



HAL
open science

Repeated Majority Voting

Antonin Macé, Rafael Treibich

► **To cite this version:**

| Antonin Macé, Rafael Treibich. Repeated Majority Voting. 2024. halshs-04610689

HAL Id: halshs-04610689

<https://shs.hal.science/halshs-04610689>

Preprint submitted on 13 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



PARIS SCHOOL OF ECONOMICS
ECOLE D'ECONOMIE DE PARIS

WORKING PAPER N° 2024 – 16

Repeated Majority Voting

Antonin Macé
Rafael Treibich

JEL Codes: D71; C73; D72

Keywords: Logrolling, Repeated games, Majority voting, Preference intensities, Consensus

anr®
agence nationale
de la recherche
AU SERVICE DE LA SCIENCE



Repeated Majority Voting*

Antonin Macé¹ and Rafael Treibich²

¹CNRS, Paris School of Economics and École Normale Supérieure-PSL

²University of Southern Denmark

June 2024

Abstract

We propose a general model of repeated voting in committees and study equilibrium behavior under alternative majority rules. We find that repetition may significantly increase the efficiency of majority voting through a mechanism of intertemporal logrolling, agents sometimes voting against their immediate preference to benefit the group's long-term interest. In turn, this affects the comparison of majority rules, which may differ significantly relative to the static setting. The model provides a rationale for the use of super-majority rules, while accounting for the prevalence of consensus in committee voting.

Keywords: logrolling, repeated games, majority voting, preference intensities, consensus.

JEL fields: D71; C73; D72.

1 Introduction

Many committees rely on majority voting to make collective decisions. Yet, majority rules are often criticized for being inefficient, as they prevent voters from expressing

*We are very grateful to Vincent Anesi, Geir Asheim, Benjamin Blumenthal, Alessandra Casella, Niall Hugues, Nicolas Klein, Barton E. Lee and Ariel Rubinstein for useful comments, and to participants of the online Social Choice and Welfare seminar, the COMSOC seminar, the workshop Aggregation Across Disciplines in Paris, the Social Choice and Welfare conference, the ACM Conference on Economics and Computation (EC'22), the OLIGO workshop in Padova, the SIOE conference, the Game Theory workshop at University of Barcelona, the European Meetings of the Econometric Society, the 2023 UniCatt Political Economy workshop, the 2024 Mediterranean Game Theory Symposium and seminars at University of Strasbourg, Paris School of Economics, Paris II Panthéon-Assas, University of Cergy, Budapest University of Technology and Economics, and Copenhagen Business School. Financial support from the ANR-17-EURE-001 is gratefully acknowledged.

the intensity of their preferences. To overcome this inefficiency, several creative mechanisms have been proposed, such as storable votes (Casella, 2012), quadratic voting (Eguia et al., 2023), or vote trading (Lazrak and Zhang, 2023). While these mechanisms are promising, their implementation in the field can be complex, and they need not always improve upon majority voting (Casella and Macé, 2021). In this paper, we argue that the inefficiency of majority voting can in fact reduce significantly when collective decisions are *repeated*. Repetition is indeed a defining feature of collective decisions in many committees; from international organizations and city councils to hiring committees and corporate boards. We show that accounting for such repetition has a substantial impact on the assessment of majority voting and, in turn, on the comparison of alternative majority rules.

In the classic setting where a committee takes a single binary decision, sincere voting is a dominant strategy and the collective decision coincides with the preference of the majority. However, when decisions are repeated, voting behavior may be conditioned both on current preferences and past decisions. As a result, agents may be incentivized to vote against their preference when it is in the group’s interest, thus increasing welfare over the long run. Our main purpose is to show how such a mechanism of *inter-temporal logrolling*, a cooperation norm whereby voters arbitrate between their own preferences and the group’s interest, emerges at equilibrium. In turn, this mechanism has implications for the design of voting rules and for the inferences that can be drawn from voting records.

In our model, an ex-ante symmetric committee makes infinitely repeated binary decisions about whether to accept or reject proposed reforms. At each stage, utilities are drawn independently from previous periods and observed publicly. Utilities reflect agents’ cardinal preferences for the reform relative to the status quo. Agents 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 under mild restrictions.

Our first result characterizes the set of equilibrium payoffs. In the repeated game, the pivotal agent may be incentivized to vote against her preference today if the net reward from complying with prescribed behavior exceeds the loss from implementing a decision she dislikes. The extent to which such non-sincere behavior can be sustained at equilibrium, i.e. the *power of intertemporal incentives*, depends on the difference between the highest and smallest equilibrium payoffs, which determines the highest admissible reward from complying with prescribed behavior. In turn, the power of

intertemporal incentives affects which decision rule can be implemented today and at which (long-term) cost, driving both the highest and smallest equilibrium payoffs. The power of intertemporal incentives thus obtains as the solution to a fixed point equation. At the optimal equilibrium, the committee follows a stationary cooperation norm, adopting the efficient decision unless the pivotal voter disagrees with a stake larger than the power of intertemporal incentives. As agents become more patient, the optimal equilibrium payoff increases, eventually reaching the first-best. We further show that the cooperation norm of the optimal equilibrium allows for a substantial level of consensual decisions, even if preferences are never unanimous. When the discount factor is large enough, any accepted reform can be approved by consensus.

To illustrate the implications of our model, we then focus on a more stylized setting where the discrepancy between average and pivotal utilities remains constant. We derive explicit formulas for the set of equilibrium payoffs under any possible majority rule and discount factor. This characterization allows us to perform two comparative statics exercises. First, we characterize the optimal voting rule, i.e. the rule maximizing the highest equilibrium payoff. While simple majority is optimal in the static benchmark, a super-majority can be optimal in the repeated game. In fact, even the unanimity rule, which performs worst in the static game, may turn out to be optimal. Second, we show that, contrary to what we would obtain in a static model, the level of consensus in the votes may be higher when committee members have more diverse preferences.

Our model assumes that agents in the committee are ex-ante symmetric. This feature is essential to ensure tractability and allows us to fully characterize the set of equilibrium payoffs in the repeated game. Yet, one may wonder if the logrolling behavior we identify at equilibrium hinges on the symmetric environment. We thus generalize the model by considering asymmetric committees and provide a necessary and sufficient condition for the existence of collectively profitable logrolling at an equilibrium of the repeated game. We then use this condition to show that logrolling is feasible in a simple model of asymmetric preferences.

Our results shed light on several aspects of majority voting. First, we argue that the inefficiency of (super-) majority rules may be exaggerated by standard voting models which do not account for repetition. Second, we show that the comparison of different majority rules in terms of welfare (such as the higher efficiency of simple majority over unanimity) can be altered, and even reversed, by accounting for repetition. Together, these arguments may explain why different majority rules are

observed in practice, even the unanimity rule which is often considered to be prone to gridlock.¹ Third, we observe that committee decisions may exhibit a higher degree of consensus under repeated voting than in the static benchmark. This finding provides a possible explanation for the widespread observation that committee decisions are often made by consensus, with or without a formal vote, even when the formal voting rule is not the unanimity rule (see [Urfalino, 2014, 2021](#) and references therein). Finally, we show that the degree of consensus may in some cases increase with the preference heterogeneity of the committee, thus drawing attention to the difficulty of inferring preferences from roll call data. We illustrate in more details our assumptions and results with the example of the Council of the European Union in [Section 8](#).

The rest of the paper is organized as follows. [Section 2](#) presents a simple example to illustrate the main intuitions behind our model. [Section 3](#) reviews the related literature. [Section 4](#) introduces the model and lays out the main assumptions. [Section 5](#) characterizes the set of equilibrium payoffs and the optimal consensus probability. [Section 6](#) revisits our general results on a more stylized model to perform comparative statics. [Section 7](#) extends the analysis to asymmetric committees by characterizing exactly when (collectively profitable) intertemporal logrolling is feasible. [Section 8](#) discusses our main assumptions and illustrates our model with the example of the Council of the European Union. [Section 9](#) concludes. All proofs are relegated to the appendices.

2 An Example

A group of three agents decides whether to accept or reject *repeated* reforms at either *unanimity* or *simple majority*. Reforms are of two possible types, A or B, each type generating a fixed distribution of utilities as described in [Table 1](#).

		Voters		
		$u_{[1]}$	$u_{[2]}$	$u_{[3]}$
Reforms	A	-3	1	5
	B	-2	-1	6

Table 1: Distribution of utilities under reforms A and B

¹See for instance [Posner and Sykes \(2014\)](#) on the variety of voting rules used in international organizations.

Agents are equally likely to occupy any rank in the associated utility distribution.² Keeping the status quo yields a utility of zero to each agent. We denote by $p \in (0, 1)$ the probability of occurrence of reforms of type A. Since both reforms A and B generate an (average) utility of 1, as opposed to 0 for the status quo, the first best consists in accepting any reform and yields an expected utility of 1.

If the committee only takes one decision, voting sincerely is a weakly dominant strategy. Under unanimity, no reform is ever accepted, yielding ex-ante utility 0. Under simple majority, only reforms A are accepted, yielding ex-ante utility $p > 0$. Simple majority thus dominates unanimity but does not achieve the first best.

If the committee takes repeated decisions, it is possible to incentivize agents to sometimes vote against their preference to achieve a higher level of utility overall. Under unanimity, consider the following strategy profile: on path, all three agents always vote in favor of reforms 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 the worst-off agent in reforms A has an incentive to vote in favor of the reform. This is the case if her inter-temporal utility from the reform being adopted, $-3(1 - \delta) + \delta$, exceeds her inter-temporal utility from the reform being rejected, 0. As a result, the first-best can be sustained under this profile if $\delta \geq 3/4$. In fact, this condition is necessary for the first-best outcome to be sustained at equilibrium.

Under simple majority, sustaining the first-best no longer requires incentivizing reform A's worst-off agent to vote in favor since reforms A are accepted under sincere voting. Consider the following strategy profile: on path, all agents but the worst-off vote in favor of reforms B, and they all vote sincerely on reforms A; any deviation is punished by a permanent reversal to the stage-game equilibrium. This profile is an equilibrium if and only if the agent with a utility of -1 for a reform B has an incentive to vote in favor of such a reform, that is $-(1 - \delta) + \delta \geq \delta p$, or equivalently $\delta \geq 1/(2 - p)$. This condition is also necessary for the first-best outcome to be sustained at equilibrium. It thus follows that when $p > 2/3$, there exists a range of discount factors, $3/4 < \delta < 1/(2 - p)$, where only unanimity can sustain the first-best.

What lessons can we draw from this simple example? First, equilibrium votes in the repeated game may be non-sincere, as we observe on reforms A and B under unanimity, and on reforms B under simple majority. Second, ex-ante utility can be higher in the repeated game, possibly reaching the first-best utility of 1, as opposed

² $u_{[1]}$, $u_{[2]}$, and $u_{[3]}$, denote the order statistics of the utility distribution associated to each type of reform.

to only 0 (under unanimity) and p (under simple majority) in the stage game. Third, some decisions can be reached by consensus in the repeated game (with probability 1 under unanimity), while they are never consensual in the stage game. Finally, the comparison between simple majority and unanimity may be reversed, the latter sometimes dominating the former in the repeated game. Why is that? In contrast to unanimity, achieving efficiency under simple majority only requires incentivizing agents with a stake $|u_i|$ 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 make efficiency harder to sustain.

3 Literature Review

Why majority voting? The literature offers various rationales for the use of majority rules in collective decisions, relying both on procedural (May, 1952; Dasgupta and Maskin, 2008) and epistemic (Condorcet, 1785) arguments. Yet, from a welfare perspective (Rae, 1969), majority voting is typically considered inefficient, as incentives constraints prevent the revelation of preference intensities (Schmitz and Tröger, 2012). To overcome these inefficiencies, several creative proposals have emerged in the literature. Casella (2012) proposes a *storable votes* mechanism whereby voters allocate a fixed budget of votes across multiple decisions, which improves over majority voting in several cases. Jackson and Sonnenschein (2007) generalize the intuition behind these results by showing that any ex-ante Pareto efficient social objective (including the utilitarian one) can be implemented if a large enough number of identical problems are linked together. Other proposals, such as quadratic voting (Eguia et al., 2023), or vote trading (Lazrak and Zhang, 2023) rely on the possibility of transfers. We contribute to this normative literature by showing how the repetition of collective decisions can significantly reduce the inefficiency of majority voting, without having to consider sophisticated mechanisms. We thus provide a new welfare rationale for majority voting.

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 exchanges between committee members, also known as *logrolling*. In the traditional logrolling model, voters have complete information on each other’s preferences over a series of binary decisions (see Casella and Macé, 2021 for a review). They exchange

votes by pairs (Riker and Brams, 1973) or within larger groups (Casella and Palfrey, 2019) to take advantage of discrepancies in their preference intensities. Yet, these exchanges are generally not welfare-improving for the group as a whole. By contrast, voters in our model only have complete information in the current stage (they are uncertain about future decisions) and they do not enter into explicit trade agreements over votes. Instead, logrolling emerges as an implicit inter-temporal agreement, which benefits all committee members. Moreover, logrolling can be decentralized as a simple cooperation norm whereby voters adopt a low-profile strategy, only voicing their disagreement when it is strong enough. This pattern of behavior is consistent with the notion of *diffuse reciprocity* in international relations (Keohane, 1986). By showing how such norm emerges from rational behavior under repeated collective decisions, we thus provide a theoretical foundation for this concept.³

A few papers from the voting literature are particularly connected to our work. Maggi and Morelli (2006) consider a setting of repeated decisions in an international organization that cannot enforce its decisions. At each step, countries vote and then decide whether to take an action, which is effective if implemented by all. The analysis focuses on voting rules that are self-enforcing (through inter-temporal incentives), and it turns out that unanimity is often optimal within that class. Our model differs as we consider enforceable decisions and we account for the heterogeneity of preference intensities, which is crucial in many environments and exacerbates the benefits of logrolling.⁴ Charroin and Vanberg (2021) compare unanimity and simple majority in a 3-voter traditional logrolling model. Relying on both simulations and laboratory experiments, they also find that unanimity can outperform simple majority. In contrast to that paper, our comparative statics results are obtained analytically for any committee size and for a larger class of voting rules, under weaker informational requirements. Finally, Fox (2006) studies a model of repeated decisions similar to ours, where one voter plays the additional role of an agenda-setter. His main objective is to elucidate situations where a coalition of voters may benefit from repetition at some equilibrium. By contrast, we characterize the complete set of equilibrium payoffs and derive comparative statics by focusing on the optimal equilibrium.

³The cooperation norm also bears some resemblance with Kocherlakota (1996)'s classic work on risk sharing, although it takes place with more than two players and without transfers. More generally, while the benefit of repetition on cooperation has long been acknowledged in the literature on repeated games (Mailath and Samuelson, 2006) and relational contracts (Levin, 2003), very few papers have explored its implications for collective decision-making.

⁴Besides, Maggi and Morelli (2006)'s argument for unanimity rule is independent from ours, and thus complementary.

Our paper contributes more broadly to the topic of dynamic collective decision-making. Most papers in this literature consider repeated decisions over a single, persistent, issue (e.g. a central bank setting the interest rate). One 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 (Baron, 1996; Kalandrakis, 2004). In a recent contribution, 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. We depart from this literature by considering repeated decisions over exogenous and independent reforms, as is the case in a committee with multiple functions and responsibilities. The repeated, as opposed to dynamic, nature of our model allows us to consider history-dependent strategies, which are essential to generate the logrolling behavior characteristic of optimal equilibria.

4 Setup

4.1 Stage game

Model - A group of agents $N = \{1, \dots, n\}$ 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 agent $i \in N$, while the status quo’s utility is normalized to 0. An agent $i \in N$ 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 $\bar{u} = (\sum_{i \in N} u_i)/n$. We say that a reform is *good* when $\bar{u} \geq 0$ and *bad* when $\bar{u} < 0$. For any rank k with $1 \leq k \leq n$, we denote by $u_{[k]}$ the k -th lowest utility, so that $u_{[1]} \leq \dots \leq u_{[n]}$.

The reform \mathbf{u} is drawn from a cumulative distribution function F . We assume that the function F is symmetric,⁵ continuous and that its support $S \subseteq \mathbb{R}^N$ is bounded and convex. The symmetry assumption reflects the fact that agents are identical ex-ante, although they may be heterogeneous ex-post. The remaining assumptions are made for ease of exposition, they ensure in particular that the distribution F has no atom.

After observing reform \mathbf{u} , each agent i votes either in favor or against the reform. A *majority rule* with threshold $m \in [1/2, 1]$ then decides the reform’s fate $d \in \{0, 1\}$.

⁵For any $\mathbf{u} \in S$ and any permutation π of N , $F(\mathbf{u}) = F(\mathbf{u}_\pi)$, where $\mathbf{u}_\pi = (u_{\pi(i)})_{i \in N}$.

If at least mn agents vote in favor, the reform is adopted, $d = 1$, and each agent i gets utility u_i . If not, the status quo remains, $d = 0$, and each agent gets utility 0. We refer to $m = 1/2$ as *simple majority* and $m = 1$ as *unanimity*.

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,⁶ every agent 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 agent with the $\lceil mn \rceil$ -th highest utility is favorable, i.e. $u_{\lceil(1-m)n\rceil} \geq 0$. This agent plays a central role in our analysis as she is a *pivot*, both in the static and in the repeated game, and she is thus denoted by p .⁷ Henceforth, we refer to her utility $u_p := u_{\lceil(1-m)n\rceil}$ as the *pivot utility*, and to her stake $|u_p|$ as the *pivot stake*. We note that the stage-game equilibrium is usually inefficient, as bad reforms may be accepted (when $u_p \geq 0$ but $\bar{u} < 0$) and good reforms may be rejected (when $u_p \leq 0$ but $\bar{u} > 0$) under sincere voting. Our main goal is to understand how such inefficiencies can be overcome when collective decisions are repeated.

Decision rules - Given majority rule m , a strategy profile $(v_i)_{i \in N}$ induces a (group) decision rule. Formally, a *decision rule* $d(\cdot)$ associates to any reform \mathbf{u} a collective decision $d(\mathbf{u}) \in \{0, 1\}$. Throughout, we focus on symmetric decision rules, and we denote the (common) expected utility attached to such a rule $d(\cdot)$ by⁸

$$U(d) = \int u_i d(\mathbf{u}) dF(\mathbf{u}).$$

We denote by $d^0(\cdot)$ the *sincere* decision rule induced by the stage-game equilibrium, $d^0(\mathbf{u}) = \mathbb{1}\{u_p \geq 0\}$. By contrast, the *efficient* rule $d^e(\cdot)$ accepts every good reform and rejects every bad reform, i.e. $d^e(\mathbf{u}) = \mathbb{1}\{\bar{u} \geq 0\}$. We denote by U^0 (resp. U^e) the expected utility under the sincere (resp. efficient) decision rule. In the sequel, we focus on the interesting case where the sincere decision rule is inefficient: $U^e := U(d^e) > U^0 := U(d^0)$.

We illustrate decision rules $d^0(\cdot)$ and $d^e(\cdot)$ on [Figure 1](#). A reform \mathbf{u} is identified graphically by its average utility \bar{u} and pivot utility u_p . The colored area represents the (projected) support S of F . Good reforms are represented in green, while bad reforms are represented in red. Collectively accepted proposals are located inside the


⁶Agents with utility $u_i = 0$ are indifferent between voting in favor or against, but this instance almost never arises because of the continuity of F . The equilibrium in undominated strategies is thus *essentially* unique.

⁷Note that the identity of the pivot is only determined ex-post, after utilities are realized. By symmetry of F , every agent is equally likely to be the pivot.

⁸The formula does not depend on the identity of $i \in N$ since both d and F are assumed to be symmetric.

black polygon. Hatched areas correspond to inefficient decisions, i.e. rejection of a good reform or acceptance of a bad reform.

Figure 1: Two decision rules.



4.2 Repeated Game

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

Timing - At each stage $t \geq 1$, a reform \mathbf{u} is drawn from probability distribution F independently of previous stages. The reform \mathbf{u} is publicly observed, then agents simultaneously vote under majority rule m 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 agent 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

We focus on subgame perfect equilibria that satisfy the following three assumptions.

Assumption 1 (Symmetry). *All agents use the same strategy.*

Assumption 2 (Anonymous Histories). *Voting behavior only depends on past anonymized realizations of utilities and decisions.*

Assumption 3 (As-if-pivotal Voting). *Taking as given future play prescribed by the strategy profile, players play as if they were pivotal.*

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

We refer to such subgame perfect equilibria as *equilibria*. *Symmetry* is a natural assumption in our context since the model is fully symmetric ex-ante. *Anonymous Histories* further constrains the set of equilibria by restricting voting behavior to depend only on anonymous information about past realizations of utilities and collective decisions. As a result, deviations from the equilibrium path may only be punished collectively. This assumption can be justified by a reluctance to single out or antagonize specific agents for their votes.¹⁰ The restriction to strategies that only condition on collective decisions, as opposed to individual votes, has been used previously in the literature to account for history-dependent strategies without losing too much tractability (Bernheim and Slavov, 2009; Anesi and Seidmann, 2015). Finally, *As-if-pivotal Voting* allows us to discard unreasonable equilibria and to pin down the voting behavior of agents who are not pivotal, which is left unconstrained by subgame perfection. Agents 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., 2023), and it derives from the game-theoretic notion of *stage-undomination* (Baron and Kalai, 1993) under the assumption of *Anonymous Histories*.

For any $\delta \in (0, 1)$, we denote by E_δ the set of equilibrium payoffs. Since E_δ is a compact interval,¹¹ we define \underline{w}_δ and \bar{w}_δ as the lowest and highest equilibrium payoffs. Note that $\underline{w}_\delta \leq U^0 \leq \bar{w}_\delta$, as playing the stage-game equilibrium after any history is an equilibrium of the repeated game. We say that a decision $d \in \{0, 1\}$ can be *implemented at equilibrium* on reform \mathbf{u} if there exists an equilibrium σ such that decision d is taken at the first stage if the first reform drawn is \mathbf{u} . We say that a decision rule $d(\cdot)$ can be *implemented at equilibrium* if there exists an equilibrium σ such that, for any reform $\mathbf{u} \in S$ drawn at the first stage, the decision $d(\mathbf{u})$ is taken.

5 Equilibrium

5.1 Intertemporal incentives

Consider a good reform \mathbf{u} that is disliked by the pivot, $u_p < 0 < \bar{u}$. In the stage game, this reform is inefficiently rejected because a blocking coalition of agents vote (sincerely) against it. By contrast, if decisions are repeated, some unfavorable voters

¹⁰Letting strategies depend on complete histories of play would imply that both rewards and punishments can be tailored individually. This would increase the incentive of the pivotal voter to abide by the prescribed decisions, and eventually allow to achieve higher levels of efficiency. From this point of view, our analysis can thus be understood as providing a lower bound on the highest equilibrium payoff.

¹¹See Lemma 3 in Appendix A.

may be convinced to vote in favor of the reform through intertemporal incentives. As a result, the reform may instead be accepted at equilibrium. To be more precise, assume that accepting the reform today triggers the best equilibrium as a reward, while rejecting the reform triggers the worst equilibrium as a punishment. Then, voting in favor of the reform yields an average intertemporal payoff of $\delta\bar{w}_\delta$, while rejecting the reform yields an average intertemporal payoff $(1 - \delta)u_i + \delta\bar{w}_\delta$. If the pivot utility is not too negative, i.e.

$$u_p \geq -\frac{\delta}{1 - \delta}(\bar{w}_\delta - \underline{w}_\delta), \quad (1)$$

then all agents with a utility greater than or equal to the pivot utility have an incentive to vote in favor of the reform. Since these agents form a winning coalition, the reform is accepted. In fact, condition (1) is necessary and sufficient for the reform to be implemented at equilibrium because (i) any other winning coalition would be harder to incentivize and (ii) the best and worst equilibrium payoffs provide the optimal intertemporal incentives for voters to comply with non-sincere behavior.

More generally, we define the cost of implementing a collective decision $d \in \{0, 1\}$ on a reform $\mathbf{u} \in S$ as,

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

It is equal to 0 if the pivot agrees with decision d and to the pivot stake $|u_p|$ if she does not. Following the same reasoning as before, we get that a decision d can be implemented on reform \mathbf{u} if and only if its cost of implementation $c(\mathbf{u}, d)$ does not exceed the *power of intertemporal incentives* κ_δ defined by,

$$\kappa_\delta := \frac{\delta}{1 - \delta}(\bar{w}_\delta - \underline{w}_\delta). \quad (2)$$

The power of intertemporal incentives κ_δ thus reflects how costly a decision can be implemented at equilibrium, given that continuation promises must themselves be equilibrium payoffs and future is discounted by δ . For any decision rule $d(\cdot)$, we define $C(d)$ and $\Delta(d)$ as, respectively, the expected and largest cost of implementing decision $d(\mathbf{u})$,

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

Note that $C(d^0) = \Delta(d^0) = 0$ since the pivot always agrees with d^0 . We denote by $c^e(\mathbf{u}) := c(\mathbf{u}, d^e(\mathbf{u}))$ the cost of implementing the efficient decision on reform \mathbf{u} . We denote by $\Delta^e := \Delta(d^e) = \sup_{\mathbf{u} \in S} c^e(\mathbf{u})$ the largest cost of implementing an efficient

decision.

For any parameter $\kappa \geq 0$, we say that a rule $d(\cdot)$ is capped by κ if the cost of implementing decision $d(\mathbf{u})$ on any reform $\mathbf{u} \in S$ does not exceed κ , i.e. $\Delta(d) \leq \kappa$. We denote by \mathcal{D}_κ the set of all decision rules capped by κ . The main observation of the above discussion is that a decision rule can be implemented at equilibrium if and only if it belongs to $\mathcal{D}_{\kappa_\delta}$.

5.2 Characterization of equilibrium payoffs

Our characterization of equilibrium payoffs relies on identifying the highest and lowest equilibrium payoffs as functions of the (endogenous) power of intertemporal incentives. For any cutoff $\kappa \in \mathbb{R}_+$, we define

$$\bar{V}(\kappa) := \sup_{d \in \mathcal{D}_\kappa} U(d) \quad (3)$$

as the highest utility that can be achieved by a decision rule capped by κ , and

$$\underline{V}(\kappa) := \inf_{d \in \mathcal{D}_\kappa} U(d) + C(d) \quad (4)$$

as the lowest sum of utility and expected cost of implementation that can be achieved by a decision rule capped by κ . As we will explain below, $\bar{V}(\kappa)$ and $\underline{V}(\kappa)$ coincide respectively with the highest and lowest equilibrium payoff when the power of intertemporal incentives is equal to κ . Intuitively, accounting for the cost of implementation is needed to construct the lowest equilibrium payoff because some inefficient decisions can only be implemented by promising rewards that have to be incurred on the equilibrium path.

Both programs (3) and (4) have straightforward solutions (to be described below), which means functions \bar{V} and \underline{V} can be computed explicitly. Furthermore, since \mathcal{D}_κ expands when κ increases, \bar{V} and \underline{V} are respectively increasing and decreasing with κ . We now characterize the set of equilibrium payoffs in the repeated game.

Theorem 1. *For any $\delta \in [0, 1]$, the set of equilibrium payoff E_δ is equal to $[\underline{V}(\kappa_\delta), \bar{V}(\kappa_\delta)]$, where the power of intertemporal incentives κ_δ is given by:*

$$\kappa_\delta = \max \left\{ \kappa \geq 0 \mid (1 - \delta)\kappa = \delta [\bar{V}(\kappa) - \underline{V}(\kappa)] \right\}. \quad (5)$$

The set of equilibrium payoffs is characterized through the power of intertemporal

incentives κ_δ , which is defined as the largest solution to fixed-point equation:

$$(1 - \delta)\kappa = \delta [\bar{V}(\kappa) - \underline{V}(\kappa)]. \quad (6)$$

Before describing the implications of [Theorem 1](#), we start by explaining the main intuition for its proof. Suppose that power of intertemporal incentives κ_δ is given. What decision rule should be implemented in the first period to achieve the *highest* equilibrium payoff? Without loss of generality, we can assume that the highest equilibrium payoff \bar{w}_δ is always used as a reward, since it maximizes both the incentives to the pivotal voter and the continuation utility that will later realize on the equilibrium path. Conversely, we can assume that the lowest equilibrium payoff \underline{w}_δ is always used as a punishment, since it provides maximal incentives without affecting the payoffs on the equilibrium path.

Under such optimal incentives, the optimal decision rule simply consists in taking the efficient decision unless its cost of implementation exceeds power κ_δ . This *capped efficient rule* $d_{\kappa_\delta}^e$ is formally defined by (see also [Figure 2](#), left panel, below),

$$\forall \mathbf{u} \in S, \quad d_{\kappa_\delta}^e(\mathbf{u}) = \begin{cases} d^e(\mathbf{u}) & \text{if } c^e(\mathbf{u}) \leq \kappa_\delta \\ d^0(\mathbf{u}) & \text{if } c^e(\mathbf{u}) > \kappa_\delta, \end{cases} \quad (7)$$

and it is such that $U(d_{\kappa_\delta}^e) = \max_{d \in \mathcal{D}_{\kappa_\delta}} U(d) = \bar{V}(\kappa_\delta)$. Since the associated path must generate the highest equilibrium payoff \bar{w}_δ , we obtain,

$$(1 - \delta)\bar{V}(\kappa_\delta) + \delta\bar{w}_\delta = \bar{w}_\delta,$$

which yields $\bar{w}_\delta = \bar{V}(\kappa_\delta)$. We conclude that $\bar{V}(\kappa_\delta)$ is the highest equilibrium payoff that can be induced by power of intertemporal incentives κ_δ .

Conversely, we ask what decision rule should be implemented in the first period to achieve the *lowest* equilibrium payoff? We first note that no matter what decision rule is played in the first period, minimizing the intertemporal equilibrium payoff requires using as small a reward as necessary, since that reward will later be realized on the equilibrium path. Given that it is still optimal to punish deviations with the lowest equilibrium payoff \underline{w}_δ , complying with decision d on reform \mathbf{u} must be rewarded with continuation utility,

$$\underline{w}_\delta + \frac{1 - \delta}{\delta} c(\mathbf{u}, d). \quad (8)$$

As a result, to achieve as low an equilibrium payoff as possible, one should implement

an inefficient decision only if its implementation cost is not too high, as this would require too large a reward, to be later realized on the equilibrium path. In fact, implementing an inefficient decision ends up increasing the average intertemporal payoff if and only if its implementation cost exceeds $|\bar{u}|$.

The lowest equilibrium payoff is thus achieved by implementing the *capped semi-inefficient rule* $d_{\kappa_\delta}^i$, formally defined by (see also [Figure 2](#), right panel, below),

$$\forall \mathbf{u} \in S, \quad d_{\kappa_\delta}^i(\mathbf{u}) = \begin{cases} 1 - d^e(\mathbf{u}) & \text{if } c^e(\mathbf{u}) \leq \min(\kappa_\delta, |\mathbf{u}|) \\ d^0(\mathbf{u}) & \text{if } c^e(\mathbf{u}) > \min(\kappa_\delta, |\mathbf{u}|). \end{cases} \quad (9)$$

This rule is such that $U(d_{\kappa_\delta}^i) + C(d_{\kappa_\delta}^i) = \min_{d \in \mathcal{D}_{\kappa_\delta}} U(d) + C(d) = \underline{V}(\kappa_\delta)$. To conclude, the lowest equilibrium payoff is achieved by implementing $d_{\kappa_\delta}^i$, using payoff (8) as a reward and \underline{w}_δ as a punishment.¹² Since the associated path must generate the lowest equilibrium payoff \underline{w}_δ , we obtain,

$$(1 - \delta)U(d_{\kappa_\delta}^i) + \delta \int \left[\underline{w}_\delta + \frac{(1 - \delta)}{\delta} c(\mathbf{u}, d_{\kappa_\delta}^i(\mathbf{u})) \right] dF(\mathbf{u}) = \underline{w}_\delta,$$

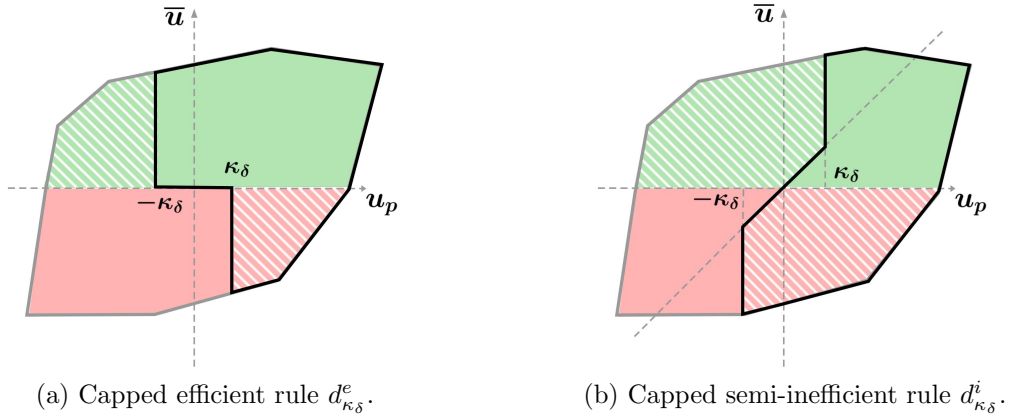
which yields $\underline{w}_\delta = U(d_{\kappa_\delta}^i) + C(d_{\kappa_\delta}^i) = \underline{V}(\kappa_\delta)$. We conclude that $\underline{V}(\kappa_\delta)$ is the lowest equilibrium payoff that can be induced by power κ_δ .

To conclude, it follows from the definition of κ_δ in (2) and from the characterization of the highest and lowest equilibrium payoffs \bar{w}_δ and \underline{w}_δ described in the above discussion that the fixed-point equation (6) must hold.

We summarize the construction of the extremal equilibrium payoffs in [Figure 2](#). The left panel illustrates the *capped efficient rule* $d_{\kappa_\delta}^e$, under which good reforms are accepted unless the pivot utility is too small ($u_p \leq -\kappa$), while bad reforms are rejected unless the pivot utility is too high ($u_p \geq \kappa$). The right panel illustrates the *capped semi-inefficient rule* $d_{\kappa_\delta}^i$, under which good reforms are rejected unless the pivot utility is too high ($u_p \geq \min(\bar{u}, \kappa)$), while bad reforms are accepted unless the pivot utility is too low ($u_p \leq \max(\bar{u}, -\kappa)$).

¹²Note that this equilibrium is non-stationary, as the reward for complying with decision rule $d_{\kappa_\delta}^i$ depends on utility realization \mathbf{u} . Another observation here is that (8) is indeed an equilibrium payoff as (i) by definition of $d_{\kappa_\delta}^i$, we always have $0 \leq c(\mathbf{u}, d_{\kappa_\delta}^i(\mathbf{u})) \leq \kappa_\delta$ and (ii) the set of equilibrium payoffs is a compact interval, as we show in the appendix (Claim 2 of Lemma 3 in the proof of [Theorem 1](#)).

Figure 2: Capped decision rules in the proof of [Theorem 1](#).



5.3 Optimal equilibrium

We now describe the main implications of [Theorem 1](#) for the outcomes of *optimal equilibria*, i.e. equilibria which generate the highest payoff \bar{w}_δ .

Corollary 1. *Any optimal equilibrium is such that the capped efficient rule $d_{\kappa_\delta}^e$ is implemented at each stage.*

[Corollary 1](#) implies that the optimal equilibrium may be understood as a (self-enforcing) *cooperation norm*, whereby voters routinely abide by collectively efficient decisions unless the pivot stake $|u_p|$ exceeds cutoff κ_δ .¹³ In turn, parameter κ_δ may be interpreted as the optimal *degree of cooperation* that is achievable in the repeated game.

Corollary 2. *The degree of cooperation κ_δ is weakly increasing with δ .*¹⁴

A consequence of [Corollary 2](#) is that the set of equilibrium payoffs expands with δ . As voters becomes more patient, 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.

We identify three cooperation regimes.

¹³To be precise, this stationary path constitutes the unique optimal equilibrium path. Yet, several punishments schemes are possible at an optimal equilibrium (in fact, all those that provide sufficient incentives for agents to abide by prescribed behavior).

¹⁴To see how [Corollary 2](#) derives from [Theorem 1](#), note that the function $G(\delta, \kappa) = \frac{1-\delta}{\delta}\kappa - (\bar{V}(\kappa) - \underline{V}(\kappa))$ is continuous, weakly decreasing in δ and such that $G(\delta, \kappa_\delta) = 0$. For any $\delta' > \delta$, we thus have $G(\delta', \kappa_\delta) \leq 0$ and, since both \bar{V} and \underline{V} are bounded (since the support of F itself is bounded), $\lim_{\kappa \rightarrow \infty} G(\delta', \kappa) = +\infty$. By the intermediate value theorem, there exists $\kappa' \geq \kappa_\delta$ such that $G(\delta', \kappa') = 0$. It then follows from [Theorem 1](#) that $\kappa_{\delta'} \geq \kappa' \geq \kappa_\delta$.

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

- (i) *if $\delta < \delta^P$, then $\bar{w}_\delta = U^0$. No cooperation is possible at equilibrium.*
- (ii) *if $\delta \in [\delta^P, \delta^F)$, then $U^0 < \bar{w}_\delta < U^e$. Partial cooperation is possible at equilibrium.*
- (iii) *if $\delta \geq \delta^F$, then $\bar{w}_\delta = U^e$. Full cooperation is possible at equilibrium.*

If the discount factor is too low, the power of intertemporal incentives is null so that the pivotal voter can never be incentivized to vote against her preferences. As a result, the highest equilibrium payoff coincides with that of the stage-game equilibrium (*no cooperation*). If the discount factor is intermediate, the power of intertemporal incentives is positive but smaller than Δ^e , the highest implementation cost of an efficient decision. In that case, the highest equilibrium payoff improves over the stage-game equilibrium without achieving the first-best, as some efficient decisions cannot be implemented (*partial cooperation*). Finally, if the discount factor is large enough, the power of intertemporal incentives exceeds Δ^e . That case is reminiscent to the folk theorem, where every efficient decision can be implemented, and the highest equilibrium payoff coincides with the first-best U^e (*full cooperation*).

5.4 Individual Voting Behavior and Consensus

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

For any stationary optimal equilibrium σ ,¹⁵ we denote by v^σ the voting profile used at each stage. The *consensus probability* under σ is defined as the probability that a reform is unanimously approved, that is,

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

This probability is the largest under the following profile,

$$\forall i \in N, \forall \mathbf{u} \in S, \quad v_i(\mathbf{u}) = \begin{cases} \mathbb{1}\{u_i \geq -\kappa_\delta\} & \text{if } \bar{u} \geq 0 \\ \mathbb{1}\{u_i \geq \kappa_\delta\} & \text{if } \bar{u} < 0. \end{cases}$$

This voting behaviour may be interpreted as agents generally “keeping a low profile”:

¹⁵Restricting attention to stationary equilibria is reasonable here as any optimal equilibrium has a stationary outcome ([Corollary 1](#)).

following the group preference unless they strongly disagree with the group, i.e. with a stake higher than κ_δ . We refer to the associated consensus probability P_δ as the *optimal consensus probability*.¹⁶

Proposition 1. *The optimal consensus probability P_δ is weakly increasing with δ . There is a threshold $\delta^C \in (0, 1)$ such that the optimal consensus probability is maximal for $\delta \geq \delta^C$. All accepted good reforms are then unanimously approved.*

As the discount factor increases, the larger power of inter-temporal 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 larger and thus more good reforms are approved.¹⁷ 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 main lesson of **Proposition 1** 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 unanimously approved. 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 unanimously approved.

6 Comparative statics

In order to illustrate the implications of repetition on majority voting, we now consider a simplified model which allows us to derive explicit formulas for equilibrium payoffs and perform comparative statics.

¹⁶We note that the highest consensus probability at a stationary 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.

¹⁷This also means that less bad reforms are approved, but as we note in the proof, these reforms never receive a consensual approval.

6.1 Preference distributions

We focus on a stylized class of preference distributions where the utility of each agent $i \in N$ writes as the weighted sum of a common component θ and a private bias β_i ,

$$u_i = \theta + \alpha\beta_i,$$

where the common component θ is drawn uniformly on $[-1, 1]$, and the ordered vector of biases is fixed across reforms: there exist $b_1 \leq \dots \leq b_n$ such that the j -th lowest bias is $\beta_{[j]} = b_j$ for all $j \in \{1, \dots, n\}$. 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 bias at any given rank of the distribution. Parameter α reflects the (relative) weight attached to the bias, which we interpret as a measure of *diversity* within the committee. For ease of exposition, we assume that n is odd, $b_1 = -1$, $b_{\frac{n+1}{2}} < \bar{b} = 0$, and $\alpha \in [0, 1]$, and we note $b_p := b_{\lfloor (1-m)n \rfloor}$.¹⁸ We refer to this class of preference distributions as the *simplified model*.

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

A reform is good if and only if the common component is positive, $\theta \geq 0$, and it is collectively accepted at the stage-game equilibrium if and only if the pivot is in favor, $\theta \geq -\alpha b_p$. For any majority rule, we have $b_p \leq b_{\frac{n+1}{2}} < 0$, so that the inefficiency of the stage-game equilibrium comes from good reforms $\theta \in [0, -\alpha b_p]$ being rejected. Under any voting rule m , the expected payoff under the efficient and sincere decision rules are now equal to:

$$U^e = \frac{1}{4} \quad \text{and} \quad U_m^0 = U^e - \frac{(\Delta_m^e)^2}{4},$$

where the largest cost of implementing an efficient decision writes $\Delta_m^e = \alpha|b_p|$. Here, the parameter Δ_m^e can also be interpreted as a measure of the *inefficiency* associated with majority rule m in the stage game. In particular, we obtain that the stage-game equilibrium payoff is decreasing with m .

¹⁸The condition $\alpha \in [0, 1]$ ensures that there always are some reforms that are good (resp. bad) for all voters. This property is not essential for the results but slightly simplifies the analysis.

Proposition 2. *In the stage game of the simplified model, the best majority rule is simple majority ($m = 1/2$) and the worst majority rule is unanimity ($m = 1$).*

We note that, although simple majority is optimal, it does not achieve the first best since $\Delta_{1/2}^e = |\alpha b_{\frac{n+1}{2}}| > 0$ and thus $U_{1/2}^0 < U^e$.

6.2 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^e , we can obtain closed-form formulas for the optimal degree of cooperation κ_δ . In turn, we obtain the partial and full cooperation thresholds δ^F and δ^P , as well as the optimal utility \bar{w}_δ . This eventually allows us to characterize the optimal voting rule, i.e. the rule which maximizes the highest equilibrium payoff \bar{w}_δ .

Theorem 2. *In the simplified model, there are thresholds $\underline{\delta}, \bar{\delta}$ with $0 < \underline{\delta} < \bar{\delta} < 1$ and right-continuous functions $m^*, m^{**} : [0, 1] \rightarrow [1/2, 1]$ such that:*

- (i) *if $\delta < \underline{\delta}$, the optimal rule is simple majority ($m = 1/2$).*
- (ii) *if $\delta = \underline{\delta}$, the optimal rules are simple majority ($m = 1/2$) and the strict (super-) majority $m^*(\underline{\delta})$.*
- (iii) *if $\underline{\delta} < \delta < \bar{\delta}$, the optimal rule(s) are the majority rule(s) $m^*(\delta)$ and $\lim_{\delta' \rightarrow \delta, \delta' < \delta} m^*(\delta')$.*
- (iv) *if $\delta \geq \bar{\delta}$, any majority rule m with $m^*(\delta) \leq m \leq m^{**}(\delta)$ is optimal.*

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

The mapping from discount factor to optimal voting rules is illustrated in [Figure 3](#).¹⁹

¹⁹Empty circles correspond to optimal voting rules that are not contained in the functions m^* and m^{**} .




Figure 3: Optimal voting rule(s) as a function of the discount factor.

Theorem 2 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 optimal voting rule coincides with the optimal rule in the stage-game, i.e. simple majority ($m = 1/2$). In the opposite benchmark where the discount factor is high ($\delta \geq \bar{\delta}$), intertemporal incentives are sufficient to achieve full cooperation for some rules ($m^*(\delta) \leq m \leq m^{**}(\delta)$), and eventually for all rules when δ is large enough.

The most interesting case of **Theorem 2** arises for intermediate discount factors ($\underline{\delta} \leq \delta < \bar{\delta}$). The optimal rule is then a supermajority rule (necessarily strict when δ is just above $\underline{\delta}$) which decreases (weakly) with the discount factor. 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 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 2**.

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

How can unanimity be optimal in the repeated game when it performs worst in the stage-game? The intuition for **Corollary 4** can be best explained by focusing on the full cooperation threshold δ_m^F , which is decreasing in majority threshold m .²⁰ In

²⁰Note that this observation alone guarantees that unanimity is the unique optimal majority rule for δ close enough

this simplified model, the largest cost of implementing an efficient decision is equal to Δ_m^e , which means that achieving full cooperation requires a power of intertemporal incentives κ_δ of at least Δ_m^e . It follows from fixed-point equation (6) that the full cooperation threshold satisfies,²¹

$$\frac{1 - \delta_m^F}{\delta_m^F} = \frac{1}{\Delta_m^e} [\bar{V}(\Delta_m^e) - \underline{V}(\Delta_m^e)].$$

The majority rule m thus has two opposing effects on threshold δ_m^F . On the one hand, since Δ_m^e is increasing in m , achieving full cooperation requires a larger power of intertemporal incentives under unanimity than under any other rule. On the other hand, the gap in equilibrium payoffs induced by power of intertemporal incentives Δ_m^e ,

$$\bar{V}(\Delta_m^e) - \underline{V}(\Delta_m^e) = \frac{3(\Delta_m^e)^2}{4}$$

is increasing in Δ_m^e and thus largest under unanimity. Here, since the latter effect (quadratic) dominates the former (linear), we get that full cooperation is easiest to achieve under unanimity.²²

6.3 Consensus

Following [Section 5.4](#), 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 κ_δ) to vote in favor of decisions that must be accepted at equilibrium. The optimal consensus probability is thus given by

$$P_\delta = \mathbb{P}(u_{[1]} + \kappa_\delta \geq 0, \bar{u} \geq 0) = \mathbb{P}(\theta \geq \alpha - \kappa_\delta, \theta \geq 0) = \frac{1 - \max(\alpha - \kappa_\delta, 0)}{2}$$

whenever $\alpha - \kappa_\delta \geq -1$. This probability can be computed explicitly in the simplified model by plugging in the formulas for κ_δ .

Proposition 3. *In the simplified model, the optimal consensus probability P_δ can be increasing with diversity α .*

[Proposition 3](#) shows that voting behavior may be more often unanimous when

to (and larger than) $\delta^F(1)$.

²¹In fact, because κ_δ is both continuous and weakly increasing in δ , the smallest discount factor for which it is possible to achieve full cooperation must generate a power of exactly Δ_m^e .

²²A similar insight that a worse institution in static terms can turn out optimal with repetition appears in [Fox and Polborn \(2022\)](#), although in a different setting. A further difference is that we focus on the optimal equilibrium rather than on a particular candidate equilibrium in the repeated game.

preferences are more diverse (α is higher). Hence, the observation of vote tallies may be a poor indicator of the preferences within a committee if agents engage in inter-temporal logrolling. To get some intuition on [Proposition 3](#), note that increasing diversity has two effects. On the one hand, it becomes more difficult to convince the worst-off agent to vote in favor of the reform when diversity is larger. On the other hand, as we show in the appendix, when partial cooperation begins, a larger diversity leads to stronger inter-temporal incentives (i.e. higher κ_δ), an effect which is directly proportional to the measure $|b_p|$ of the majority rule's static inefficiency. When the majority rule is inefficient enough, the latter effect dominates, and the optimal consensus probability increases with α .

7 Asymmetric Committees

We have maintained up to now the assumption that committee members are ex-ante symmetric. This stylized hypothesis ensures the possibility of (collectively profitable) inter-temporal logrolling while providing a tractable framework for understanding its consequences. Yet, one may wonder whether this mechanism carries over to asymmetric committees. In this section, we provide a necessary and sufficient condition for the feasibility of (collectively profitable) inter-temporal logrolling, and show that it is satisfied in a simple model of asymmetric preferences.

7.1 Collectively Profitable Logrolling

We extend the model as follows: the c.d.f. F of the reform distribution is now continuous but possibly asymmetric, with a bounded and convex support $S \subseteq \mathbb{R}^N$. In the repeated game, an *equilibrium* σ is a subgame-perfect equilibrium satisfying *Anonymous Histories* and *As-If-Pivotal Voting*. We say that an equilibrium σ is *outcome-stationary* if the decision rule it induces at each stage always remains the same on the equilibrium path. For instance, as in the symmetric model, sincere voting is an outcome-stationary equilibrium. We say that a probability distribution F and a majority threshold m jointly satisfy the *Gains from Trade Condition* if there exist $K \geq 1$, a family of reforms $(\mathbf{u}^k)_{1 \leq k \leq K} \in S^K$ and a family of weights $(w^k)_{1 \leq k \leq K} \in (\mathbb{R}_+)^K$ such that:

$$\forall i \in N, \quad \sum_{k=1}^K w^k u_i^k (1 - 2d^0(\mathbf{u}^k)) > 0. \quad (10)$$

Intuitively, the condition reflects the possibility of reversing a certain number of sincere

collective decisions in a mutually beneficial way.²³ The condition precisely characterizes when (collectively profitable) intertemporal logrolling can occur.

Theorem 3. *There exists a stationary-outcome equilibrium which strictly Pareto-dominates the sincere voting equilibrium for some $\delta \in (0, 1)$ if and only if (F, m) satisfies the Gains from Trade Condition.*

This result is powerful as it provides a tractable condition which can be checked to investigate if a given preference model allows (or not) for (collectively profitable) intertemporal logrolling.²⁴ Note that when the condition is satisfied, if there exists a strictly Pareto-improving stationary-outcome equilibrium for some $\delta \in (0, 1)$, then such an equilibrium also exists for any $\delta' \geq \delta$. Moreover, the Pareto-dominating equilibrium we construct in the proof also generates more consensus than sincere voting, in line with the analysis of the previous sections.

In order to illustrate [Theorem 3](#), we now consider a simple model of asymmetric preferences and discuss when the Gains from Trade Condition is satisfied.

7.2 Example

We consider the following simple model of asymmetric preferences, adapted from [Dzuida and Loeper \(2018\)](#).²⁵ Each collective decision features a left proposal ($d = L$) and a right proposal ($d = R$). Each agent i 's utility for the right proposal writes as the sum of a common component θ , a private bias β_i and a noise term ε_i ,

$$x_i^R = \theta + \beta_i + \varepsilon_i,$$

while the utility of the left proposal x_i^L is normalized to 0. Private biases are fixed and such that $\beta_1 < \dots < \beta_n$. On each reform, common component θ is drawn from a distribution whose support includes $(-\beta_n - \eta, -\beta_1 + \eta)$ for some $\eta > 0$. Each voter-specific noise term ε_i is drawn from a distribution with support $[-\bar{\varepsilon}, \bar{\varepsilon}]$, with $\bar{\varepsilon} \geq 0$, independently from $(\varepsilon_j)_{j \neq i}$ and from θ .

On any decision, the status-quo may be either a left proposal (the reform is then a right proposal) or a right proposal (the reform is then a left proposal). Formally, each

²³While we focus here on cooperation at the level of the whole group, we acknowledge that cooperation by subgroups can be relevant in several contexts too. See for instance [Fox \(2006\)](#) and [Eguia \(2011\)](#) for an application to political parties in legislatures and [Kikuchi and Koriyama \(2023\)](#) for an application to the U.S. Electoral College.

²⁴The robustness exercise conducted in this section bears some similarity to the analysis of [Fox \(2006\)](#). Yet, by contrast to [Fox \(2006\)](#)'s Proposition 1, we provide a necessary and sufficient condition for the existence of logrolling, which, as we show, is sufficiently tractable to be directly tested in applied models.

²⁵In their model, the status quo is not stochastic but inherited from the previous collective decision.

agent i 's utility for a reform writes as $u_i = (\theta + \beta_i + \varepsilon_i)s$, where variable $s \in \{-1, 1\}$ is drawn from a non-degenerate distribution, independently from θ and $\{\varepsilon_i\}_{i \in N}$.²⁶ We refer to this class of preference distributions as the *simplified asymmetric model*.

Proposition 4. *In the simplified asymmetric model, when n is odd:*

- for any $\bar{\varepsilon} \geq 0$, the Gains from Trade Condition is satisfied for any strict super-majority rule $m > \frac{1}{2} + \frac{1}{n}$,
- there exists $\bar{\varepsilon}^* > 0$ with $\bar{\varepsilon}^* \leq \frac{1}{2}(\beta_{\frac{n+3}{2}} - \beta_{\frac{n-1}{2}})$ such that: the Gains from Trade Condition is satisfied under simple majority rule ($m = 1/2$) if and only if $\bar{\varepsilon} \geq \bar{\varepsilon}^*$.

The result shows that inter-temporal logrolling is indeed possible on a large domain of asymmetric preferences, under any strict super-majority rule. While the median voter can never gain from logrolling absent noise under simple majority, a small amount of noise is sufficient to make inter-temporal logrolling possible even under that rule. More generally, the example illustrates that [Theorem 3](#) provides a tractable test to assess the feasibility of logrolling for any preference model.

8 Discussion

8.1 Assumptions

Malevolent Behavior. Our characterization of equilibrium payoffs relies on the possibility of punishing agents for voting in favor of efficient decisions (see [Figure 2](#), right panel, dashed areas with $u_p \bar{u} > 0$). One may consider such punishment schemes to be unrealistic, or undesirable. How much cooperation could be achieved at equilibrium without having to rely on these kinds 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 lowest equilibrium payoff necessarily coincides with the stage-game equilibrium payoff U^0 . The set of equilibrium payoffs is then equal to $[U^0, \bar{V}(\kappa'_\delta)]$, where cutoff $\kappa'_\delta \leq \kappa_\delta$ is given by,

$$\kappa'_\delta = \max \left\{ \kappa \geq 0 \mid (1 - \delta)\kappa = \delta [\bar{V}(\kappa) - U^0] \right\}.$$

The restriction to non-malevolent strategies weakens the optimal punishment (i.e. the lowest equilibrium payoff), thus reducing the highest equilibrium payoff. As a

²⁶Here, $s = 1$ refers to the status quo being L and $s = -1$ refers to the status-quo being R .

result, the set of equilibrium payoffs under non-malevolence is included in the set of equilibrium payoffs characterized in [Theorem 1](#). The inclusion is strict unless there is either no or full cooperation. However, the characterization of optimal voting rules in the simplified model ([Theorem 2](#)) remains valid.

Incomplete information. Our model assumes complete information in the stage-game, which is crucial to maintain tractability. While this assumption may seem a priori restrictive, it is in fact consistent with the main application of our model. Our analysis is concerned with repeated collective decisions in a *standing committee*. In such a setting, it is natural to think that committee members learn to know each other well, which makes the complete information assumption plausible. Assuming complete information is also less demanding than it may initially appear. First, we only require complete information at the current stage, thus preserving uncertainty in the shadow of future decisions, making the model significantly less demanding than traditional models of logrolling. Second, the cooperation norm characteristic of the optimal equilibria can be implemented with rather minimal informational requirements. As discussed in [Section 5.4](#), voters only need to observe their own utility and the average utility (an aggregate statistic) to implement the optimal equilibrium on path.²⁷

One may still wonder if the norm of inter-temporal logrolling could emerge under incomplete information. In such environments, if we allow for communication before each vote, equilibria can indeed be found which combine the cooperation norm together with a linking mechanism ([Jackson and Sonnenschein, 2007](#)). This mechanism is not imposed by a designer but rather implemented on path by the voters through a threat of punishment, and it allows voters to reveal their preference intensities. We illustrate this idea in a simple example in [Appendix B](#).

8.2 Application to the Council of the E.U.

We conclude by discussing how our assumptions and results apply to the Council of the European Union, which is the topic of an extensive literature in political science. This committee, which gathers ministers from each Member State, is one of the E.U.'s main decision-making body, and takes numerous decisions on many different policy areas. While the formal voting rule varies across areas (from unanimity to qualified

²⁷Note that a single other aggregate statistic, the pivot utility, is required to implement the optimal equilibrium off path.

majority), a common observation is that most decisions are taken by consensus, and that negative votes are rarely observed when a formal vote takes place (Mattila and Lane, 2001).

The Council of the E.U. fits well our model’s main assumptions. First, committee decisions are prepared by a small group of representatives with strong ties (COREPER), who “know each other extremely well” and show a “sense of understanding of the other delegations’ interest” (Lewis, 2000), ensuring that “member states are well aware both of each other’s policy preferences and of how important particular decisions are for them” (Heisenberg, 2005). This makes the assumption of complete information about current preferences particularly relevant. Second, since decisions are made in many different policy areas, coalitions in favor of a given reform change frequently (Mattila and Lane, 2001),²⁸ which lends support to the assumption that preferences are independent across periods.

The norm of *diffuse reciprocity*, by which Member States only voice their disagreement in extreme cases, is often put forward to explain the prevalence of consensus in the Council. Warntjen (2010) notes that delegates “apply an informal rule which allows mutual concessions for mutual benefit in long-standing relationships’, while Lewis (2000) writes that “the norm of diffuse reciprocity can also instill a sense of self-restraint in the kinds of demands and outcomes a delegation pushes for”.²⁹ By formalizing this norm as a mechanism of inter-temporal logrolling and by showing how it emerges at the optimal equilibrium of a repeated voting game, our work thus provides a theoretical foundation for this important empirical observation.³⁰

9 Conclusion

Majority voting is often criticized for not reflecting the intensity of individual preferences, leading to inefficient collective decisions. In this paper, we argue that these inefficiencies can be significantly reduced when decisions are repeated. We propose a general model of repeated collective decisions and characterize equilibrium outcomes

²⁸“Such is the nature of Community Affairs, however, that coalitions change dramatically according to the item under discussion” (Mattila and Lane, 2001).

²⁹Lewis (2000) also writes that “the balancing of concessions and the expectations that all will profit in the long run is spread across a range of issue-areas and functional contexts”.

³⁰The rationales for consensus in the Council are sometimes framed as divided between rationalist (logic of consequences, goal-oriented) and constructivist (logic of appropriateness, norm-oriented) explanations, see for instance Warntjen (2010). Our work illustrates that the two approaches are not exclusive, in line with Fearon and Wendt (2002).

under various majority rules. We find that repetition may significantly increase the efficiency of majority voting through a mechanism of inter-temporal logrolling, agents sometimes voting against their immediate preferences to benefit the group’s long-term interest. In turn, this mechanism can substantially affect the comparison of alternative majority rules, and may explain the prevalence of supermajority rules. Repetition also allows for higher levels of consensus, a feature of collective-decision making that is often observed in practice. While the model applies to an ex-ante symmetric environment, we show that the main mechanism extends to asymmetric committees and identify the exact conditions under which (collectively profitable) intertemporal logrolling is feasible.

References

- Ali, S. N., Bernheim, B. D., Bloedel, A. W., and Battilana, S. C. (2023). Who controls the agenda controls the legislature. *American Economic Review*, 113(11):3090–3128.
- 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*. University of Michigan Press, Ann Arbor.
- Casella, A. (2012). *Storable votes: protecting the minority voice*. Oxford University Press.
- 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*.
- Condorcet (1785). *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Paris: Imprimerie Royale.
- Dasgupta, P. and Maskin, E. (2008). On the robustness of majority rule. *Journal of the European Economic Association*, 6(5):949–973.
- Dzuida, W. and Loeper, A. (2018). Dynamic pivotal politics. *American Political Science Review*, 112(3):580–601.
- Eguia, J. X. (2011). Endogenous parties in an assembly. *American Journal of Political Science*, 55(1):16–26.
- Eguia, J. X., Immorlica, N., Lalley, S. P., Ligett, K., Weyl, G., and Xefteris, D. (2023). Efficiency in collective decision-making via quadratic transfers. *arXiv preprint arXiv:2301.06206*.
- Fearon, J. and Wendt, A. (2002). Rationalism and constructivism in international relations theory. *Handbook of International Relations Theory*. London, Sage Publications, pages 52–72.
- Fox, J. (2006). Legislative cooperation among impatient legislators. *Journal of Theoretical Politics*, 18(1):68–97.
- Fox, J. and Polborn, M. (2022). Veto institutions, hostage-taking, and tacit cooperation. *American Journal of Political Science*.
- Heisenberg, D. (2005). The institution of ‘consensus’ in the european union: Formal versus informal decision-making in the council. *European Journal of Political Research*, 44(1):65–90.
- Jackson, M. O. and Sonnenschein, H. F. (2007). Overcoming incentive constraints by linking decisions. *Econometrica*, 75(1):241–257.
- Kalandrakis, A. (2004). A three-player dynamic majoritarian bargaining game. *Journal of Economic Theory*, 116(2):294–14.
- Keohane, R. O. (1986). Reciprocity in international relations. *International Organization*, 40(1):1–27.

- Kikuchi, K. and Koriyama, Y. (2023). The winner-take-all dilemma. *Theoretical Economics*, 18(3):917–940.
- Kocherlakota, N. R. (1996). Implications of efficient risk sharing without commitment. *The Review of Economic Studies*, 63(4):595–609.
- Lazrak, A. and Zhang, J. (2023). Democratic policy decisions with decentralized promises contingent on vote outcome. *arXiv preprint arXiv:2304.08008*.
- Levin, J. (2003). Relational incentive contracts. *American Economic Review*, 93(3):835–857.
- Lewis, J. (2000). The methods of community in eu decision-making and administrative rivalry in the council’s infrastructure. *Journal of European Public Policy*, 7(2):261–289.
- 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 reputations: long-run relationships*. Oxford University Press.
- Mattila, M. and Lane, J.-E. (2001). Why unanimity in the council? a roll call analysis of council voting. *European Union Politics*, 2(1):31–52.
- May, K. O. (1952). A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica*, 20(4):680–684.
- Posner, E. A. and Sykes, A. O. (2014). Voting rules in international organizations. *Chicago Journal of International Law*, 15(1):195.
- Rae, D. W. (1969). Decision-rules and individual values in constitutional choice. *American Political Science Review*, 63(1):40–56.
- Riker, W. H. and Brams, S. J. (1973). The paradox of vote trading. *American Political Science Review*, 67(4):1235–1247.
- Schmitz, P. W. and Tröger, T. (2012). The (sub-) optimality of the majority rule. *Games and Economic Behavior*, 74(2):651–665.
- Urfalino, P. (2014). The rule of non-opposition: Opening up decision-making by consensus. *Journal of Political Philosophy*.

Urfalino, P. (2021). *Décider ensemble: la fabrique de l'obligation collective*. Seuil.

Warntjen, A. (2010). Between bargaining and deliberation: decision-making in the council of the european union. *Journal of European Public Policy*, 17(5):665–679.

A Proofs

A.1 Proof of **Theorem 1**

Proof. In the sequel, we denote by \mathbf{u}^t the reform drawn at stage t , by d^t the decision taken at stage t , and by $H = \cup_{T \geq 1} \{(\mathbf{u}^t, d^t)_{1 \leq t \leq T}\}$ the set of histories in the repeated game.

Definition 1. A payoff $w \in \mathbb{R}$ is (symmetrically) decomposable on set $W \subset \mathbb{R}$ if there exist three symmetric functions on S : a decision rule $d^w : S \mapsto \{0, 1\}$ and continuation payoffs $r^w : S \mapsto W$ (reward) and $q^w : S \mapsto W$ (punishment) such that:

$$w = \mathbb{E}[(1 - \delta)d^w(\mathbf{u})\bar{u} + \delta r^w(\mathbf{u})] \quad (11)$$

and

$$\forall \mathbf{u} \in S, \quad (1 - \delta)d^w(\mathbf{u})u_p + \delta r^w(\mathbf{u}) \geq (1 - \delta)(1 - d^w(\mathbf{u}))u_p + \delta q^w(\mathbf{u}). \quad (12)$$

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

In words, w is decomposable by (d^w, r^w, q^w) if (i) it is the expected average payoff given by the rule d^w and on-path reward r^w and (ii) for any utility realization \mathbf{u} , implementing the action $d^w(\mathbf{u})$ is rational for the pivot given a reward $r^w(\mathbf{u})$ and a punishment $q^w(\mathbf{u})$. For the sequel, it is useful to note that the rationality condition (12) is equivalent to:

$$\forall \mathbf{u} \in S, \quad (1 - \delta)(2d^w(\mathbf{u}) - 1)u_p + \delta(r^w(\mathbf{u}) - q^w(\mathbf{u})) \geq 0. \quad (13)$$

We now state and prove three lemmata before concluding the proof.

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 $\tau : W \times S \times \{0, 1\} \rightarrow W$, such that

$$\tau(w, \mathbf{u}, d) = \begin{cases} r^w(\mathbf{u}) & \text{if } d = d^w(\mathbf{u}) \\ q^w(\mathbf{u}) & \text{otherwise.} \end{cases}$$

We extend the transition function from $W \times S \times \{0, 1\}$ to $W \times H$, by recursively defining $\tau(w, \emptyset) = w$ and

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

We then define the (history-dependent) decision rule d by $d(h, \mathbf{u}) = f(\tau(w^0, h))(\mathbf{u})$ and the reward r and punishment q by $r(h, \mathbf{u}) = r^{\tau(w^0, h)}(\mathbf{u})$ and $q(h, \mathbf{u}) = q^{\tau(w^0, h)}(\mathbf{u})$. Let $D_i(h, \mathbf{u})$ be the utility difference in state $\tau(w^0, h)$ between following the decision $d(h, \mathbf{u})$ and deviating for an agent with utility u_i (assuming that the agent is pivotal). Formally:

$$\begin{aligned} D_i(h, \mathbf{u}) &= (1 - \delta)d(h, \mathbf{u})u_i + \delta r(h, \mathbf{u}) - \left((1 - \delta)(1 - d(h, \mathbf{u}))u_i + \delta q(h, \mathbf{u}) \right) \\ &= (1 - \delta)(2d(h, \mathbf{u}) - 1)u_i + \delta(r(h, \mathbf{u}) - q(h, \mathbf{u})). \end{aligned}$$

By construction, as W is self-generating, by application of (13), we must have $D_p(h, \mathbf{u}) \geq 0$ for the pivot, for any $h \in H$ and $\mathbf{u} \in S$.

Now, define individual strategy σ_i by:

$$\forall h \in H, \forall \mathbf{u} \in S, \quad \sigma_i(h, \mathbf{u}) = \begin{cases} d(h, \mathbf{u}) & \text{if } D_i(h, \mathbf{u}) \geq 0 \\ 1 - d(h, \mathbf{u}) & \text{otherwise.} \end{cases} \quad (14)$$

Observe first that the strategy profile $\sigma = (\sigma_i)_{i \in N}$ indeed implements the decision rule d :

- if $d(h, \mathbf{u}) = 1$, then for any i with $u_i \geq u_p$, we have $D_i(h, \mathbf{u}) \geq D_p(h, \mathbf{u}) \geq 0$ and thus $\sigma_i(h, \mathbf{u}) = d(h, \mathbf{u}) = 1$. Therefore, the strategy profile σ implements the decision $d = 1 = d(h, \mathbf{u})$ after history h for reform \mathbf{u} .
- if $d(h, \mathbf{u}) = 0$, then for any i with $u_i \leq u_p$, we have $D_i(h, \mathbf{u}) \geq D_p(h, \mathbf{u}) \geq 0$ and

thus $\sigma_i(h, \mathbf{u}) = d(h, \mathbf{u}) = 0$. Therefore, the strategy profile σ implements the decision $d = 0 = d(h, \mathbf{u})$ after history h for reform \mathbf{u} .

Then, as W is self-generating and F is symmetric, we have by application of (11):

$$\forall i \in N, \quad U_i(\sigma) = \mathbb{E}[(1 - \delta)d^{w^0}(\mathbf{u})u_i + \delta r^{w^0}(\mathbf{u})] = \mathbb{E}[(1 - \delta)d^{w^0}(\mathbf{u})\bar{u} + \delta r^{w^0}(\mathbf{u})] = w^0.$$

To conclude, let us show that σ is an equilibrium. By definition, the profile σ satisfies Symmetry. It also satisfies Anonymous Histories since the transition function τ only depend on the anonymized reform $(u_{[k]})_{1 \leq k \leq n}$ and the decision d , but not on individual votes (hence $\sigma_i(h, \mathbf{u})$ depends on h only through past anonymized reforms and decisions). Finally, by definition of σ_i in (14), it is clear that σ is a subgame-perfect equilibrium of the repeated game that satisfies the criterion of As-If-Pivotal Voting. \square

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

Proof. This proof is adapted from standard techniques exposed in [Mailath and Samuelson \(2006\)](#). By the [Lemma 1](#) above, it suffices to show that the set of equilibrium payoffs E is a self-generating set. For any payoff $w \in E$ with associated equilibrium σ , we let:

- the decision rule d^w be such that $d^w(\mathbf{u})$ is the decision implemented at the first stage under the profile σ for reform \mathbf{u} .
- the reward r^w be $r^w(\mathbf{u}) = U(\sigma \mid \mathbf{u}, d^w(\mathbf{u}))$
- the punishment q^w be $q^w(\mathbf{u}) = U(\sigma \mid \mathbf{u}, 1 - d^w(\mathbf{u}))$

As σ is a subgame-perfect equilibrium, the promises r^w and q^w take values in E , as desired. Moreover, we have $w = U(\sigma) = \mathbb{E}[(1 - \delta)d^w(\mathbf{u})\bar{u} + \delta r^w(\mathbf{u})]$ by definition of d^w and r^w . Finally, it is easy to see that (12) is satisfied since σ is an equilibrium:

- if $d^w(\mathbf{u}) = 1$, then there are at least m agents (where m is the majority rule threshold) for which

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

and thus

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

We thus obtain that (12) is satisfied.

- if $d^w(\mathbf{u}) = 0$, 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 $\lambda \in (0, 1)$, the payoff $w' = \lambda w + (1 - \lambda)U^0$ is also decomposable on W . Let w be decomposable by a triplet (d, r, q) on W .

We may write:

$$\begin{aligned} w &= U^0 + (w - U^0) \\ &= U^0 + \mathbb{E} \left[\underbrace{(1 - \delta)d(\mathbf{u})\bar{u} + \delta r(\mathbf{u}) - (1 - \delta)d^0(\mathbf{u})\bar{u} - \delta U^0}_{z(\mathbf{u})} \right] \\ &= U^0 + \int_S z(\mathbf{u})dF(\mathbf{u}). \end{aligned}$$

As F has no atom by assumption, there must exist a symmetric subdomain³¹ $S' \subset S$ such that $\int_{S'} z(\mathbf{u})dF(\mathbf{u}) = \lambda \int_S z(\mathbf{u})dF(\mathbf{u})$. Let us define (d', r', q') by:

$$d'(\mathbf{u}) = \begin{cases} d(\mathbf{u}) & \text{if } \mathbf{u} \in S' \\ d^0(\mathbf{u}) & \text{otherwise} \end{cases}, r'(\mathbf{u}) = \begin{cases} r(\mathbf{u}) & \text{if } \mathbf{u} \in S' \\ U^0 & \text{otherwise} \end{cases}, q'(\mathbf{u}) = \begin{cases} q(\mathbf{u}) & \text{if } \mathbf{u} \in S' \\ U^0 & \text{otherwise} \end{cases}.$$

As S' is symmetric, the functions d' , r' and q' are symmetric. As $U^0 \in W$, it is clear that both r' and q' take values in W . Moreover, since d^0 is the decision rule implemented by the stage-game equilibrium, it is clear by construction that the rationality condition (12) is satisfied by (d', r', q') . Finally, applying (11), the expected payoff is given by:

$$\begin{aligned} w' &= \int_{S'} ((1 - \delta)d(\mathbf{u})\bar{u} + \delta r(\mathbf{u})) dF(\mathbf{u}) + \int_{S \setminus S'} ((1 - \delta)d^0(\mathbf{u})\bar{u} + \delta U^0) dF(\mathbf{u}) \\ &= \int_{S'} z(\mathbf{u})dF(\mathbf{u}) + U^0 = \lambda(w - U^0) + U^0 = \lambda w + (1 - \lambda)U^0, \end{aligned}$$

³¹We say that $A \subset \mathbb{R}^N$ is symmetric if for any $\mathbf{u} \in A$, for any permutation $\pi : N \rightarrow N$, we have $\mathbf{u}_\pi \in A$.

as desired. Hence $w' = \lambda w + (1 - \lambda)U^0$ 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 $W = [\underline{w}, \bar{w}]$ be a compact interval. Let $\kappa = \frac{\delta}{1-\delta}(\bar{w} - \underline{w})$. We define $(d^{\max}, r^{\max}, q^{\max})$ by:

$$d^{\max}(\mathbf{u}) = d_{\kappa}^e(\mathbf{u}), \quad r^{\max}(\mathbf{u}) = \bar{w}, \quad q^{\max}(\mathbf{u}) = \underline{w}.$$

In words, d^{\max} is the κ -capped efficient rule defined in (7), i.e. it is efficient whenever the pivot stake is below κ , and coincides with the sincere rule otherwise (see also Figure 2, left panel). The reward r^{\max} is always the highest possible and the punishment q^{\max} is always the lowest possible. Clearly, by definition of κ , the triplet $(d^{\max}, r^{\max}, q^{\max})$ satisfies the rationality condition (12). The payoff induced by the triplet $(d^{\max}, r^{\max}, q^{\max})$ is given by:

$$w^{\max} = (1 - \delta)U(d_{\kappa}^e) + \delta\bar{w}.$$

Now, suppose that there is a triplet (d', r', q') which generates a higher payoff $w' > w^{\max}$. For any $\mathbf{u} \in S$, we have $r'(\mathbf{u}) \leq \bar{w} = r^{\max}(\mathbf{u})$, since W is upper bounded by \bar{w} . For any $\mathbf{u} \in S$ with $|u_p| \leq \kappa$, we have $d'(\mathbf{u})\bar{u} \leq d^{\max}(\mathbf{u})\bar{u} = d^e(\mathbf{u})\bar{u}$, since d^{\max} is efficient for such reforms. Thus, to have $w' > w^{\max}$, there must exist \mathbf{u} with $|u_p| > \kappa$ and such that $d'(\mathbf{u})\bar{u} > d^{\max}(\mathbf{u})\bar{u} = d^0(\mathbf{u})\bar{u}$. Suppose without loss of generality that $\bar{u} > 0$. Then we have $d^0(\mathbf{u}) = 0$, and thus $u_p < -\kappa$ and $d'(\mathbf{u}) = 1$. Now, we have:

$$\begin{aligned} (1 - \delta)(2d'(\mathbf{u}) - 1)u_p + \delta(r'(\mathbf{u}) - q'(\mathbf{u})) &\leq (1 - \delta)u_p + \delta(\bar{w} - \underline{w}) \\ &< \delta(\bar{w} - \underline{w}) - (1 - \delta)\kappa = 0. \end{aligned}$$

Thus, the triplet (d', r', q') fails the rationality condition (13). We conclude that w^{\max} is the highest payoff that can be decomposed on W .

Consider now the triplet $(d^{\min}, r^{\min}, q^{\min})$ defined by:

$$d^{\min}(\mathbf{u}) = d_{\kappa}^i(\mathbf{u}), \quad r^{\min}(\mathbf{u}) = \underline{w} + \frac{1 - \delta}{\delta}c(\mathbf{u}, d_{\kappa}^i(\mathbf{u})), \quad q^{\min}(\mathbf{u}) = \underline{w}.$$

In words, d^{\min} is the κ -capped semi-inefficient rule defined in (9), i.e. it is inefficient whenever the pivot stake is below κ and the pivot utility is of the same sign as the average utility but with a lower absolute value, and it coincides with the sincere rule otherwise (see also Figure 2, right panel). The reward r^{\min} is the minimum needed to provide incentives to the pivot to vote against her will when the prescribed decision

is inefficient and insincere, it is the lowest possible payoff otherwise. The punishment q^{\min} is always the lowest possible.

Clearly, by construction, the triplet $(d^{\min}, r^{\min}, q^{\min})$ satisfies the rationality condition (12), and by definition of d_{κ}^i , we know that for any $\mathbf{u} \in S$, we have $0 \leq c(\mathbf{u}, d_{\kappa}^i(\mathbf{u})) \leq \kappa$, so that $r^{\min}(\mathbf{u}) \in W$. Now, suppose that there is a triplet (d', r', q') which generates a lower payoff $w' < w^{\min}$, where w^{\min} is the expected payoff generated by $(d^{\min}, r^{\min}, q^{\min})$, that is:

$$w^{\min} = \mathbb{E}[(1 - \delta)d^{\min}(\mathbf{u})\bar{u} + \delta r^{\min}(\mathbf{u})] = (1 - \delta)(U(d_{\kappa}^i) + C(d_{\kappa}^i)) + \delta \underline{w}.$$

Then, there must exist $\mathbf{u} \in S$ such that the following condition holds:

$$(1 - \delta)d'(\mathbf{u})\bar{u} + \delta r'(\mathbf{u}) < (1 - \delta)d^{\min}(\mathbf{u})\bar{u} + \delta r^{\min}(\mathbf{u}). \quad (15)$$

Condition (15) can only be satisfied if at least one of the two following inequalities hold: $d^{\min}(\mathbf{u}) = d^e(\mathbf{u})$ or $r^{\min}(\mathbf{u}) > \underline{w}$. We thus consider two cases:

- if $r^{\min}(\mathbf{u}) > \underline{w}$, then by construction r^{\min} , we must have $c(\mathbf{u}, d_{\kappa}^i(\mathbf{u})) > 0$. We thus have $|u_p| < \kappa$ and (abusing notation, as u_p and \bar{u} can be either both positive or both negative) we have $u_p \in (0, \bar{u})$. Under these conditions, we know that $d^{\min}(\mathbf{u}) = 1 - d^e(\mathbf{u}) = 1 - d^0(\mathbf{u})$. It follows that $d'(\mathbf{u})\bar{u} \geq d^{\min}(\mathbf{u})\bar{u}$, so that condition (15) implies that $r'(\mathbf{u}) < r^{\min}(\mathbf{u}) = \underline{w} + \frac{1-\delta}{\delta}|u_p|$. There are two subcases to consider:
 - if $d'(\mathbf{u}) = d^{\min}(\mathbf{u}) = 1 - d^0(\mathbf{u})$, we obtain a contradiction by observing that, since $r'(\mathbf{u}) < r^{\min}(\mathbf{u})$ and $q'(\mathbf{u}) \geq \underline{w} = q^{\min}(\mathbf{u})$, we must have $r'(\mathbf{u}) - q'(\mathbf{u}) < r^{\min}(\mathbf{u}) - q^{\min}(\mathbf{u})$. Therefore, the triplet (d', r', q') cannot satisfy the rationality condition (12) at \mathbf{u} (recall that by construction, $r^{\min}(\mathbf{u}) - q^{\min}(\mathbf{u})$ is the smallest difference between continuations payoffs that can incentivize the pivot agent to deviate from sincere voting).
 - if $d'(\mathbf{u}) = d^0(\mathbf{u}) = d^e(\mathbf{u})$, we obtain a contradiction with (15), as since $u_p \in (0, \bar{u})$, we have:

$$\begin{aligned} (1 - \delta)d'(\mathbf{u})\bar{u} + \delta r'(\mathbf{u}) &\geq (1 - \delta)d^e(\mathbf{u})\bar{u} + \delta \underline{w} \\ &> (1 - \delta)((1 - d^e(\mathbf{u}))\bar{u} + |u_p|) + \delta \underline{w} \\ &> (1 - \delta)d^{\min}(\mathbf{u})\bar{u} + \delta r^{\min}(\mathbf{u}). \end{aligned}$$

- if $d^{\min}(\mathbf{u}) = d^e(\mathbf{u})$, then we have (by construction of the capped semi-inefficient

rule d^{\min}) that $c(\mathbf{u}, d^{\min}(\mathbf{u})) = 0$ and thus $r^{\min}(\mathbf{u}) = \underline{w}$ and $d^0(\mathbf{u}) = d^e(\mathbf{u})$. For (15) to hold, as $r'(\mathbf{u}) \geq \underline{w}$, we must have $d'(\mathbf{u}) = 1 - d^e(\mathbf{u}) = 1 - d^0(\mathbf{u})$. We consider two subcases:

- if $|u_p| \leq \kappa$, then we also have (by construction of the capped semi-inefficient rule d^{\min}) that $|u_p| > |\bar{u}|$. For the triplet (d', r', q') to satisfy the rationality condition (12), we must have $r'(\mathbf{u}) - q'(\mathbf{u}) \geq \frac{1-\delta}{\delta}|u_p|$, and thus $r'(\mathbf{u}) \geq \underline{w} + \frac{1-\delta}{\delta}|u_p|$. We obtain a contradiction with (15), as since $|u_p| > |\bar{u}|$, we have:

$$\begin{aligned} (1 - \delta)d'(\mathbf{u})\bar{u} + \delta r'(\mathbf{u}) &\geq (1 - \delta)((1 - d^e(\mathbf{u}))\bar{u} + |u_p|) + \delta \underline{w} \\ &> (1 - \delta)d^e(\mathbf{u})\bar{u} + \delta \underline{w} \\ &> (1 - \delta)d^{\min}(\mathbf{u})\bar{u} + \delta r^{\min}(\mathbf{u}). \end{aligned}$$

- if $|u_p| > \kappa$, then we obtain a contradiction by observing that, since $r'(\mathbf{u}) - q'(\mathbf{u}) \leq \bar{w} - \underline{w} = \frac{1-\delta}{\delta}\kappa$, the triplet (d', r', q') cannot satisfy the rationality condition (12) at \mathbf{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 \bar{I} is also self-generating.

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

To conclude, we know that the set of equilibrium payoffs E contains U^0 (repeating the stage-game equilibrium is an equilibrium of the repeated game). 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. This concludes the proof of Lemma 3. \square

We are now equipped to finish the proof of Theorem 1. We know from the previous lemma that $E = [\underline{w}, \bar{w}]$. Let w^{\max} the maximal payoff that can be decomposed on E . As E is self-generating, we have that $w^{\max} \geq \bar{w}$. Now, if $w^{\max} > \bar{w}$, then this would contradict the fact that E must be the largest self-generating set. Hence, we must

have $w^{\max} = \bar{w}$. Similarly, if we note w^{\min} the minimal payoff that can be decomposed on E , we have $w^{\min} = \underline{w}$.

The formula for w^{\max} in the proof is $w^{\max} = (1 - \delta)U(d_\kappa^e) + \delta\bar{w}$. We thus obtain $w^{\max} = U(d_\kappa^e) = \bar{V}(\kappa)$. The formula for w^{\min} in the proof is $w^{\min} = (1 - \delta)(U(d_\kappa^i) + C(d_\kappa^i)) + \delta\underline{w}$. We thus obtain $w^{\min} = U(d_\kappa^i) + C(d_\kappa^i) = \underline{V}(\kappa)$.

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

A.2 Proof of Proposition 1

Proof. First, it is straightforward that the highest consensus will be achieved at an optimal equilibrium if, for any reform \mathbf{u} with prescribed decision $d_{\kappa_\delta}^e(\mathbf{u}) = 1$, the reward will be the highest possible, equal to \bar{w}_δ , and the punishment will be the lowest possible, equal to \underline{w}_δ . Thus, such an optimal equilibrium will be stationary, implementing a stage-game profile $v = (v_i)_{i \in N}$ with $v_i(\mathbf{u}) = \mathbb{1}\{u_i \geq -\kappa_\delta\}$ whenever $d_{\kappa_\delta}^e(\mathbf{u}) = 1$.

As mentioned in the text below Proposition 1, it is straightforward that more good reforms will be adopted and that each such reform is more likely to receive consensus when δ increases. Therefore, to show that P_δ is increasing, it suffices to show that a bad reform never receives consensus (even if approved).

Let \mathbf{u} be a bad reform, i.e $\bar{u} \leq 0$ and assume that it is approved at an optimal equilibrium, i.e. $d_{\kappa_\delta}^e(\mathbf{u}) = 1$. Then, by definition of the capped efficient rule, it must be that $u_p > \kappa_\delta$. Now observe that we can upper bound the lowest utility $u_{[1]}$ by lower bounding the average utility

$$0 \geq \bar{u} \geq \lceil mn \rceil u_p + (n - \lceil mn \rceil) u_{[1]} > \lceil mn \rceil \kappa_\delta + (n - \lceil mn \rceil) u_{[1]}.$$

We thus have $u_{[1]} < -\frac{\lceil mn \rceil}{n - \lceil mn \rceil} \kappa_\delta \leq -\kappa_\delta$, since $m \geq 1/2$. It follows that reform \mathbf{u} cannot be unanimously approved at an optimal equilibrium.

We have thus shown that P_δ is (weakly) increasing. As the support S of reform distribution F is bounded, while the function κ_δ is increasing and unbounded (by application of (2), since $\bar{w}_\delta > U^0 \geq \underline{w}_\delta$ for $\delta > \delta^F$), there must exist $\delta^C \in (0, 1)$ above

which we have $\kappa_\delta \geq -\min_{\mathbf{u} \in S} u_{[1]}$. Then, consensus is maximal and any good reform is adopted with consensus. □

A.3 Proof of **Theorem 2**

Proof. The first step of the proof consists in providing closed-form formulas and comparative statics results for all variables of interest (κ_δ , \bar{w}_δ , discount factor thresholds) for any possible values of discount factor δ and inefficiency gap Δ (throughout the proof we abuse notation and write Δ for Δ^e). It will be useful to divide our inquiry into two cases, corresponding to two regions of values for Δ , that is $\Delta \in (0, 1/2]$, and then $\Delta \in (1/2, 1]$.

A.3.1 Preliminaries

Note that for any $\kappa \in [0, \Delta]$, we may write

$$\bar{V}(\kappa) = \int_{u_c \geq \kappa} \bar{u} dF(\mathbf{u}) = \int_{\theta \geq \Delta - \kappa} \theta \frac{d\theta}{2} = \frac{1 - (\kappa - \Delta)^2}{4}.$$

Hence, we obtain the general formula

$$\bar{V}(\kappa) = \begin{cases} \frac{1 - (\kappa - \Delta)^2}{4} & \text{if } 0 \leq \kappa \leq \Delta \\ \frac{1}{4} & \text{if } \kappa \geq \Delta. \end{cases} \quad (16)$$

Note that when $\Delta \leq 1 - \kappa$, we have

$$\underline{V}(\kappa) = \int_{\Delta}^{\Delta + \kappa} (\theta - \Delta) \frac{d\theta}{2} + \int_{\Delta + \kappa}^1 \theta \frac{d\theta}{2} = \frac{1 - \Delta(\Delta + 2\kappa)}{4},$$

while when $1 - \kappa \leq \Delta \leq 1$, we have

$$\underline{V}(\kappa) = \int_{\Delta}^1 (\theta - \Delta) \frac{d\theta}{2} = \frac{(1 - \Delta)^2}{4}.$$

Hence, we obtain the general formula

$$\underline{V}(\kappa) = \begin{cases} \frac{1 - \Delta(\Delta + 2\kappa)}{4} & \text{if } \kappa \leq 1 - \Delta \\ \frac{(1 - \Delta)^2}{4} & \text{if } \kappa \geq 1 - \Delta. \end{cases} \quad (17)$$

A.3.2 Equilibrium Payoffs

We separate our analysis between Case 1, $\Delta \in (0, 1/2]$, and Case 2, $\Delta \in (1/2, 1]$.

Case 1: $\Delta \in (0, 1/2]$


In that case, we have $\Delta \leq 1 - \Delta$. By application of (16) and (17), we obtain

$$\bar{V}(\kappa) - \underline{V}(\kappa) = \begin{cases} \kappa \left(\Delta - \frac{\kappa}{4} \right) & \text{if } \kappa \leq \Delta \\ \frac{\Delta(\Delta + 2\kappa)}{4} & \text{if } \Delta \leq \kappa \leq 1 - \Delta \\ \frac{\Delta(2 - \Delta)}{4} & \text{if } \kappa \geq 1 - \Delta. \end{cases}$$

Observe that the function $\bar{V}(\kappa) - \underline{V}(\kappa)$ is strictly concave for $\kappa \in [0, \Delta]$, then affine for $\kappa \in [\Delta, 1 - \Delta]$, and finally constant for $\kappa \in [1 - \Delta, +\infty)$. Moreover, its derivative is continuous at $k = \Delta$. Equation (6) thus admits no solution (if δ is too low) or a unique solution (if δ is high enough).

The determination of κ_δ in Case 1 is illustrated in Figure 4 below. The colored curve corresponds to the difference $\bar{V}(\kappa) - \underline{V}(\kappa)$, while the grey lines represent functions $\kappa \mapsto (1 - \delta)\kappa/\delta$, for different values of the discount factor δ . As κ increases, both the highest and the lowest sustainable payoff increase. Once κ reaches Δ , the highest sustainable payoff reaches its maximum $U^e = 1/4$, while the lowest sustainable payoff keeps decreasing (green part). For κ larger than $1 - \Delta$, the lowest sustainable payoff reaches its minimum $(1 - \Delta)^2/4$ (blue part).

Figure 4: Determination of κ_δ in Case 1.



The thresholds δ^P and δ^F are then defined by,

$$\frac{1 - \delta^P}{\delta^P} = \frac{\partial (\bar{V}(\kappa) - \underline{V}(\kappa))}{\partial \kappa} \Big|_{\kappa=0} = \Delta, \quad \frac{1 - \delta^F}{\delta^F} = \frac{(\bar{V}(\Delta) - \underline{V}(\Delta))}{\Delta} = \frac{3\Delta}{4}. \quad (18)$$

We identify three regimes of cooperation as a function of δ .

Regime 1: $\delta \leq \delta^P$. The grey line lies above the colored curve for any positive value of κ , and we get that $\kappa_\delta = 0$; there is no cooperation and the unique equilibrium payoff is U^0 .

Regime 2: $\delta \in (\delta^P, \delta^F)$. The grey line intersects the colored curve for a value of $\kappa_\delta \in (0, \Delta)$ (red part) defined by,

$$\kappa_\delta = 4 \left(\Delta - \frac{1 - \delta}{\delta} \right). \quad (19)$$

Applying equation (16), the associated highest equilibrium payoff is given by,

$$\bar{w}_\delta = \frac{1 - (\kappa - \Delta)^2}{4} = U^e - 4 \left(\frac{1 - \delta}{\delta} - \frac{3\Delta}{4} \right)^2. \quad (20)$$

We then get a regime of partial cooperation as $\bar{w}_\delta \in (U^0, U^e)$.

Regime 3: $\delta \geq \delta^F$. The grey line intersects either the green or the blue line for a value of κ greater than Δ . In that case, we get full cooperation as $\bar{w}_\delta = U^e$.

Case 2: $\Delta \in (1/2, 1]$


In that case, we have $\Delta \geq 1 - \Delta$. By application of (16) and (17), we obtain

$$\bar{V}(\kappa) - \underline{V}(\kappa) = \begin{cases} \kappa \left(\Delta - \frac{\kappa}{4} \right) & \text{if } \kappa \leq 1 - \Delta \\ \frac{-\kappa^2 + 2\Delta(1 + \kappa) - 2\Delta^2}{4} & \text{if } 1 - \Delta \leq \kappa \leq \Delta \\ \frac{\Delta(2 - \Delta)}{4} & \text{if } \kappa \geq \Delta. \end{cases}$$

Observe that the function $\bar{V}(\kappa) - \underline{V}(\kappa)$ is strictly concave for $\kappa \in [0, 1 - \Delta]$, then strictly concave for $\kappa \in [1 - \Delta, \Delta]$, and finally constant for $\kappa \in [\Delta, +\infty)$. Moreover, its derivative is (discontinuously) decreasing at $k = 1 - \Delta$. Equation (6) thus admits no solution (if δ is too low) or a unique solution (if δ is high enough).

The determination of κ_δ in Case 2 is illustrated in Figure 5 below. As κ increases, both the highest and the lowest sustainable payoff increase. Once κ reaches $1 - \Delta$, the lowest sustainable payoff reaches its minimum $(1 - \Delta)^2/4$, while the highest sustainable payoff keeps increasing (green part). For κ larger than Δ , the highest sustainable payoff reaches its maximum $1/4$ (blue part).

Figure 5: Determination of κ_δ in Case 2.



Let threshold δ^P be as in (18), and define new thresholds δ^R and δ^F by,

$$\begin{cases} \frac{1-\delta^R}{\delta^R} = \frac{(\bar{V}(1-\Delta) - \underline{V}(1-\Delta))}{\Delta} = \frac{5\Delta - 1}{4}, \\ \frac{1-\delta^F}{\delta^F} = \frac{(\bar{V}(\Delta) - \underline{V}(\Delta))}{\Delta} = \frac{2 - \Delta}{4}. \end{cases} \quad (21)$$

We identify four regimes of cooperation as a function of δ .

Regime 1: $\delta \leq \delta^P$. The grey line lies above the colored curve for any positive value of κ , and we get that $\kappa_\delta = 0$.

Regime 2: $\delta \in (\delta^P, \delta^R)$. The grey line intersects the colored curve for a value $\kappa_\delta < 1 - \Delta$ (red part) defined by equation (19), as in Case 1. The highest equilibrium payoff \bar{w}_δ is given by formula (20).

Regime 3: $\delta \in (\delta^R, \delta^F)$. The grey line intersects the colored curve for a value $\kappa_\delta \in (1 - \Delta, \Delta]$ (green part) defined by,

$$(1 - \delta)\kappa_\delta = \delta \left(\frac{-\kappa^2 + 2\Delta(1 + \kappa_\delta) - 2\Delta^2}{4} \right).$$

Simplifying, we obtain the quadratic equation,

$$\kappa_\delta^2 + 4 \left(\frac{1-\delta}{\delta} - \frac{\Delta}{2} \right) \kappa_\delta + 2\Delta(\Delta - 1) = 0. \quad (22)$$

Since $\Delta - 1 < 0$, equation (22) admits a unique positive solution, given by,

$$\kappa_\delta = \Delta - 2 \frac{1-\delta}{\delta} + \sqrt{\left(\Delta - 2 \frac{1-\delta}{\delta} \right)^2 + 2\Delta(1-\Delta)}. \quad (23)$$

The highest equilibrium payoff is now given by,

$$\bar{w}_\delta = \frac{1 - (\Delta - \kappa_\delta)^2}{4}. \quad (24)$$

Regime 4: $\delta \geq \delta^F$. The grey line intersects the colored curve for a value of κ greater than Δ (blue part). In that case, we get full cooperation as $\bar{w}_\delta = U^e$.

A.3.3 Comparative Statics

We now provide general comparative statics on cooperation thresholds and welfare levels, considering all possible values of $\Delta \in [0, 1]$ and $\delta \in (0, 1)$.

By application of equations (18) and (21), we have that thresholds δ^P and δ^R are decreasing with respect to Δ , while threshold δ^F is decreasing with Δ when $\Delta \leq 1/2$ and increasing with Δ when $\Delta \geq 1/2$.

When $\delta \leq \delta^P$ (regime 1 of both cases 1 and 2), we have $\kappa_\delta = 0$, so that $\bar{w}_\delta = U^e - \Delta^2/4$ is constant in δ and decreasing in Δ .

When $\delta^P < \delta < \min(\delta^R, \delta^F)$, we are in the regime 2 of both cases 1 and 2. Then, by equations (16) and (20), we obtain that \bar{w}_δ is increasing with δ and also increasing with Δ as,

$$\frac{\partial \bar{w}_\delta}{\partial \Delta} = 6 \left(\frac{1-\delta}{\delta} - \frac{3\Delta}{4} \right) > 0.$$

This last inequality is obtained by observing that,

$$\begin{aligned} \frac{1-\delta}{\delta} &> \frac{1-\delta^F}{\delta^F} = \frac{3\Delta}{4} && \text{when } \Delta \leq 1/2, \\ \frac{1-\delta}{\delta} &> \frac{1-\delta^R}{\delta^R} = \frac{5\Delta-1}{4} \geq \frac{3\Delta}{4} && \text{when } \Delta \geq 1/2. \end{aligned}$$

When $\Delta \geq 1/2$ and $\delta^R < \delta < \delta^F$, we are in the regime 3 of case 2. In that case, \bar{w}_δ is increasing with δ but decreasing with Δ . The first comparative statics is obtained

by observing that κ_δ is increasing in δ (Corollary 2) and applying equation (24). The second comparative statics is obtained by applying equation (24) and differentiating (22), i.e.

$$\frac{\partial(\Delta - \kappa)}{\partial\Delta} = \frac{\Delta + 2\left(\frac{1-\delta}{\delta}\right) - 1}{\kappa - \Delta + 2\left(\frac{1-\delta}{\delta}\right)}.$$

The numerator can be shown to be positive as $\delta < \delta^F$ implies, by application of (21), that,


$$2\frac{1-\delta}{\delta} > 1 - \frac{\Delta}{2} \geq 1 - \Delta.$$

The denominator is also positive by definition of κ_δ in (23). Hence, $\frac{\partial\bar{w}_\delta}{\partial\Delta} < 0$.

Finally, for $\delta \geq \delta^F$, we have $\bar{w}_\delta = U^e$, which is constant in both δ and Δ .

Figure 6 below summarizes the cooperation regimes, which are separated by the various discount thresholds, represented as functions of the inefficiency gap Δ . For each regime, the comparative statics of \bar{w}_δ as a function of δ and Δ are provided. We note that partial cooperation is easier when Δ increases, while full cooperation is most easily achieved when $\Delta = 1/2$.

Figure 6: Discount factor thresholds, cooperation regimes and comparative statics.



A.3.4 Optimal voting rule(s).

We now characterize the set of optimal voting rules, i.e. the voting rules which maximize the highest equilibrium payoff \bar{w}_δ .

For given values of $\mathbf{b} = (b_1, \dots, b_n)$ and α , the set of relevant voting rules $\mathcal{M}_n := \{m \in [1/2, 1] \mid mn \in \mathbb{N}\}$ can be described by the set of feasible inefficiency gaps $\Delta = \{\Delta_m \mid m \in \mathcal{M}_n\} = \{\alpha|b_1|, \dots, \alpha|b_{\frac{n+1}{2}}|\}$ with $\alpha|b_{\frac{n+1}{2}}| \leq \dots \leq \alpha|b_1|$.

When $\delta \leq \min_{\Delta \in \Delta} \delta^P(\Delta)$, the optimal voting rule is the one with the lowest inefficiency gap Δ , i.e. the simple majority rule $m = 1/2$.

When $\min_{\Delta \in \Delta} \delta^P(\Delta) < \delta < \min_{\Delta \in \Delta} \delta^F(\Delta)$, there are essentially two candidate optimal rules: the optimal rule among those that do not allow for partial cooperation, call it $m^{NC}(\delta)$, and the optimal rule among those that do allow for partial cooperation, call it $m^{PC}(\delta)$. The rule $m^{NC}(\delta)$, if it exists, is simply the one with the lowest inefficiency Δ , by application of the comparative results represented in [Figure 6](#). The rule $m^{PC}(\delta)$, if it exists, is either the one with the highest $\Delta \leq (\delta^R)^{-1}(\delta)$, call it $m^{R-}(\delta)$, or the one with the lowest $\Delta \geq (\delta^R)^{-1}(\delta)$, call it $m^{R+}(\delta)$.³²

Observe that the payoff under $m^{NC}(\delta)$ is constant in δ and strictly lower than U^e , while the payoff under $m^{PC}(\delta)$ is increasing in δ , and tends to U^e when δ approaches $\min_{\Delta \in \Delta} \delta^F(\Delta)$. Thus, there exists $\underline{\delta}$ such that: only $m^{NC}(\delta)$ is optimal when $\delta < \underline{\delta}$, only $m^{PC}(\delta)$ is optimal when $\delta > \underline{\delta}$, and both are optimal when $\delta = \underline{\delta}$.

To show that $m^{PC}(\delta)$ is (weakly) decreasing in δ , observe first that $m^{R-}(\delta)$ and $m^{R+}(\delta)$ are both (weakly) decreasing in δ . As we always have $m^{R-}(\delta) \leq m^{R+}(\delta)$ and $n(m^{R+}(\delta) - m^{R-}(\delta)) \leq 1$, the only way to have $m^{PC}(\cdot)$ increasing between two discount factors δ and $\delta' > \delta$ is to have $m^{PC}(\delta) = m^{R-}(\delta) = m^{R-}(\delta') < m^{PC}(\delta') = m^{R+}(\delta') = m^{R+}(\delta)$. Let us show that this is not possible.

If that were true, we would have two rules Δ and $\Delta' > \Delta$ with $\Delta \leq (\delta^R)^{-1}(\delta') \leq (\delta^R)^{-1}(\delta) \leq \Delta'$ and $\bar{w}_\delta(\Delta) > \bar{w}_\delta(\Delta')$ and $\bar{w}_{\delta'}(\Delta) < \bar{w}_{\delta'}(\Delta')$. This would imply by equation (24), which holds in both partial cooperation regimes, that $\kappa_{\delta'}(\Delta') - \kappa_{\delta'}(\Delta) > \Delta' - \Delta > \kappa_\delta(\Delta') - \kappa_\delta(\Delta)$. For that to be true we would need to have $\frac{\partial \kappa_\delta}{\partial \delta}(\Delta', \delta_0) > \frac{\partial \kappa_\delta}{\partial \delta}(\Delta, \delta_0)$ for at least some $\delta_0 \in (\delta, \delta')$. Deriving expression (22) for the

³²In all rigor, if $\delta \geq \delta^F(\Delta = 1/2)$, then $m^{PC}(\delta)$ is either the rule with the highest $\Delta \leq 1/2$ or with the highest $\Delta \geq 1/2$. Yet, the same reasoning applies for that case.

first derivative and expression (19) for the second one, this is equivalent to,

$$\frac{2\kappa_{\delta_0}(\Delta')}{\delta_0^2 \left(\kappa_{\delta_0}(\Delta') + 2\frac{1-\delta_0}{\delta_0} - \Delta' \right)} > \frac{4}{\delta_0^2}.$$

Applying the expression of $\kappa_{\delta_0}(\Delta')$ in (23), we obtain,

$$\sqrt{(\Delta' - 2\frac{1-\delta_0}{\delta_0})^2 + 2\Delta'(1-\Delta')} < \Delta' - 2\frac{1-\delta_0}{\delta_0},$$

which is impossible, hence a contradiction.

When $\delta \geq \min_{\Delta \in \Delta} \delta^F(\Delta) := \bar{\delta}$, there is at least one voting rule which allows for full cooperation and hence full efficiency. This rule cannot be higher than the rule $m^{PC}(\delta)$ for δ just below the threshold $\bar{\delta}$, by application of the previous argument and by observing that utilities are continuous in δ in this model. As δ increases, the set of optimal voting rules is convex and expanding, since $(\delta^F)^{-1}$ is quasi-convex (see Figure 6). Hence, the lowest optimal voting rule (weakly) decreases while the highest optimal voting rule (weakly) increases. Finally, when δ is high enough, all rules are optimal by the folk theorem (Corollary 3). □

A.4 Proof of Proposition 3

Proof. Consider the optimal degree of cooperation κ_δ in the first partial cooperation regime of cases 1 and 2 in the previous proof of Theorem 2. In this regime, $\kappa_\delta = 4(\Delta - \frac{1-\delta}{\delta}) = 4(\alpha|b_p| - \frac{1-\delta}{\delta})$. We thus have $\frac{\partial \kappa_\delta}{\partial \alpha} = 4|b_p|$. The optimal consensus probability is given by $P_\delta = \frac{1 - \max(\alpha - \kappa_\delta, 0)}{2}$. For δ small enough in the partial cooperation regime, we have $\kappa_\delta < \alpha$ and thus $\frac{\partial P_\delta}{\partial \alpha} = \frac{1}{2}(\frac{\partial \kappa_\delta}{\partial \alpha} - 1)$. In such a case, we obtain that $\frac{\partial P_\delta}{\partial \alpha} > 0$ whenever $4|b_p| > 1$, that is, when $|b_p| > 1/4$. □

A.5 Proof of Theorem 3

Proof. Assume first that the Gains from Trade Condition is satisfied, with a family of reforms $(\mathbf{u}^k)_{1 \leq k \leq K}$ and a family of weights $(w^k)_{1 \leq k \leq K}$. Observe first that we may assume without loss of generality that $u_p^k \neq 0$ for each k . Indeed, if it is the case that $u_p^k = 0$ for some value(s) of k , we can pick vector(s) \mathbf{x}^k arbitrarily close to \mathbf{u}^k such that $x_p^k \neq 0$, equation (10) remains satisfied, and $\mathbf{x}^k \in S$, since F is continuous. Similarly, we may assume without loss of generality that each reform \mathbf{u}^k belongs to the interior

of the support S .

Now, as the decision rule d^0 is constant in a small enough neighborhood of each \mathbf{u}^k , we obtain that there exists $\varepsilon > 0$ such that $\mathcal{B}(\mathbf{u}^k, \varepsilon) \subset S$ for each k and for any family $(\mathbf{x}^k)_{1 \leq k \leq K}$ with $\mathbf{x}^k \in \mathcal{B}(\mathbf{u}^k, \varepsilon)$ for each k , the following holds: $x_p^k u_p^k > 0$ and

$$\forall i \in N, \quad \sum_{k=1}^K w^k x_i^k (1 - 2d^0(\mathbf{x}^k)) > 0. \quad (25)$$

Therefore, there exists a family $(X_k)_{1 \leq k \leq K}$ of measurable sets such that $X_k \subseteq \mathcal{B}(\mathbf{u}^k, \varepsilon)$ for each k and for any k, k' , we have $\frac{F(X_k)}{F(X_{k'})} = \frac{w^k}{w^{k'}}$. Consider the decision rule d^* defined by:

$$\forall \mathbf{u} \in S, \quad d^*(\mathbf{u}) = \begin{cases} 1 - d^0(\mathbf{u}) & \text{if } \mathbf{u} \in \cup_{k=1}^K X_k \\ d^0(\mathbf{u}) & \text{if } \mathbf{u} \notin \cup_{k=1}^K X_k. \end{cases}$$

To show that d^* strictly Pareto-improves over d^0 , we write:

$$U_i(d^*) - U_i(d^0) = \int_S u_i(d^*(\mathbf{u}) - d^0(\mathbf{u})) dF(\mathbf{u}) = \sum_{k=1}^K \int_{X_k} u_i(1 - 2d^0(\mathbf{u})) dF(\mathbf{u}). \quad (26)$$

Next, observe that for any k , since the function d^0 is constant on X_k (as $x_p^k u_p^k > 0$ for any $\mathbf{x}^k \in X_k$), the function $\mathbf{u} \mapsto u_i(1 - 2d^0(\mathbf{u}))$ is continuous on X_k , so that, by the mean value theorem, there exists $\mathbf{x}^k \in X_k$ such that

$$\int_{X_k} u_i(1 - 2d^0(\mathbf{u})) dF(\mathbf{u}) = F(X_k) x_i^k (1 - 2d^0(\mathbf{x}^k)).$$

By combining equations (26) and (25), we obtain that $U_i(d^*) > U_i(d^0)$ for any $i \in N$. Finally, consider the strategy profile defined $\sigma_i(h, \mathbf{u}) = d^*(\mathbf{u})$ if $\mathbf{u} \in \cup_{k=1}^K X_k$ and $\sigma_i(h, \mathbf{u}) = \mathbb{1}\{u_i > 0\}$ if $\mathbf{u} \notin \cup_{k=1}^K X_k$ for any h on the equilibrium path (i.e. such that d^* has been implemented at each previous stage), and $\sigma_i(h, \mathbf{u}) = \mathbb{1}\{u_i > 0\}$ for any h off the equilibrium path. When δ is close enough to 1, it is clear that σ is an outcome-stationary equilibrium, and it achieves the payoff vector $(U_i(d^*))_{i \in N}$, which strictly Pareto-dominates $(U_i(d^0))_{i \in N}$.

Conversely, suppose that for some $\delta \in (0, 1)$, an outcome-stationary equilibrium exists which strictly Pareto-improves over sincere voting. Let d^* be the decision rule implemented at each stage by this equilibrium. Let $X = \{\mathbf{u} \in S \mid d^*(\mathbf{u}) \neq d^0(\mathbf{u})\}$. By assumption, we have for any $i \in N$, $\int_S u_i(d^*(\mathbf{u}) - d^0(\mathbf{u})) dF(\mathbf{u}) > 0$, which is equivalent to $\int_X u_i(1 - 2d^0(\mathbf{u})) dF(\mathbf{u}) > 0$. We may assume without loss of generality that there

is $\eta > 0$ such that $|u_p| > \eta$ for any $\mathbf{u} \in X$, as the previous set of strict inequalities is preserved if we remove from X all vectors such that $|u_p| \leq \eta$ for η small enough.

Let $\varepsilon = \min(\eta, \min_{i \in N} \int_X u_i(1 - 2d^0(\mathbf{u}))dF(\mathbf{u})) > 0$. Let $(X_k)_{1 \leq k \leq K}$ be a family of measurable sets such that $X = \cup_{k=1}^K X_k$ is a partition and $\forall k \in \{1, \dots, K\}$ there exists $\mathbf{u}^k \in V$ such that $X_k \subset \mathcal{B}(\mathbf{u}^k, \varepsilon)$. Then, as $\varepsilon \leq \eta$, the function $\mathbf{u} \mapsto u_i(1 - 2d^0(\mathbf{u}))$ is continuous on X_k , and we obtain:

$$\begin{aligned} \forall i \in N, \quad \int_X u_i(1 - 2d^0(\mathbf{u}))dF(\mathbf{u}) &= \sum_{k=1}^K \int_{X_k} u_i(1 - 2d^0(\mathbf{u}))dF(\mathbf{u}) \\ &= \sum_{k=1}^K F(X_k)x_i^{k,i}(1 - 2d^0(\mathbf{x}^{k,i})) \end{aligned}$$

for some vectors $\mathbf{x}^{k,i} \in X_k \subset \mathcal{B}(\mathbf{u}^k, \varepsilon)$. Let $w^k = F(X_k) \geq 0$, and note that $\sum_{k=1}^K w^k \leq 1$. As $|1 - 2d^0(\mathbf{u})| \leq 1$ for any $\mathbf{u} \in X$, we obtain:

$$\begin{aligned} \forall i \in N, \quad \sum_{k=1}^K w^k u_i^k(1 - 2d^0(\mathbf{u}^k)) &= \sum_{k=1}^K w^k u_i^k(1 - 2d^0(\mathbf{x}^{k,i})) \\ &\geq \sum_{k=1}^K w^k \left(x_i^{k,i}(1 - 2d^0(\mathbf{x}^{k,i})) - \varepsilon \right) \\ &\geq \left(\sum_{k=1}^K w^k x_i^{k,i}(1 - 2d^0(\mathbf{x}^{k,i})) \right) - \varepsilon \\ &\geq \int_X u_i(1 - 2d^0(\mathbf{u}))dF(\mathbf{u}) - \varepsilon > 0. \end{aligned}$$

Thus, the Gains from Trade Condition is satisfied. \square

A.6 Proof of Proposition 4

Proof. Let $m > \frac{1}{2} + \frac{1}{n}$ be a strict super-majority rule and fix $\xi > 0$. We will show that the Gains from Trade Condition is satisfied for $\bar{\varepsilon} = 0$, which implies that it is satisfied for any $\bar{\varepsilon} \geq 0$. We will show that the condition is satisfied with a family of two reforms $(\mathbf{u}^1, \mathbf{u}^2)$. To construct the first reform, let $s^1 = 1$ (the status quo is Left) and $\theta^1 = -\beta_{\lfloor(1-m)n\rfloor} - \xi$, so that for any $i \in N$, $u_i^1 = \beta_i - \beta_{\lfloor(1-m)n\rfloor} - \xi$. For that reform, agent $\lfloor(1-m)n\rfloor$ (with utility $-\xi$) is pivotal, so that $d^0(\mathbf{u}^1) = 0$, i.e. the reform is rejected under sincere voting. Symmetrically, let $s^2 = -1$ (the status quo is Right) and $\theta^2 = -\beta_{\lceil mn \rceil} + \xi$, so that for any $i \in N$, $u_i^2 = \beta_{\lceil mn \rceil} - \beta_i - \xi$. We have $d^0(\mathbf{u}^2) = 0$.

To conclude, we have for any $i \in N$, $u_i^1 + u_i^2 = \beta_{\lceil mn \rceil} - \beta_{\lfloor (1-m)n \rfloor} - 2\xi$. By taking ξ small enough, we obtain that reforms \mathbf{u}^1 and \mathbf{u}^2 belong to the support of F and that for any $i \in N$, $u_i^1 + u_i^2 > 0$. Hence, (F, m) satisfies the Gains from Trade Condition.

Consider now the simple majority rule $m = 1/2$. Assume first that $\bar{\varepsilon} = 0$. On any reform, the pivotal agent is then the median $\frac{n+1}{2}$, i.e. the one with the median bias $\beta_{\frac{n+1}{2}}$. This agent has her favorite alternative accepted under the sincere voting rule d^0 , for any reform $\mathbf{u} \in S$. It is thus impossible to strictly improve her utility by reversing some sincere collective decisions. Therefore (F, m) violates the Gains from Trade Condition when $\bar{\varepsilon} = 0$.

Assume now that $\bar{\varepsilon} \geq \frac{1}{2}(\beta_{\frac{n+3}{2}} - \beta_{\frac{n-1}{2}})$. To obtain the result, it will be sufficient to show that the Gains from Trade Condition is satisfied for this $\bar{\varepsilon}$ (indeed, observe that the condition is then satisfied for any $\bar{\varepsilon}' \geq \bar{\varepsilon}$). We fix $\xi > 0$. It will be useful to note $i_1 := \frac{n-1}{2}$, $i_2 := \frac{n+1}{2}$ and $i_3 := \frac{n+3}{2}$ and to define for any $i, j \in \{i_1, i_2, i_3\}$, $\beta_{ij} := \frac{\beta_i + \beta_j}{2}$. By assumption, we have $\bar{\varepsilon} \geq \frac{1}{2}(\beta_{i_3} - \beta_{i_1})$, and we may further assume without loss of generality that $\beta_{i_2} \geq \beta_{i_1 i_3}$ (the opposite case can be treated similarly). We will show that the Gains from Trade Condition is satisfied with a family of three reforms $(\mathbf{u}^1, \mathbf{u}^2, \mathbf{u}^3)$, weighted by $\mathbf{w} = (2, 1, 1)$:

- reform 1 is such that $s = -1$, $\varepsilon_i = 0$ for $i \neq i_2, i_3$, $\tilde{\beta}_{i_2} = \tilde{\beta}_{i_3} = \beta_{i_2 i_3}$ and $\theta = \xi - \beta_{i_2 i_3}$. It follows that $u_{i_2}^1 = u_{i_3}^1 = -\xi$ and $u_i^1 = \beta_{i_2 i_3} - \beta_i - \xi$ for $i \neq i_2, i_3$.
- reform 2 is such that $s = 1$, $\varepsilon_i = 0$ for $i \neq i_1, i_2, i_3$, $\tilde{\beta}_{i_1} = \tilde{\beta}_{i_3} = \beta_{i_1 i_3}$, $\tilde{\beta}_{i_2} = \beta_{i_2} + \bar{\varepsilon}$ and $\theta = -\xi - \beta_{i_1 i_3}$. It follows that $u_{i_1}^2 = u_{i_3}^2 = -\xi$, $u_{i_2}^2 = \beta_i + \bar{\varepsilon} - \beta_{i_1 i_3} - \xi$ and $u_i^2 = \beta_i - \beta_{i_1 i_3} - \xi$ for $i \neq i_1, i_2, i_3$.
- reform 3 is such that $s = 1$, $\varepsilon_i = 0$ for $i \neq i_1, i_2$, $\tilde{\beta}_{i_1} = \tilde{\beta}_{i_2} = \beta_{i_1 i_2}$ and $\theta = -\xi - \beta_{i_1 i_2}$. It follows that $u_{i_1}^3 = u_{i_2}^3 = -\xi$ and $u_i^3 = \beta_i - \beta_{i_1 i_2} - \xi$ for $i \neq i_1, i_2$.

Note that the inequality $\bar{\varepsilon} \geq \frac{1}{2}(\beta_{i_3} - \beta_{i_1})$ guarantees that each reform \mathbf{u}^k belongs to S . Note also that each reform \mathbf{u}^k is rejected under sincere voting, i.e. $d^0(\mathbf{u}^k) = 0$. We thus compute for each agent the (weighted) expected benefit from accepting the three reforms. We obtain:

$$\sum_{k=1}^3 w^k u_{i_1}^k = 2(\beta_{i_2 i_3} - \beta_{i_1}) - 4\xi, \quad \sum_{k=1}^3 w^k u_{i_2}^k = \beta_{i_2} - \beta_{i_1 i_3} + \bar{\varepsilon} - 4\xi, \quad \sum_{k=1}^3 w^k u_{i_3}^k = \beta_{i_3} - \beta_{i_1 i_2} - 4\xi$$

and

$$\forall i \notin \{i_1, i_2, i_3\}, \quad \sum_{k=1}^3 w^k u_i^k = 2\beta_{i_2 i_3} - \beta_{i_1 i_3} - \beta_{i_1 i_2} - 4\xi = \beta_{i_2 i_3} - \beta_{i_1} - 4\xi.$$

For ξ small enough, all those sums are positive. Hence, (F, m) satisfies the Gains from Trade Condition. □

B Cooperation under incomplete information: an example

To illustrate how cooperation can emerge under incomplete information, we focus on the following simple example. We assume that $n = 2$ and $m = 1$, that is, two voters decide under the unanimity rule. There are two possible types of reforms, A and B , each type generating a fixed distribution of utilities as described in [Table 2](#).

		Voters	
		$u_{[1]}$	$u_{[2]}$
Reforms	A	-5	2
	B	-5	8

Table 2: Distribution of utilities under reforms A and B

Agents are equally likely to occupy any rank in the associated utility distribution. Moreover, both types of reforms are equally likely. Crucially, we assume here that agents observe their own utility but not that of their fellow committee member, and that the informational environment is common knowledge. The first best consists in only accepting reforms of type B , yielding an ex-ante utility $U^e = \frac{3}{4}$ for each agent. In the stage game, the (essentially) unique equilibrium in undominated strategies consists in voting sincerely, yielding an ex-ante utility $U^0 = 0$ as all reforms are then rejected. In the repeated game, absent any communication technology, agents cannot reach a payoff higher than $U^0 = 0$ at equilibrium.³³

Consider now an augmented stage game where voters can (simultaneously and publicly) report their own utility before voting. As without communication, sincere voting is a weakly dominant strategy for each voter in the voting step, so that $U^0 = 0$. To illustrate how repetition can generate welfare gains, we consider the following

³³Observe that the strategy profile consisting in accepting all reforms also yields an ex-ante utility of 0.

strategy profile, inspired from the linking mechanism of [Jackson and Sonnenschein \(2007\)](#).

Time is divided in consecutive *blocs* of four consecutive stages. At the each of first three stages of a given bloc, each voter reports her true utility u_i^t , unless $u_i^t = 8$ and she has already reported a utility of 8 in the current bloc. In the latter case, she reports 2. At the fourth stage, a voter with negative utility reports her utility truthfully, while a voter with positive utility reports a utility of 8 if she has not already reported such a utility in the current bloc, and a utility of 2 otherwise. In the voting step of any stage, voters vote sincerely unless at least one committee member reported 8 in that stage. In the latter case, both voters vote in favor of the reform. Finally, if a deviation is ever detected (either because a voter reports 8 twice within a bloc or because a voter does not comply with prescribed voting behavior), then all voters revert to the stage-game equilibrium (sincere voting) in all subsequent stages.

To analyse this strategy profile, it will be useful to introduce the following automaton. The set of states is $\Omega = \{0+, 1+, 1-, 2-\}$, where the number of a state ω indicates the number of voters who have already reported a utility of 8 in the current bloc, while the sign of a state ω indicates whether a given voter has not yet sent a utility of 8 in the current bloc (positive sign +), or has already sent it (negative sign -). From the point of view of any voter, a bloc starts at the initial state $0+$ and evolves along the stages of the bloc according to the automaton represented below.




Figure 7: Evolution of the state and instantaneous utilities for stages $t \in \{1, 2, 3\}$ and 4 in the candidate equilibrium profile.

On [Figure 7](#), we describe the transition probabilities as well as the instantaneous utilities (U_ω^t) incurred by the player at each state ω and stage t of the bloc. From

these data, we can compute the forward utilities $(V_\omega^t)_{\omega \in \Omega}^{t \geq 0}$ defined by,³⁴

$$\forall t \geq 0, \forall \omega \in \Omega, \quad V_\omega^t = (1 - \delta) \sum_{t'=t}^{+\infty} \delta^{t'} \sum_{\omega' \in \Omega} \mathbb{P}(\omega^{t'} = \omega' \mid \omega^t = \omega) U_{\omega'}^{t'}.$$

For the strategy profile under consideration to be an equilibrium, two sets of constraints need to be satisfied. First, voters must have incentives to report their utilities truthfully in the positive states $\Omega^+ = \{0+, 1+\}$, when their utility is positive, for all of the first three stages of a bloc. Denoting by $-\omega$ the function $0+ \mapsto 1-$ and $1+ \mapsto 2-$, these constraints can be written,

$$\forall t \in \{1, 2, 3\}, \forall \omega \in \Omega^+, \quad \begin{cases} (1 - \delta)8 + \delta V_{-\omega}^{t+1} \geq 0 + \delta V_\omega^{t+1} \\ (1 - \delta)2 + \delta V_{-\omega}^{t+1} \leq 0 + \delta V_\omega^{t+1}, \end{cases}$$

or equivalently,

$$\forall t \in \{1, 2, 3\}, \forall \omega \in \Omega^+, \quad 2 \leq \frac{\delta}{1 - \delta} (V_\omega^{t+1} - V_{-\omega}^{t+1}) \leq 8. \quad (27)$$

Second, voters must comply with prescribed voting strategies, in particular when they are asked to vote insincerely. These constraints can be written: $\forall t \in \{1, 2, 3\}, \forall \omega \in \Omega$, $(1 - \delta)(-5) + \delta V_\omega^{t+1} \geq 0 + \delta 0$ and $(1 - \delta)(-5) + \delta V_{0+}^1 \geq 0 + \delta 0$, or equivalently,

$$V_{0+}^1 \geq \frac{5(1 - \delta)}{\delta} \quad \text{and} \quad \forall t \in \{2, 3, 4\}, \forall \omega \in \Omega, \quad V_\omega^t \geq \frac{5(1 - \delta)}{\delta}. \quad (28)$$

Numerically, we find that constraints (27) and (28) are jointly satisfied when $\delta = 0.96$, yielding an ex-ante utility $U^* := V_{0+}^1 \approx 0.45 > U^0 = 0$. This example thus illustrates how, even under incomplete information, a cooperation norm can emerge to reach outcomes that Pareto-dominate (in expectation) the sincere voting benchmark.

³⁴Note that $\omega^{4t+1} = 0+$ for each t so that $V_{\omega^{4t+1}}^{4t+1} = V_{0+}^1$ for each t . It follows that all forward utilities can be computed in a few lines of code.