

The set W is self-generating *if any $w \in W$ is decomposable on W.*

In words, w is decomposable by (d^w, v_+^w, v_-^w) if (i) it is the expected average payoff given by the rule d^w and on-path continuation promise v_+^w and (ii) for any utility realization $\tilde{u}$, implementing the action $d^w(\tilde{u})$ is rational for the critical individual given a continuation payoff $v_+^w(\tilde{u})$ and a punishment $v_-^w(\tilde{u})$. For the sequel, it is useful to note that the rationality condition (7) is equivalent to:

$$\forall \tilde{u} = (u_c, \overline{u}) \in \tilde{S}, \qquad (1-\delta)(2d^w(\tilde{u})-1)u_c + \delta(v_+^w(\tilde{u}) - v_-^w(\tilde{u})) \geq 0 \tag{8}$$

Lemma 1. *If W is self-generating, then any $w \in W$ is an equilibrium payoff.*

Proof. This proof is adapted from standard techniques exposed in [Mailath and Samuelson, 2006].

Let W be self-generating and let $w^0 \in W$. Consider the automaton defined by:

- the set of states W
- the initial state w^0
- the output function $f : w \mapsto d^w$
- a transition function $\tau : W \times \tilde{S} \times \{0,1\} \to W$, such that

$$\tau(w, \tilde{u}, D) = \begin{cases} v_+^w(\tilde{u}) & \text{if } D = d^w(\tilde{u}) \\ v_-^w(\tilde{u}) & \text{otherwise} \end{cases}$$

Extend the transition function from $W \times \tilde{S} \times \{0,1\}$ to $W \times \tilde{H}$, where $\tilde{H}$ denotes the set of (ex-ante) reduced histories, by recursively defining $\tau(w, \emptyset) = w$ and

$$\tau(w, h^t) = \tau(\tau(w, h^{t-1}), \tilde{u}^t, D^t).$$

Then define the group-strategy γ by $\gamma(h, \tilde{u}) = f(\tau(w^0, h))(\tilde{u})$ and the collective continuation payoff v_+ and punishment v_- by $v_+(h, \tilde{u}) = v_+^{\tau(w^0,h)}(\tilde{u})$ and $v_-(h, \tilde{u}) = v_-^{\tau(w^0,h)}(\tilde{u})$.

Let $\Delta_i(h, \tilde{u})$ be the utility difference in state $\tau(w^0, h)$ between following the group-strategy $\gamma(h, \tilde{u})$ and deviating for an individual with utility u_i (assuming that the individual is pivotal). Formally:

$$\begin{aligned} \Delta_i(h, \tilde{u}) &= (1-\delta)\gamma(h, \tilde{u})u_i + \delta v_+(h, \tilde{u}) - \Big((1-\delta)(1-\gamma(h, \tilde{u}))u_i + \delta v_-(h, \tilde{u})\Big) \\ &= (1-\delta)(2\gamma(h, \tilde{u}) - 1)u_i + \delta(v_+(h, \tilde{u}) - v_-(h, \tilde{u})). \end{aligned}$$

By construction, as W is self-generating, we must have $\Delta_c(h, \tilde{u}) \geq 0$ for any $h \in \tilde{H}$ and $\tilde{u} \in \tilde{S}$.

Now, define individual strategy σ_i by:

$$\forall h \in \tilde{H},\ \forall \tilde{u} \in \tilde{S}, \qquad \sigma_i(h, \tilde{u}) = \begin{cases} \gamma(h, \tilde{u}) & \text{if } \Delta_i(h, \tilde{u}) \geq 0 \\ 1 - \gamma(h, \tilde{u}) & \text{otherwise.} \end{cases}$$

Observe first that the strategy profile $\sigma = (\sigma_i)_{i\in N}$ indeed implements the group-strategy γ:

- if $\gamma(h,\tilde{u}) = 1$, then for any i with $u_i \geq u_c$, we have $\Delta_i(h,\tilde{u}) \geq \Delta_c(h,\tilde{u}) \geq 0$ and thus $\sigma_i(h,\tilde{u}) = \gamma(h,\tilde{u}) = 1$. Therefore $d^\sigma(h)(\tilde{u}) = 1 = \gamma(h,\tilde{u})$.

- if $\gamma(h,\tilde{u}) = 0$, then for any i with $u_i \leq u_c$, we have $\Delta_i(h,\tilde{u}) \geq \Delta_c(h,\tilde{u}) \geq 0$ and thus $\sigma_i(h,\tilde{u}) = \gamma(h,\tilde{u}) = 0$. Therefore $d^\sigma(h)(\tilde{u}) = 0 = \gamma(h,\tilde{u})$.

Then, as W is self-generating, we have by construction that $U(\sigma) = w^0$. It remains to show that σ is an equilibrium. It is easy to see that by construction σ satifies the sufficient conditions for equilibrium □

Lemma 2. *The set of equilibrium payoffs is the largest self-generating set.*

Proof. This proof is adapted from standard techniques exposed in [Mailath and Samuelson, 2006].

By the previous lemma, it suffices to show that the set of equilibrium payoffs $E \subseteq \mathbb{R}$ is a self-generating set. For any payoff $w \in E$ with associated equilibrium σ, we let:

- the decision rule d^w be $d^w(\tilde{u}) = d^\sigma(\tilde{u})$

- the continuation payoff v_+^w be $v_+^w(\tilde{u}) = U(\sigma \mid \tilde{u}, d^\sigma(\tilde{u}))$

- the punishment payoff v_-^w be $v_-^w(\tilde{u}) = U(\sigma \mid \tilde{u}, 1 - d^\sigma(\tilde{u}))$

As σ is a subgame-perfect equilibrium, the promises v_+^w and v_-^w take values in E, as desired. Moreover, we have $w = U(\sigma) = \mathbb{E}[(1-\delta)d^w(\tilde{u})\overline{u} + \delta v_+^w(\tilde{u})]$ by definition of d^w and v_+^w. Finally, it is easy to see that (7) is satisfied since σ is an equilibrium:

- if $d^\sigma(\tilde{u}) = 1$, then there are at least k individuals (where k is the majority rule threshold) for which

$$(1-\delta)u_i + \delta U(\sigma \mid \tilde{u}, 1) \geq \delta U(\sigma \mid \tilde{u}, 0)$$

  and thus

$$(1-\delta)u_c + \delta U(\sigma \mid \tilde{u}, 1) \geq \delta U(\sigma \mid \tilde{u}, 0).$$

  We thus obtain that (7) is satisfied.

- if $d^\sigma(\tilde{u}) = 0$, the proof is similar and thus omitted.

□

Lemma 3. *The set of equilibrium payoffs is a compact interval.*

Proof. **Claim 1:** If $W \subset \mathbb{R}$ is a self-generating set such that $U^0 \in W$, then any payoff $w \in (\inf W, \sup W)$ can be decomposed on W.

To show that, it suffices to show that for any payoff w decomposable on W and any $\alpha \in (0,1)$, the payoff $w' = \alpha w + (1-\alpha)U^0$ is also decomposable on W. Let w be decomposable by a triplet (d, v_+, v_-) on W.
We may write:

$$\begin{aligned} w &= U^0 + (w - U^0) \\ &= U^0 + \mathbb{E}\left[\underbrace{(1-\delta)d(\tilde{u})\overline{u} + \delta v_+(\tilde{u}) - (1-\delta)d^0(\tilde{u})\overline{u} - \delta U^0}_{z(\tilde{u})}\right] \\ &= U^0 + \int_{\tilde{S}} z(\tilde{u}) d\tilde{G}(\tilde{u}). \end{aligned}$$

As G has no atom by assumption, there must exist a subdomain $\tilde{S}' \subset \tilde{S}$ such that $\int_{\tilde{S}'} z(\tilde{u}) d\tilde{G}(\tilde{u}) = \alpha \int_{\tilde{S}} z(\tilde{u}) d\tilde{G}(\tilde{u})$. Let us define (d', v'_+, v'_-) by:

$$d'(\tilde{u}) = \begin{cases} d(\tilde{u}) & \text{if } \tilde{u} \in \tilde{S}' \\ d^0(\tilde{u}) & \text{otherwise} \end{cases}, v'_+(\tilde{u}) = \begin{cases} v_+(\tilde{u}) & \text{if } \tilde{u} \in \tilde{S}' \\ U^0 & \text{otherwise} \end{cases}, v'_-(\tilde{u}) = \begin{cases} v_-(\tilde{u}) & \text{if } \tilde{u} \in \tilde{S}' \\ U^0 & \text{otherwise.} \end{cases}$$

As $U^0 \in W$, it is clear that both v'_+ and v'_- take values in W. Moreover, since d^0 is the stage-game equilibrium, it is clear by construction that the rationality condition (7) is satisfied by (d', v'_+, v'_-). Finally, applying (6), the expected payoff is given by:

$$\begin{aligned} w' &= \int_{\tilde{S}'} \left((1-\delta)d(\tilde{u})\overline{u} + \delta v_+(\tilde{u})\right) d\tilde{G}(\tilde{u}) + \int_{\tilde{S}\setminus\tilde{S}'} \left((1-\delta)d^0(\tilde{u})\overline{u} + \delta U^0\right) d\tilde{G}(\tilde{u}) \\ &= \int_{\tilde{S}'} z(\tilde{u}) d\tilde{G}(\tilde{u}) + U^0 = \alpha(w - U^0) + U^0 = \alpha w + (1-\alpha)U^0, \end{aligned}$$

as desired. Hence $\alpha w + (1-\alpha)U^0$ is decomposable on W, this concludes the proof of Claim 1.

Claim 2: The set of payoffs generated by a compact interval is a compact interval.

Let $W = [\underline{w}, \overline{w}]$ be a compact interval. Let $b = \frac{\delta}{1-\delta}(\overline{w} - \underline{w})$. We define $(d^{\max}, v_+^{\max}, v_-^{\max})$ by:

$$d^{\max}(\tilde{u}) = \mathbb{1}_{\{|u_c| \le b\}} d^e(\tilde{u}) + \mathbb{1}_{\{|u_c| > b\}} d^0(\tilde{u}), \qquad v_+^{\max}(\tilde{u}) = \overline{w}, \qquad v_-^{\min}(\tilde{u}) = \underline{w}.$$

In words, the decision rule is efficient whenever the critical stake is below b, and coincides with the sincere rule otherwise. The continuation payoff is always the highest possible and the punishment payoff is always the lowest possible. Clearly, by definition of b, the triplet $(d^{\max}, v_+^{\max}, v_-^{\max})$ satisfies the rationality condition (7). The payoff induced by the triplet $(d^{\max}, v_+^{\max}, v_-^{\max})$ is given by:

$$w^{\max} = (1-\delta)\left(\int_{|u_c| \le b} d^e(\tilde{u})\overline{u} d\tilde{G}(\tilde{u}) + \int_{|u_c| > b} d^0(\tilde{u})\overline{u} d\tilde{G}(\tilde{u})\right) + \delta \overline{w}.$$

Now, suppose that there is a triplet (d', v'_+, v'_-) which generates a higher payoff $w' > w^{\max}$. As we have for any $\tilde{u} \in \tilde{S}$ that $d'(\tilde{u})\overline{u} \le d^e(\tilde{u})\overline{u}$ and $v'_+(\tilde{u}) \le \overline{w}$, there must exists $\tilde{u}$ with $|u_c| > b$ and such that $d'(\tilde{u})\overline{u} > d^0(\tilde{u})\overline{u}$. Suppose without loss of generality that $\overline{u} > 0$. Then we have $d^0(\tilde{u}) = 0$, and thus $u_c < -b$ and $d'(\tilde{u}) = 1$. Now, we have:

$$(1-\delta)(2d'(\tilde{u}) - 1)u_c + \delta(v'_+(\tilde{u}) - v'_-(\tilde{u})) \le (1-\delta)u_c + \delta(\overline{w} - \underline{w}) < \delta(\overline{w} - \underline{w}) - (1-\delta)b = 0.$$

Thus, the triplet (d', v'_+, v'_-) fails (7). We conclude that $w^{\max}$ is the highest payoff that can be decomposed on W. Consider now the triplet $(d^{\min}, v_+^{\min}, v_-^{\min})$ defined by:[17]

$$d^{\min}(\tilde{u}) = \mathbb{1}_{\{|u_c| \le b\} \cup \{u_c \in (0, \overline{u})\}}(1 - d^e(\tilde{u})) + \mathbb{1}_{\{|u_c| > b\} \cup \{u_c \notin (0, \overline{u})\}} d^0(\tilde{u})$$

and

$$v_+^{\min}(\tilde{u}) = \mathbb{1}_{\{|u_c| \le b\} \cup \{u_c \in (0, \overline{u})\}}\left(\underline{w} + \frac{1-\delta}{\delta}|u_c|\right) + \mathbb{1}_{\{|u_c| > b\} \cup \{u_c \notin (0, \overline{u})\}} \underline{w}, \qquad v_-^{\min}(\tilde{u}) = \underline{w}.$$

In words, the decision rule is inefficient whenever the critical stake is below b and the

[17] In the formula, as in the sequel, we abuse notation by writing $u_c \in (0, \overline{u})$ for $u_c \in (\min(0, \overline{u}), \max(0, \overline{u}))$. Note in particular that $u_c \in (0, \overline{u})$ implies that u_c and $\overline{u}$ have the same sign, i.e. $d^e(\tilde{u}) = d^0(\tilde{u})$.

critical utility is of the same sign as the average utility but with a lower magnitude, and it coincides with the sincere rule otherwise. The continuation payoff is the minimum needed to provide incentives to the critical voter to vote against her will when the prescribed decision is inefficient and insincere, it is the lowest possible payoff otherwise. The punishment payoff is always the lowest possible.

Clearly, by construction, the the triplet $(d^{\min}, v_+^{\min}, v_-^{\min})$ satisfies the rationality condition (7). Now, suppose that there is a triplet (d', v'_+, v'_-) which generates a lower payoff $w' < w^{\min}$, where $w^{\min}$ is the expected payoff generated by $(d^{\min}, v_+^{\min}, v_-^{\min})$. Then, there must exist $\tilde{u}$ such that the following condition holds:

$$(1-\delta)d'(\tilde{u})\overline{u} + \delta v'_+(\tilde{u}) < (1-\delta)d^{\min}(\tilde{u})\overline{u} + \delta v_+^{\min}(\tilde{u}). \tag{9}$$

Condition (9) may only hold if either $d^{\min}(\tilde{u}) = d^e(\tilde{u})$ or $v_+^{\min}(\tilde{u}) > \underline{w}$. We thus consider two cases:

- if $v_+^{\min}(\tilde{u}) > \underline{w}$, then by construction of $d^{\min}$ and $v_+^{\min}$, we have $|u_c| < b$, $u_c \in (0, \overline{u})$ and $d^{\min}(\tilde{u}) = 1 - d^e(\tilde{u}) = 1 - d^0(\tilde{u})$. We thus have $d'(\tilde{u})\overline{u} \geq d^{\min}(\tilde{u})\overline{u}$, so that condition (9) implies that $v'_+(\tilde{u}) < v_+^{\min}(\tilde{u}) = \underline{w} + \frac{1-\delta}{\delta}|u_c|$. There are two subcases to consider:
  - if $d'(\tilde{u}) = d^{\min}(\tilde{u}) = 1 - d^0(\tilde{u})$, we obtain a contradiction by observing that, since $v'_+(\tilde{u}) < v_+^{\min}(\tilde{u})$ and $v'_-(\tilde{u}) \geq \underline{w}$, the triplet $(d^{\min}, v_+^{\min}, v_-^{\min})$ cannot satisfy the rationality condition (7) at $\tilde{u}$ (recall that by construction, $v_+^{\min}(\tilde{u})$ is the smallest continuation payoff that can incentivize a deviation, given the lowest possible punishment of $\underline{w}$).
  - if $d'(\tilde{u}) = d^0(\tilde{u}) = d^e(\tilde{u})$, we obtain a contradiction with (9), as since $u_c \in (0, \overline{u})$, we have:

$$\begin{aligned}(1-\delta)d'(\tilde{u})\overline{u} + \delta v'_+(\tilde{u}) &\geq (1-\delta)d^e(\tilde{u})\overline{u} + \delta\underline{w} \\ &> (1-\delta)\left((1-d^e(\tilde{u}))\overline{u} + |u_c|\right) + \delta\underline{w} = (1-\delta)d^{\min}(\tilde{u})\overline{u} + \delta v_+^{\min}(\tilde{u}).\end{aligned}$$

- if $d^{\min}(\tilde{u}) = d^e(\tilde{u})$, then we have (by construction) $v_+^{min}(\tilde{u}) = \underline{w}$. For (9) to hold, as $v'_+(\tilde{u}) \geq \underline{w}$, we must have $d'(\tilde{u}) = 1 - d^e(\tilde{u})$. We consider two subcases:

- if $|u_c| \leq b$, then we also have (by construction) $u_c \notin (0, \overline{u})$. For the triplet (d', v'_+, v'_-) to satisfy the rationality condition (7), we must have $v'_+(\tilde{u}) - v'_-(\tilde{u}) \geq \frac{1-\delta}{\delta}|u_c|$, and thus $v'_+(\tilde{u}) \geq \underline{w} + \frac{1-\delta}{\delta}|u_c|$. We obtain a contradiction with (9), as since $u_c \notin (0, \overline{u})$, we have:

$$(1-\delta)d'(\tilde{u})\overline{u} + \delta v'_+(\tilde{u}) \geq (1-\delta)\left(d^e(\tilde{u})\overline{u} + |u_c|\right) + \delta \underline{w}$$
$$> (1-\delta)(1-d^e(\tilde{u}))\overline{u} + \delta \underline{w} = (1-\delta)d^{\min}(\tilde{u})\overline{u} + \delta v_+^{\min}(\tilde{u}).$$

- if $|u_c| > b$, then we obtain a contradiction by observing that, since $v'_+(\tilde{u}) - v'_-(\tilde{u}) \leq \overline{w} - \underline{w} = \frac{1-\delta}{\delta}b$, the triplet (d', v'_+, v'_-) cannot satisfy the rationality condition (7) at $\tilde{u}$.

We conclude that $w^{\min}$ is the lowest payoff that can be decomposed on W.

Claim 3: If a bounded interval I is self-generating, then its closure $\overline{I}$ is also self-generating.

If I is self-generating, then any payoff in I is decomposable on I, and thus also on $\overline{I}$. The set of payoffs that can be decomposed on $\overline{I}$ is thus a compact interval (by Claim 2) which contains I. Thus any payoff in $\overline{I}$ can be decomposed on $\overline{I}$, i.e. $\overline{I}$ is self-generating.

To conclude, we know that the set of equilibrium payoffs E contains U^0. As E is self-generating, it must be a (bounded) interval by Claim 1. As E is the largest self-generating set by the previous lemma, it must be closed by Claim 3, and thus compact.

Before stating the next lemma, we introduce the functions $\overline{U}(b)$ and $\underline{U}(b)$ which correspond respectively to the highest and lowest payoffs that can be achieved at equilibrium given a maximum difference of b between continuation promises and punishments (b reflects the power of intertemporal incentives):

$$\overline{U}(b) = \left(\int_{|u_c| \leq b} d^e(\tilde{u})\overline{u} d\tilde{G}(\tilde{u}) + \int_{|u_c| > b} d^0(\tilde{u})\overline{u} d\tilde{G}(\tilde{u})\right)$$

and

$$\underline{U}(b) = \int_{|u_c| \leq b, u_c \in (0, \overline{u})} \left((1-d^e(\tilde{u}))\overline{u} + |u_c|\right) d\tilde{G}(\tilde{u}) + \int_{|u_c| > b \text{ or } u_c \notin (0, \overline{u})} d^0(\tilde{u})\overline{u} d\tilde{G}(\tilde{u}).$$

Note that $\overline{U}$ is increasing while $\underline{U}$ is decreasing.

Lemma 4. *The set of equilibrium payoff can be written as* $E = [\underline{U}(b^*), \overline{U}(b^*)]$ *where* b^* *is characterized by:*

$$b^* = \max\{b \geq 0 \mid (1-\delta)b = \delta(\overline{U}(b) - \underline{U}(b))\}.$$

Proof. We know from the previous lemma that $E = [\underline{w}, \overline{w}]$. Let $w^{\max}$ the maximal payoff that can be decomposed on E. As E is self-generating, we have that $w^{\max} \geq \overline{w}$. Now, if $w^{\max} > \overline{w}$, then this would contradict the fact that E must be the largest self-generating set. Hence, we must have $w^{\max} = \overline{w}$. Similarly, if we note $w^{\min}$ the minimal payoff that can be decomposed on E, we have $w^{\min} = \underline{w}$.
Following the proof of the previous lemma, noting $b = \frac{\delta}{1-\delta}(\overline{w} - \underline{w})$ we have:

$$\begin{aligned}
w^{\max} &= (1-\delta)\left(\int_{|u_c|\leq b} d^e(\tilde{u})\overline{u}d\tilde{G}(\tilde{u}) + \int_{|u_c|>b} d^0(\tilde{u})\overline{u}d\tilde{G}(\tilde{u})\right) + \delta\overline{w} \\
&= \left(\int_{|u_c|\leq b} d^e(\tilde{u})\overline{u}d\tilde{G}(\tilde{u}) + \int_{|u_c|>b} d^0(\tilde{u})\overline{u}d\tilde{G}(\tilde{u})\right) \\
&= \overline{U}(b).
\end{aligned}$$

Similarly, the expected payoff associated to the triplet $(d^{\min}, v_+^{\min}, v_-^{\min})$ can be written:

$$\begin{aligned}
w^{\min} = &\int_{|u_c|\leq b, u_c\in(0,\overline{u})} \left((1-\delta)(1-d^e(\tilde{u}))\overline{u} + \delta(\underline{w} + \frac{1-\delta}{\delta}|u_c|)\right) d\tilde{G}(\tilde{u}) \\
&+ \int_{|u_c|>b \text{ or } u_c\notin(0,\overline{u})} \left((1-\delta)d^0(\tilde{u})\overline{u} + \delta\underline{w}\right) d\tilde{G}(\tilde{u}) \\
= &(1-\delta)\left(\int_{|u_c|\leq b, u_c\in(0,\overline{u})} \left((1-d^e(\tilde{u}))\overline{u} + |u_c|\right) d\tilde{G}(\tilde{u}) + \int_{|u_c|>b \text{ or } u_c\notin(0,\overline{u})} d^0(\tilde{u})\overline{u}d\tilde{G}(\tilde{u})\right) + \delta\underline{w}. \\
= &\int_{|u_c|\leq b, u_c\in(0,\overline{u})} \left((1-d^e(\tilde{u}))\overline{u} + |u_c|\right) d\tilde{G}(\tilde{u}) + \int_{|u_c|>b \text{ or } u_c\notin(0,\overline{u})} d^0(\tilde{u})\overline{u}d\tilde{G}(\tilde{u}) \\
= &\underline{U}(b).
\end{aligned}$$

We thus obtained that $E = [\underline{U}(b), \overline{U}(b)]$ for $b \geq 0$ such that $(1-\delta)b = \delta(\overline{U}(b) - \underline{U}(b))$. To finish, suppose that there is $b' > b$ such that $(1-\delta)b' = \delta(\overline{U}(b') - \underline{U}(b')$. Following the arguments in the proof of the previous lemma, we obtain that $[\underline{U}(b'), \overline{U}(b')]$ is self-generating, but this set strictly contains E, this provides a contradiction with the fact that E is the largest self-generating set.

□

□

A.2 Proof for the Optimal Voting Rule (to be updated)

Proof. First, note that since $e_{[1/2]} < 0$, we have that $\Delta(m) = |e_{[1-m]}|$ is increasing in m. The argument relies on a few preliminary observations:

- the partial and full cooperation thresholds, i.e. $\delta^P(m)$ and $\delta^F(m)$, are both decreasing in $\Delta(m)$ and thus in m. Indeed, we have $\delta^P(m) = \dfrac{1}{1+\Delta(m)}$ and $\delta^F(m) = \dfrac{1}{1+\frac{3\Delta(m)}{4}}$.
- the optimal utility under partial cooperation is increasing in δ and in $\Delta(m)$, and thus in m. Indeed, for $\delta^P(m) < \delta < \delta^F(m)$, we have

$$U^*(\delta, m) = U^e - 4\left(\frac{1}{\delta} - \frac{1}{\delta^F(m)}\right)^2.$$

  As $\delta < \delta^F(m)$, this function is clearly increasing in δ, decreasing in $\delta^F(m)$ and thus increasing in m.

- the utility of the stage-game equilibrium is decreasing in $\Delta(m)$ and thus in m. Indeed,

$$\begin{aligned} U^0(m) &= U^e - 4\left(\frac{1}{\delta^P(m)} - \frac{1}{\delta^F(m)}\right)^2 \\ &= U^e - \frac{\Delta(m)^2}{4}. \end{aligned}$$

- the threshold δ^O is such that $\delta^P(1) < \delta^O < \delta^F(1)$ and $\delta^O < \delta^P(1/2)$. The former inequalities are true by definition, and the latter one is obtained by observing that $U(\delta^0, 1) = U^0(1/2)$ implies $\frac{1-\delta^0}{\delta^0} - \frac{1-\delta^F(1)}{\delta^F(1)} = \dfrac{\Delta(1/2)^2}{4}$, which gives:

$$\frac{1-\delta^O}{\delta^O} = \frac{3\Delta(1) + \Delta(1/2)}{4} > \Delta(1/2) = \frac{1-\delta^P(1/2)}{\delta^P(1/2)}.$$

To conclude, we distinguish three cases:

- If $\delta \geq \delta^F(1)$, $m = 1$ achieves full cooperation. As a result, a majority rule m is optimal if and only if it achieves full cooperation, i.e. $\delta \geq \delta^F(m)$.
- If $\delta \in (\delta^O, \delta^F(1))$, unanimity achieves partial cooperation, and reaches utility $U^*(\delta, 1)$. Consider any other rule $m < 1$. If it achieves no cooperation, it cannot be better than unanimity since $U^0(m) \leq U^0(1/2) = U^*(\delta^O, 1) < U^*(\delta, 1)$. If it achieves partial cooperation, it cannot be better than unanimity since $U^*(\delta, m) < U^*(\delta, 1)$.
- If $\delta < \delta^O$, there is no cooperation under simple majority, and the utility is $U^0(1/2)$. Consider any other rule $m > 1/2$. If it achieves no cooperation, then we have $U^0(m) < U^0(1/2)$, so that m cannot be optimal. If it achieves partial cooperation, then we have $U^*(\delta, m) \leq U^*(\delta, 1) < U^*(\delta^O, 1) = U^0(1/2)$, so that m cannot be optimal.

□