Policy Experimentation and Redistribution: Optimality of Non-Unanimity Rules*

Vincent Anesi†
T. Renee Bowen‡

November 16, 2017

PRELIMINARY AND INCOMPLETE

Abstract

We study conditions under which optimal policy experimentation can be implemented by a legislature. We consider a dynamic legislative bargaining game in which, each period, legislators choose to implement a risky reform or maintain a known policy. We first show that when no redistribution is allowed the unique equilibrium outcome is generically inefficient. When legislators are allowed to redistribute resources (even a small amount) and can amend the status quo sufficiently frequently, there always exists an equilibrium that supports optimal experimentation for any non-unanimity voting rule. We show that with unanimity, optimal policy experimentation is possible only with a sufficient amount of redistribution. In this sense, non-unanimity rules dominate unanimity rule.

*We thank Joel Watson, Laura Doval and participants in the Stanford Political Economy Theory seminar, the Summer Workshop in Political Economy at Stony Brook University and the UCSD theory mini conference, for helpful comments and suggestions.

†University of Nottingham

‡Stanford University

1 Introduction

One objective of a well-functioning government is to embark on a reform if the expected aggregate benefit exceeds the aggregate cost, and, conversely, quit an existing reform program when the likelihood that promised benefits will materialize becomes sufficiently small. Yet examples abound where reforms expected to yield net benefits are not undertaken (for example labor market reforms, or trade reforms), and numerous reform programs continue well beyond the point where they are generally accepted as "failed" (as in the case of import-substituting policies in many developing countries). In many democracies, the decision to begin or end a reform program is made by legislatures acting in the interest of their constituents. Thus, in spite of the promise of aggregate benefits, if a sufficient number of legislative districts do not directly gain, reforms may fail to be implemented. Similarly, if a sufficiently large coalition of vested interests support a failed reform, it may persist at the the expense of everyone else. To avoid such issues, it might be expected that the gains to winners should be redistributed in such a way that losers are fully compensated and the right level of reform experimentation occurs, but such redistribution does not occur in practice. In fact, as Acemoglu et al., 2015 points out, there can be a number of (exogenous) direct constraints on redistribution. Even if some redistribution is allowed, legislatures may still face another constraint: potential winners from a reform may not be able to *commit* to redistribute future benefits in order to compensate losers. This paper asks if (in the absence of commitment) socially efficient reform experimentation can be implemented by legislatures if benefits can be redistributed. If yes, is there some minimal level of redistribution necessary for this to happen and how do these answers depend on voting rules?

To answer these questions, we present a dynamic legislative bargaining model that combines policy experimentation and redistribution. A legislature meets each period to decide policy. Policy has three components – the choice to implement a risky reform or revert to a safe (known) policy, a tax rate, and the choice of how to distribute available resources. Available resources are determined by an exogenous constraint on redistribution. This exogenous constraint may be thought of as a tax rate that is chosen separately from the choice of reforms and redistribution.[1] Other possible interpretations are constitutional constraints, or the threat of capital flight (see Acemoglu et al., 2015). Policies are assumed to continue unless changed by the legislature, and, in this sense, exhibits an *endogenous status quo* feature.[2] Bargaining follows the protocol developed in Diermeier et al. (2017) .

The choice between a risky reform and a safe policy is modeled as a bandit problem in the spirit of Keller, Rady and Cripps (2005), but with heterogeneous payoffs across legislators. The safe policy generates a certain benefit if selected, and these benefits are heterogeneous across legislative districts. If the reform is good, then it generates benefits stochastically. These benefits are also heterogeneous across legislative districts. If the reform is not good, it never generates benefits. There is a prior belief that the reform is good that is common to all legislators. With each failed attempt at reform, all legislators update their belief about whether the reform is good or not according to Bayes' rule. Thus, with each failed reform attempt, the belief that the reform is good decreases, and the expected payoff from the reform also decreases. In the absence of legislative bargaining, a utilitarian social planner follows a stopping rule in which the reform is attempted until the belief that it is good reaches the optimal cutoff. We call this the *optimal stopping rule.*

We first find that when no redistribution is allowed, the legislature, in general, does not implement the optimal stopping rule. We show that with no redistribution the same outcome is implemented in any equilibrium. Either experimentation never occurs, or experimentation continues until *right pivot* wishes to stop. Experimentation only begins if the *left pivot* is in agreement.[3] Thus, if experimentation begins, it ends when the posterior belief that the reform is good is at the right pivot's ideal cutoff belief. Thus the optimal

[1] For example, U.S. tax legislation is primarily driven by the U.S. Treasury and the office of the President, whereas other policies related to reforms and redistribution may come from the House of Representatives or the Senate.

[2] This endogenous status quo feature of dynamic policymaking as been used in a number of recent papers and in a variety of policy settings. These include Kalandrakis (2004) (pure redistribution), Bowen, Eraslan and Chen (2014) (entitlement programs), Piguillem and Riboni (2015) (public spending), Dzuida and Loeper (2015) (unidimensional policy).

[3] We use terminology developed in Dzuida and Loeper (2017) applied to our setting. For any q-majority rule, the *right pivot* is legislator q, and the *left pivot* is legislator $n-q+1$.

stopping rule is only implemented if the right pivot's ideal coincides with the optimal cut-off. In some cases experimentation with reform may not begin even though all legislators wish to experiment to some degree.

In contrast, we show that when redistribution is allowed, the legislature may implement the optimal stopping rule depending on the voting rule, and how much redistribution is allowed. With a nonunanimity voting rule, socially efficient experimentation is attainable as long as the exogenous constraint on redistribution is relaxed (even marginally). With unanimity, a minimum level of redistribution must be permitted to sustain the optimal stopping rule. Non-unanimity voting rules thus permit socially efficient redistribution for a larger set of parameters than does unanimity, and thus non-unanimity dominates unanimity in the sense of Bouton, Llorente-Saguer and Malherbe (forthcoming).

Related literature This project is most closely related to the literature on collective experimentation and voting rules, including Strulovici (2010), and Messner and Polborn (2012). Like us, these papers study how various voting rules affect incentives to experiment in committees. In particular, Strulovici (2010) shows that efficient policy experimentation cannot be sustained with voting. This occurs as voters learn whether or not the policy will be beneficial to them. In contrast to Strulovici (2010), we assume that agents know their potential future benefit form experimentation, but there is a common uncertainty about whether the reform is good. In this sense, a conflict exists between voters prior to beginning experimentation. We show that this conflict can be mitigated with a sufficient level of redistribution. Strulovici (2010) does not consider redistribution. Other papers considering policy experimentation and politics include Majumdar and Mukand (2004), Volden et al. (2008), Cai and Treisman (2009), Callander (2011), and Callander and Hummel (2014). These papers do not consider dynamic legislative bargaining and policy experimentation.

We consider that policies, once implemented, can only be changed with a new round of voting, and hence what we do relates to the literature on bargaining with an endogenous status quo. This literature was pioneered by Baron (1996). A degenerate version of our model, in which policy is purely private, extends the divide-the-dollar framework studied in Kalandrakis (2004, 2010), Battaglini and Palfrey (2012), Bowen and Zahran (2012), Bowen and Baron (2014), Nunnari (2014), Richter (2014), and Anesi and Seidmann (2015).

This paper is related to the substantial body of political economy research studying political failures, which was first articulated by Besley and Coate (1998). More closely related are papers by Fernandez and Rodrik (1991) and Dziuda and Loeper (2016). Both these papers explore inefficient policy persistence, but do not consider how this might be affected by the ability of a legislature to distribute resources. Dewatripont and Roland (1992) examined gradualism in reforms, however, did not consider how this is affected by redistribution.[4]

The remainder of the paper is organized as follows. In Section 2 we present our model of dynamic policy making in a legislature. In Section 3 we analyze two important benchmarks - the optimal stopping rule and the equilibria in the case of no redistribution. In Section 4 we consider that redistribution is allowed and consider non-unanimity and unanimity rules separately. We conclude with a discussion of the results in Section 5.

2 Model

Players, policies and preferences. We present a stylized model of dynamic policy making by a legislature, which consists of $n \geq 3$ legislators: $N \equiv \{1, \ldots, n\}$. Time is divided into discrete periods of length $\Delta > 0$, with the legislature meeting again at the beginning of each period. We will subsequently focus on the limiting case as Δ becomes arbitrarily short.[5]

[4] Tornell (1998) also provides a theory of reform, but does not focus on the impact of legislative redistribution.

[5] This approach to "discretizing" dynamic games is commonly used in the experimentation literature, including for example Murto and Välimäki (2011) and Hörner and Samuelson (2013). It permits, in particular, to analyze how heterogeneous agents collectively trade off exploration and exploitation in experimentation, while avoiding the standard difficulties inherent in formulating continuous-time games with history-dependent behavior - for example, strategies may not be well-defined (e.g., Bergin and MacLeod, 1993). The main results of the paper hold with standard discrete time, however, the experimentation-

In each period t the legislature has to choose a policy p^t that has three components. The first component of the policy a^t is a choice to engage in a risky reform R or implement a known safe policy S. The reform R is either good or bad. There are two possible outcomes if it is implemented, success or failure. The probability of success is $\gamma\Delta$ if the reform is good and the probability of success is 0 if it is bad. Legislator i values a success at $r_i > 0$ and a failure at 0. If implemented, the safe alternative S gives legislator i a per-period benefit of $\Delta s_i > 0$ with probability one. Let $\bar{r} \equiv \sum_{i \in N} r_i$ and $\bar{s} \equiv \sum_{i \in N} s_i$, and assume that $\gamma\bar{r} > \bar{s} > 0$, so that the reform, if good, is better than the safe policy in expectation. The second component of p^t is a tax rate on individual benefits τ^t. We assume that there is an exogenous upper bound $\tau^{\max} \in [0, 1]$ on τ^t, so that $\tau^t \in [0, \tau^{\max}]$. This upper bound $\tau^{\max}$ represents an exogenous constraint on redistribution. The third component of p^t, denoted x^t, is a choice of how to redistribute the tax revenues raised in period t and hence $x^t \in X \equiv \big\{(x_1, \ldots, x_n) \in [0, 1]^n \colon \sum_{i \in N} x_i = 1\big\}$.

If policy $p^t = (a^t, \tau^t, x^t)$ is implemented at the start of period t and legislator i believes that the reform is good with probability α, then her (per-period) expected payoff is given by

$$w_i(a^t, \tau^t, x^t \mid \alpha) \equiv \begin{cases} \alpha\gamma\Delta\big[(1-\tau^t) r_i + \tau^t x_i^t \bar{r}\big] & \text{if } a^t = R\ , \\ \Delta\big[(1-\tau^t) s_i + \tau^t x_i^t \bar{s}\big] & \text{if } a^t = S\ . \end{cases}$$

Legislators discount at the continuously compounded rate ρ — so that the common discount factor is $\delta = e^{-\rho\Delta}$ — and seek to maximize their average discounted sums of payoffs.

Policy making. We model policy-making using a dynamic bargaining framework with an endogenous status quo. Each period t begins with a status quo policy p^{t-1}, inherited from the previous period. An order of proposers $(\pi_1, \ldots, \pi_n)$ is randomly selected from the set Π of all permutations of N, with each permutation in Π having a positive probability of being selected.[6] Proposer π_1 then makes the first proposal $p = (a, \tau, x) \in \{R, S\} \times [0, \tau^{\max}] \times X$; once the proposal is made legislators vote sequentially (in an arbitrary order) over whether

exploitation trade-off with the reform is not present with discrete time, implying that legislator's choices only depend on whether they benefit from the reform or not, and not the *degree* to which they benefit.

[6] We maintain this assumption throughout the text for greater clarity, but could obtain similar results using more general protocols. In particular, our results remain intact if, as do Diermeier et al. (forthcom-

to accept it.[7] If the proposal is accepted by q legislators, $q \in \left\{ \lceil \frac{n+1}{2} \rceil, \ldots, n \right\}$, it is implemented, payoffs accrue and the game transitions to the next period, where the new status quo is $p^t = p$. Otherwise, proposer π_2 is call upon to make a proposal and the same process is repeated. If the n proposers all make unsuccessful proposals, then the status quo p^{t-1} is implemented and remains the status quo in period $t+1$. Each of the proposal rounds takes a negligible amount of time.

The game begins with the exogenously given status quo $p^0 \equiv (a^0, \tau^0, x^0)$, where $a^0 \equiv S$, $\tau^0 \equiv 0$ and $x_i^0 \equiv s_i/\bar{s}$ for all $i \in N$. That is, the initial status quo consists of the safe alternative and no redistribution of individual benefits.

Learning. The initial probability that the risky reform R is good is given by $\alpha_0 \in (0,1)$. Legislators update their (common) belief about R's type through the sequence of policy choices using Bayes' rule. The first successful trial of R reveals to all legislators that it is good. In the event that $k \in \mathbb{N}$ trials are unsuccessful, the belief is

$$\alpha_k \equiv \frac{\alpha_0(1-\gamma\Delta)^k}{\alpha_0(1-\gamma\Delta)^k + 1 - \alpha_0} \ .$$

Let $A \equiv \left\{ \alpha_k \colon k = 0, 1, 2, \ldots \right\} \cup \{1\}$ be the set of possible values for the belief.

Equilibrium. As stated, our objective is to explore the institutional mechanisms that support socially efficient experimentation. To do so, we follow closely the approach taken by Acemoglu et al. (2008), studying conditions under which efficient experimentation can be sustained by renegotiation-proof (pure strategy) perfect Bayesian equilibria (PBEs). Note that, in a PBE of this game, legislators' beliefs are necessarily given by Bayes' rule. A PBE is said to be renegotiation-proof if, after any public history, there does not exist another PBE that can make all players weakly better off (and some strictly better off).[8]

ing), we allow Nature to select any finite list of players $(\pi_1, \ldots, \pi_m)$ (possibly with repetition) such that the members of this list form a blocking coalition. Our main results also remain intact with an individual random proposer as is common in the legislative bargaining literature with an endogenous status quo.

[7] Sequential voting is the standard approach with non-Markovian equilibrium concepts — e.g., Cho and Duggan (JET 2009). This ensures that agents always vote as if pivotal.

[8] Though the game has imperfect information, all players are symmetrically informed at every history and, therefore, the usual interpretation of renegotiation-proofness applies.

However, we will not impose any refinement of PBE when stating our negative results (those claiming that efficient policies cannot be sustained) in order to make them stronger. In order to limit the number of possible cases (without affecting the paper's conclusions), we will also assume that in case of a tie, a player will prefer to continue rather than to stop experimenting. Henceforth we will refer to a PBE with this tie-breaking rule as simply an *equilibrium*, and we will refer to a renegotiation-proof PBE with this tie-breaking rule as a *renegotiation-proof equilibrium.*

3 Benchmarks

To highlight the normative implications of redistribution on experimentation outcomes, we employ two benchmarks to which later results can be compared: in the first, we characterize the policy sequence set by a utilitarian social planner; in the second, we characterize equilibrium outcomes for the legislative bargaining game when there is no redistribution (i.e., $\tau^{\max} = 0$). Comparison of these benchmarks reveals that in the absence of redistribution, socially efficient experimentation typically cannot be sustained in equilibrium.

3.1 The Optimal Stopping Rule

Consider the problem of a social planner whose objective is to maximize aggregate payoffs. This is a standard Markov decision problem with the planner's belief as a state variable. When the belief is α_k and the planner implements the risky reform R she obtains, in addition to the expected aggregate revenue $\alpha_k \gamma \Delta \bar{r}$, some information that she uses to update her beliefs. When she implements the safe alternative, S, she only obtains the aggregate revenue $\bar{s}\Delta$ and her belief remains unchanged. Therefore, if the optimal solution requires that S be implemented in a given period t, then the belief will remain the same and S will also be implemented in all future periods. As is standard in the literature on experimentation, it follows that the optimal solution is a stopping rule: there exists a $k^* \in \mathbb{N}$ such that, after k^* unsuccessful trials of R, the belief is so low that the planner suspends experimentation and implements the safe alternative only. Formally, let $V^*(\alpha)$ be the planner's average discounted value from a period that begins with a belief α. That is, if the planner applies the optimal stopping rule, then $V^*(\alpha)$ is the expected sum of the legislators' average discounted payoffs from the resulting outcome path. As $\gamma\bar{r} > \bar{s}$, we evidently have $V^*(1) = \gamma\Delta\bar{r}$. If the social planner chooses to continue experimenting when the belief is α_k, then her expected payoff is equal to $\big[1-\delta(1-\gamma\Delta)\big]\alpha_k\gamma\Delta\bar{r} + \delta(1-\alpha_k\gamma\Delta)V^*(\alpha_{k+1})$. If she chooses to stop, then her payoff is $\bar{s}\Delta$. Hence,

$$V^*(\alpha_k) = \max\Big\{\big[1-\delta(1-\gamma\Delta)\big]\alpha_k\gamma\Delta\bar{r} + \delta(1-\alpha_k\gamma\Delta)V^*(\alpha_{k+1}), \bar{s}\Delta\Big\}.$$

Denote α^* as the optimal cutoff. Recall that periods are discrete and thus α^* must be an element of the set of feasible beliefs A. The optimal cutoff α^* may thus be strictly smaller than the belief that makes the social planner indifferent between continuing with the reform for one more period and switching to the safe alternative. Noting that if $\alpha_k = \alpha^*$, then $V^*(\alpha_{k+1}) = \bar{s}\Delta$, we obtain

$$\alpha^* \equiv \min\left\{\alpha_0, \max\left\{\alpha \in A \colon \alpha < \frac{(1-\delta)}{\gamma\big[(1-\delta(1-\gamma\Delta))(\bar{r}/\bar{s}) - \delta\Delta\big]}\right\}\right\}.$$

We will say that an equilibrium *sustains the optimal stopping rule* if the legislature applies the optimal stopping rule on the equilibrium path.

To make things interesting, we wish to study situations where social optimality dictates to experiment for at least one period, and thus $\alpha^* < \alpha_0$, for Δ very small. Note that, as $\Delta \to 0$, the social planner's ideal cutoff converges to

$$\min\left\{\alpha_0, \frac{\rho}{\gamma\big[(\rho+\gamma)(\bar{r}/\bar{s}) - 1\big]}\right\}.$$

Imposing $\alpha_0 > \rho/(\gamma\big[(\rho+\gamma)(\bar{r}/\bar{s}) - 1\big])$ guarantees that there is a $\widehat{\Delta} > 0$ such that, for all $\Delta < \widehat{\Delta}$, we have $\alpha^* < \alpha_0$. Throughout, we maintain this assumption.

Assumption A1. $\alpha_0 > \rho/(\gamma\big[(\rho+\gamma)(\bar{r}/\bar{s}) - 1\big])$.

3.2 Policy Experimentation without Redistribution

We now return to the analysis of the bargaining game introduced in Section 2. Throughout this section, we assume that no redistribution is permitted, i.e., $\tau^{\max} = 0$.

To begin we must establish some notation. By the same logic as above, each legislator i's ideal experimentation plan is a stopping rule with cutoff

$$\hat{\alpha}_i \equiv \min\left\{\alpha_0, \max\left\{\alpha \in A \colon \alpha < \frac{(1-\delta)}{\gamma\big[(1-\delta(1-\gamma\Delta))(r_i/s_i) - \delta\Delta\big]}\right\}\right\} .$$

This is the rule that legislator i would implement if she were the sole decision maker. Note that $\hat{\alpha}_i$ is (weakly) decreasing in the ratio r_i/s_i. That is, each legislator's incentive to experiment increases with the extent to which she values the reform over the safe alternative. Henceforth, without loss of generality, we order legislators such that $r_i/s_i \leq r_{i+1}/s_{i+1}$ and thus legislator 1 wishes to cease experimenting first. Moreover, we refer to $\hat{\alpha}_i$ as legislator i's *ideal cutoff,* and to r_i/s_i as her *benefit ratio.*

Next, let $\overline{V}_i(\alpha_k)$ be the dynamic payoff to legislator i at the start of the game induced by a stopping rule with cutoff $\alpha_k \in A \setminus \{1\}$. When the cutoff is α_k, in each of the first k periods, player i receives $(1-\delta)r_i$ if the risky alternative is good and successful. This occurs with probability $\alpha_0\gamma\Delta$. Starting from period $k+1$ onward there are two possibilities. Either R succeeded in at least one of the first k periods, and it is implemented from period $k+1$ on. This occurs with probability $\alpha_0[1-(1-\gamma\Delta)^k]$ and yields a per-period expected payoff of $(1-\delta)\gamma\Delta r_i$. Or, instead, each of the first k trials is unsuccessful, and the safe alternative S is implemented from period $k+1$ on. This occurs with the complementary probability $[1-\alpha_0\big(1-(1-\gamma\Delta)^k\big)]$ and yields a per-period payoff of $(1-\delta)\Delta s_i$. Thus legislator i's dynamic payoff for cutoff α_k is given by

$$\begin{aligned}\overline{V}_i(\alpha_k) \;&\equiv\; \big(1-\delta^k\big)\alpha_0\gamma\Delta r_i + \delta^k\Big[\alpha_0\big[1-(1-\gamma\Delta)^k\big]\gamma\Delta r_i + \big[1-\alpha_0\big(1-(1-\gamma\Delta)^k\big)\big]\Delta s_i\Big] \\ &=\; \big[1-\delta^k(1-\gamma\Delta)^k\big]\alpha_0\gamma\Delta r_i + \delta^k\big[1-\alpha_0\big(1-(1-\gamma\Delta)^k\big)\big]\Delta s_i \;.\end{aligned}$$

Our first result gives a complete characterization of the equilibria for the bargaining game in terms of the voting rule and the distribution of ideal cutoffs:

Proposition 1. *Suppose* $\tau^{\max} = 0$. *There is a unique PBE outcome that takes the form of a stopping rule with cutoff:*

$$\bar{\alpha} = \begin{cases} \hat{\alpha}_q & \text{if } \overline{U}_{n-q+1}\big(\hat{\alpha}_q\big) \geq s_{n-q+1}\Delta \;, \\ \alpha_0 & \text{otherwise.} \end{cases}$$

An immediate consequence of Proposition 1 is that if the voting rule is simple majority $q = (n+1)/2$, then the stopping rule with median legislator's ideal cutoff is implemented. This is reminiscent of the well-known median voter theory, where an odd number of voters with static, single-peaked preferences must collectively choose a policy from a unidimensional choice space. However, the logic behind Proposition 1 is more subtle, not only because this is a dynamic setting with evolving status quo and beliefs, but mainly because policy preferences in each period are *endogenously* determined by equilibrium behavior in future periods. To see this, we discuss the intuition for the proof below.

Although equilibrium continuation values in a PBE could intricately depend on the previous history of play, we first show that if the belief becomes smaller than or equal to legislator n's ideal cutoff, $\hat{\alpha}_n$, the safe alternative must be implemented in every period of every continuation game. Therefore, when the belief is α_k with $\alpha_{k+1} = \hat{\alpha}_n$, the members of the winning coalition $\{1, \ldots, q\}$ can play in accordance with their own preferences without risking to trigger adverse decisions in future periods: they will always agree to switch from R to S if R is the status quo, and will always reject any proposal to change S to R if S is the status quo.

Applying the same logic recursively, we obtain that whenever the belief α_k is smaller than $\hat{\alpha}_q$, the equilibrium outcome of every continuation game is unique and has to be a stopping rule with cutoff α_k. Now consider any belief $\alpha_k \geq \hat{\alpha}_q$. All agents that wish to experiment more prefer to do so right away rather than wait to experiment because by waiting, expected future benefits are further discounted, thus turning the dynamic bargaining problem into the choice between two options: implementing the stopping rule with cutoff $\hat{\alpha}_q$, or maintaining the safe alternative. If $\overline{U}_{n-q+1}(\hat{\alpha}_q) \geq s_{n-q+1}\Delta = \overline{U}_{n-q+1}(\alpha_0)$, then legislator $n-q+1$ and, by single-peakedness, all the other members of the winning coalition $\{n-q+1, \ldots, n\}$ prefer the first option. The stopping rule with cutoff $\hat{\alpha}_q$ must therefore be the unique equilibrium outcome. If instead $\overline{U}_{n-q+1}(\hat{\alpha}_q) < s_{n-q+1}\Delta$, then legislator $n-q+1$ and all the other members of the blocking coalition $\{1, \ldots, n-q+1\}$ prefer to maintain the initial status quo S and, consequently, experimentation never occurs — i.e., the unique equilibrium cutoff is α_0.

An immediate consequence of Proposition 1 is that, with simple majority voting, the legislature always over-experiments if $\hat{\alpha}_q < \alpha^*$, and always under-experiments if $\hat{\alpha}_q > \alpha^*$. With any super-majoritarian voting rule, it over-experiments if $\hat{\alpha}_q < \alpha^*$ and $\overline{U}_{n-q+1}(\hat{\alpha}_q) \geq s_{n-q+1}\Delta$, and it under-experiments either if $\hat{\alpha}_q > \alpha^*$ and $\overline{U}_{n-q+1}(\hat{\alpha}_q) \geq s_{n-q+1}\Delta$, or if $\overline{U}_{n-q+1}(\hat{\alpha}_q) < s_{n-q+1}\Delta$. Since $\lim_{\Delta \to 0} \hat{\alpha}_q \neq \alpha^*$ whenever $r_q/s_q \neq \bar{r}/\bar{s}$, this implies the following:

Corollary 1. *Suppose $\tau^{\max} = 0$. If $r_q/s_q \neq \bar{r}/\bar{s}$, then there exists $\Delta_0 > 0$ such that, whenever $\Delta < \Delta_0$, all equilibria fail to sustain the optimal stopping rule.*

Corollary 1 states that, in the absence of redistribution, every equilibrium fails to sustain the optimal stopping rule (in the limit as $\Delta \to 0$) whenever legislator q's benefit ratio differs from $\bar{r}/\bar{s}$, which is generically the case. We conclude that efficient experimentation is typically impossible without redistribution.

We conclude this section with a remark on Pareto inefficiency. Proposition 1 and the subsequent corollary show how under-experimentation may happen under any voting rule, yielding socially inefficient outcomes in equilibrium (in a Utilitarian sense). But such equilibrium outcomes may even be Pareto dominated. This is illustrated by the following example. Suppose $q > (n+1)/2$ (so that $n - q + 1 < q$), $\rho = \gamma = \alpha_0 = 2/3$, $r_i = 2$ for all i, $s_i = 1$ for all $i < q$, and $s_i = \epsilon$ for all $i \geq q$, where $\epsilon > 0$ is arbitrarily small. It is readily checked that, under these assumptions, $\lim_{\Delta \to 0} \hat{\alpha}_i = 3/5 < \alpha_0$ for all $i < q$, and $\lim_{\Delta \to 0} \hat{\alpha}_i = 3\epsilon/(8 - 3\epsilon)$ for all $i \geq q$. Hence, for arbitrarily small $\Delta > 0$, legislator q's optimal stopping rule converges to perpetual experimentation, so that $\overline{U}_{n-q+1}(\hat{\alpha}_q) \to \alpha_0 \gamma \Delta r_{n-q+1} = 8\Delta/9 < \Delta = s_{n-q+1}\Delta$, as $\epsilon \to 0$. It therefore follows from Proposition 1 that, for sufficiently small ϵ (and sufficiently small $\Delta > 0$), the reform is never implemented in equilibrium, although all legislators would be better off experimenting for a positive number of periods. The reason this happens is that, while legislators in the blocking coalition $\{1, \ldots, n - q + 1\}$ would like to experiment (since $\hat{\alpha}_i < \alpha_0$ for all $i < q$), because of the endogeneity of the status quo, they fear that experimentation will go on for too long, and so they prefer not to begin experimenting at all.

4 Policy Experimentation and Redistribution

We saw in the previous section that the optimal stopping rule is typically not sustainable in equilibrium without redistribution. In this section, we ask whether this is still true when it is possible to redistribute the revenues from experimentation among legislators and, if the answer is negative, how much redistribution is needed to attain efficient experimentation. The answer critically turns on the voting rule. We first consider non-unanimity voting rules and then unanimity rule.

4.1 Non-unanimity Voting Rules

In stark contrast to the case of no redistribution, our next proposition states that with non-unanimity voting rules the optimal stopping rule can be sustained by a renegotiation-proof equilibrium for *any* level of redistribution. In fact, something stronger is true: although stationary Markov strategies sharply constrain the ability to punish and reward legislators for past behavior, the renegotiation-proof equilibrium that sustains the optimal stopping rule can be taken to be stationary Markov.[9] In addition, if there are no limits to redistribution, then all renegotiation-proof equilibrium outcomes are arbitrarily close to the optimal stopping rule in terms of the payoff vectors they generate. We conclude that, without bounds on redistribution, collective experimentation must yield efficient, or nearly efficient outcomes. To state this formally, let $U_i(\sigma)$ denote legislator i's average discounted payoff at the beginning of the game under any strategy profile σ.

Proposition 2. *Suppose $q < n$. Then:*
(i) For every upper bound $\tau^{\max} > 0$, there exists $\overline{\Delta} > 0$ such that, for all $\Delta < \overline{\Delta}$, the optimal stopping rule is sustained by a (stationary Markov) renegotiation-proof equilibrium; and
(ii) if $\tau^{\max} = 1$ then, for all $\varepsilon > 0$, there exists $\widetilde{\Delta} > 0$ such that the following holds for all $\Delta < \widetilde{\Delta}$: $\sum_{i=1}^{n} U_i(\sigma) > V^*(\alpha_0) - \varepsilon$ for every renegotiation-proof equilibrium σ.

[9] Of course, renegotiation-proofness already constitutes an obstacle to the construction of efficient equilibria, as it may reduce the severity of the off-path "punishments" available to support the appropriate incentives (which must themselves be renegotiation-proof).

The first part of Proposition 2 shows that if the voting rule is less than unanimity, then efficient experimentation can be supported within a legislature without resorting to punishment strategies that condition on individual behavior. The proof is constructive; we exploit the approach of Anesi and Seidmann (2015) and Baron and Bowen (2015) but push this further to obtain an efficient Markovian equilibrium in a non-stationary environment where the legislators' policy preferences evolve with the endogenous belief. The idea of the construction is simple: it is built around a collection of $(n-1)$-member coalitions — the potential "*governing coalitions*." These governing coalitions have the property that each legislator belongs to at least one of these coalitions but not to all of them. The same inevitably happens from any history (both on and off the path): the optimal stopping rule is implemented and, in every period, the members of a given governing coalition equally share the sum of the expected aggregate revenues that can be redistributed, i.e., $\tau^{\max}\alpha_k\gamma\Delta\bar{r} > 0$ if the belief α_k exceeds α^*, and $\tau^{\max}\bar{s}\Delta > 0$ otherwise. At the start of every period, the status quo policy and the current belief — which are payoff relevant — reveal to the legislature whether play in the previous period was consistent with the optimal stopping rule, and whether a governing coalition formed (i.e., equally shared the entire transferable benefits among its members). If this is the case, then the same governing coalition forms again and continues to implement the optimal stopping rule; otherwise, the (randomly selected) first proposer successfully offers to form a ruling coalition and to follow the optimal stopping rule. Given the inevitability of this process, the best possible scenario for any legislator is to form or be a member of the governing coalition that will share the transferable revenues from experimentation in every future period. For any member i of such a coalition, the potential benefits of a one-period deviation vanish as the period length Δ becomes arbitrarily small, whereas the long-run cost does not: a deviation would trigger the formation of a new governing coalition in the next period, which legislator i might not be a member of. Though this would have no impact on the proportion $(1-\tau^{\max})$ of her future payoffs that cannot be redistributed (since the optimal stopping rule is implemented in any case), she would potentially lose a proportion $\tau^{\max} > 0$ of her share of future aggregate revenues as a member of the governing coalition. As no member of a governing coalition is prepared to run such a risk and governing coalitions are winning ($q < n$), profitable deviations from the prescribed path are impossible. Moreover, this equilibrium is renegotiation-proof since it generates payoff vectors in the Pareto frontier both on and off the path.

To see the main idea behind the second part of Proposition 2, observe that the intuition above remains intact if we use any sharing rule that gives every legislator a higher share of aggregate revenues when she is a member of the governing coalition than when she is not. Varying sharing rules thus allows us to slide along the relative interior of the Pareto frontier so as to generate efficient equilibria that Pareto dominate those equilibria that are too far away from the Pareto frontier. By the renegotiation-proofness criterion, the latter equilibria can be eliminated.

4.2 Unanimity Voting Rule

Given the result obtained for non-unanimity voting rules in the previous subsection, it is natural to ask whether the optimal stopping rule is also sustainable with any level of redistribution under unanimity rule. A little reflection suggests the answer is no. Indeed, efficient experimentation notably requires two sets of incentive constraints to be met. The first set ensures that the legislators unanimously agree to change the initial status quo policy to some (R, τ, x), $x \in X$ and $\tau \in [0, \tau^{\max}]$; the second ensures that if the belief becomes equal to α^*, then they unanimously agree to stop experimenting and implement some policy (S, τ', y), $y \in X$ and $\tau' \in [0, \tau^{\max}]$. Formally, the first constraint requires that each legislator i's equilibrium continuation value from implementing (R, τ, x) in the first period is greater than or equal to her payoff from maintaining the initial status quo in all future periods. In the benchmark case where $\tau^{\max} = 0$, this is equivalent to $\overline{U}_i(\alpha^*) \geq \Delta s_i$ for all $i \in N$ (where $\overline{U}_i(\cdot)$ is defined as in Section 4). Rearranging terms and taking the limit as Δ goes to zero, an application of l'Hôpital's rule gives

$$\left[1 - e^{\psi(1+\rho\gamma)}\right]\alpha_0\gamma r_i + e^{\psi\rho\gamma}(1 - \alpha_0 + e^{\psi}\alpha_0)s_i \geq s_i \ ,$$

where $\psi \equiv \log\left[\frac{\rho\bar{s}}{(\rho+\gamma)(\gamma\bar{r}-\bar{s})}\right] < 0$. Therefore, if legislator 1's benefit ratio satisfies

$$\frac{r_1}{s_1} < \frac{1 - e^{\psi\rho\gamma}(1-\alpha_0+e^{\psi}\alpha_0)}{\left[1-e^{\psi(1+\rho\gamma)}\right]\alpha_0\gamma} \ , \tag{1}$$

then there exists $T > 0$ such that, for every $\tau^{\max} \in [0, T)$, her incentive constraint is always violated for arbitrarily small Δ's. Intuitively, when $\tau^{\max} < T$, the permitted level of redistribution is not sufficiently large to compensate legislator 1's loss from experimenting and, consequently, the optimal stopping rule is not sustainable in equilibrium.

The second key incentive constraint in the construction of efficient equilibria formally requires that if the belief becomes equal to α^*, then each legislator i's payoff from implementing policy (S, τ', y), $\left[(1-\tau')s_i + \tau' y_i \bar{s}\right]\Delta$, exceeds her equilibrium continuation value from rejecting it.[10] For every status quo policy (R, τ'', z), the latter value is bounded below by $\alpha^*\gamma\Delta\left[(1-\tau'')r_i + \tau'' z_i \bar{r}\right]$. Hence, an obvious necessary condition for the existence of an equilibrium that supports the optimal stopping rule is that $(1-\tau')s_i + \tau' y_i \bar{s} \geq \alpha^*\gamma\left[(1-\tau'')r_i + \tau'' z_i \bar{r}\right]$ for some $(\tau', y), (\tau'', z) \in [0, \tau^{\max}] \times X$ and all $i \in N$. Setting $\tau^{\max} = 0$ and letting Δ go to zero, we can rewrite it as

$$s_i \geq \frac{\rho r_i \bar{s}}{(\rho+\gamma)\bar{r} - \bar{s}} \ ,$$

for all $i \in N$. Thus, if legislator n's benefit ratio is large relative to the social planner's, i.e., if

$$(\rho+\gamma)\frac{\bar{r}}{\bar{s}} < 1 + \rho\frac{r_n}{s_n} \ , \tag{2}$$

then there exists $T > 0$ such that, for every $\tau^{\max} \in [0, T)$ and arbitrarily small Δ, legislator n's incentive constraint cannot hold. Because of the legislature's imperfect ability to redistribute the benefits from ending experimentation towards legislator n, the latter must reject any proposal to stop experimenting when the belief is α^* in equilibrium.

The discussion above is summarized in the following proposition.

[10] As the payoff vector $\Delta\big((1-\tau')s_i + \tau_i y_i \bar{s}\big)_{i\in N}$ belongs to the Pareto frontier when the belief is α^* (and each legislator can reject all future proposals to amend (S, τ', y) once it has been implemented), $\left[(1-\tau')s_i + \tau' y_i \bar{s}\right]\Delta$ must be each legislator i's equilibrium continuation value from implementing (S, τ', y).

Proposition 3. *Suppose $q = n$. If either condition (1) or (2) is satisfied, then there exists $T > 0$ such that the following holds for all $\tau^{\max} \in [0, T)$: there is $\widehat{\Delta} > 0$ such that, whenever $\Delta < \widehat{\Delta}$, every equilibrium fails to sustain the optimal stopping rule.*

This proposition shows that institutional details matter: in contrast to any quota rule short of unanimity, *efficient experimentation may not be attainable under unanimity rule if not enough redistribution is permitted.* Under the premises of the proposition, even unrefined equilibria all fail to support the optimal stopping rule.

Explicit analytical results characterizing the exact minimum value of $\tau^{\max}$ needed to avoid this negative conclusion for every parametric configuration of the model are hard to come by: a general characterization of the set of (renegotiation-proof) equilibria, though possible,[11] would provide relatively little if any analytical purchase on the problem at hand. Instead, we establish that sufficiently high levels of redistribution permit efficient experimentation under unanimity rule, irrespective of the details of the underlying policy preferences and information structure.

To gain some intuition on how redistribution allows the legislature to appropriately adjust benefit ratios, consider again the two key incentive constraints discussed above: One must first ensure that all players (including those who prefer the safe alternative to a good reform) optimally agree to implement the risky reform R at the start of the game, and second that they all agree to revert to the safe alternative if the belief attains the optimal cutoff α^*. If the legislators can credibly commit in equilibrium to redistribute revenues in such a way that their benefit ratios all coincide with the social planner's, $\bar{r}/\bar{s}$, then the second requirement will be met. More precisely, suppose a legislator proposes to change the status quo (R, τ, x) to some policy (S, τ', y) when the belief becomes equal to

[11] It is possible to show that, in continuation games with arbitrarily small beliefs, the set of renegotiation-proof equilibrium payoff vectors (for small enough Δ) is the simplex $\{(w_1, \ldots, w_n) \in \prod_{i=1}^{n} [(1 - \tau^{\max})\Delta s_i, \bar{s}\Delta] : \sum_{i=1}^{n} w_i = \bar{s}\Delta\}$ if the initial status quo is of the form (R, τ, x), and the singleton $\{[(1-\tau)s_i + \tau x_i \bar{s}]\Delta\}_{i \in N}$ if the status quo is of the form (S, τ, x). A backward-induction argument then gives the sets of equilibrium payoff vectors for higher beliefs.

α^*. Then, as long as

$$\frac{(1-\tau)r_i + \tau x_i \bar{r}}{(1-\tau')s_i + \tau' y_i \bar{s}} = \frac{\bar{r}}{\bar{s}} ,$$

all legislators are better off accepting the proposal. This condition is trivially satisfied if $\tau = \tau' = 1$ and $x = y$. The requirement that all legislators agree to change the initial status quo $(S, 0, x^0)$ to $(R, \tau, x) = (R, 1, x)$ in the first period can then be written as

$$\Big[1 - \delta^{k^*-1}(1-\gamma\Delta)^{k^*-1}\Big]\alpha_0 \gamma x_i \bar{r} + \delta^{k^*-1}\big[1 - \alpha_0 + (1-\gamma\Delta)^{k^*-1}\alpha_0\big] x_i \bar{s} \geq x_i \bar{s}$$

or, equivalently, $V^*(\alpha^*) \geq \bar{s}\Delta$, which holds for sufficiently small Δ by assumption. Thus, full redistribution ensures that the two key incentive constraints are satisfied if legislators can commit to use the appropriate redistributive policies in equilibrium. Though legislators face an infinite number of additional incentive constraints, full redistribution allows great flexibility in creating "rewards" and "punishments" that can support the incentives to implement such policies.

Proposition 4. *Suppose $q = n$ and $\tau^{\max} = 1$. Then, there exists $\widehat{\Delta} > 0$ such that, for all $\Delta < \widehat{\Delta}$, renegotiation-proof equilibria exist and all of them sustain the optimal stopping rule.*

5 Discussion

We analyze a model of policy making in which the benefits of reform may or may not be redistributed. We show that, except for nongeneric cases, socially efficient experimentation necessitates redistribution. That is, when no redistribution is possible, the optimal stopping rule is achieved if and only if the median gain from reform is equal to the average gain from reform. We show that arbitrarily small amounts of redistribution suffice to support socially efficient experimentation under non-unanimity voting rules. We further show that with unanimity, the optimal stopping rule can be sustained with a sufficient amount of redistribution. Since non-unanimity supports efficient experimentation for a larger set of parameters than unanimity, we say that non-unanimity rules dominate unanimity.

In future work we ask how the "progressiveness" of redistribution may affect the ability to sustain the optimal stopping rule in experimentation. That is, we consider what levels of taxation and redistribution will be endogenously chosen by legislatures, for some exogenously specified division of resources.

The redistributive tool we have used to support efficient policy experimentation in the case of non-unanimity voting rules can be used in other settings. In the case of non-unanimity we have supported efficient experimentation by linking the equilibrium of the experimentation game to an equilibrium of the redistribution game with a structure of dynamic coalitions (as in Anesi and Seidman (2015) and Baron and Bowen (2015)). Such "linking" of policies can be used in other settings to support efficient policymaking. We also explore this in future work.

Appendix

A Proof of Proposition 1

We assume throughout this section that $\tau^{\max} = 0$, so that distributions have no impact on players' payoffs. To lighten the notation, we simply represent policies as elements of $\{R, S\}$, omitting the irrelevant sharing-rule component. To prove Proposition 1, it is useful to consider a class of games $\big\{\Gamma(a, \alpha) \colon a \in \{R, S\} \;\&\; \alpha \in A\big\}$, where $\Gamma(a, \alpha)$ is the same game as that described in Section 2, except that it begins with an initial status quo alternative a (possibly equal to R) and a probability α (possibly different from α_0) that Alternative R is good.

It is useful to define the set of losers from the reform as $L \equiv \{i \in N \colon \gamma r_i < s_i\}$ and the set of winners as $W \equiv \{i \in N : \gamma r_i \geq s_i\}$.

Lemma A1. *Suppose $\tau^{\max} = 0$. Then,*
(i) $\Gamma(R, 1)$ has a unique equilibrium outcome: Alternative R is implemented in every period if $\big|L\big| < q$, and alternative S is implemented in every period otherwise;
(ii) for all $a \in \{R, S\}$ and $\alpha_k \leq \hat{\alpha}_n$, $\Gamma(a, \alpha_k)$ has a unique equilibrium outcome: Alternative S is implemented in every period.

Proof. *(i)* To see that there exists an equilibrium with the proposed outcome, consider the following (stationary Markov) strategy profile:

- Whenever the status quo is R, all proposers pass (i.e., propose R), and each voter i accepts proposal S if and only if $i \in L$;

- whenever the status quo is S, each proposer i proposes R if $i \notin L$ and passes otherwise, and each voter i accepts proposal R if and only if $i \notin L$.

It is easy to check that this strategy profile constitutes an equilibrium. (In particular, proposers who prefer S to R do not deviate and propose to amend status quo R because they anticipate that such a proposal would be rejected.)

Next we show that this is the unique equilibrium outcome. Our proof shares some of the intuitions of Shaked and Sutton's (1984) proof of equilibrium uniqueness in the Rubinstein model. Let the set of PBEs of $\Gamma(R,1)$ be denoted by $\mathcal{E}(R,1)$. In $\Gamma(R,1)$, legislator i's expected payoff in every period t is a convex combination of $\gamma\Delta r_i$ and Δs_i. Therefore, for every strategy profile σ her average discounted payoff is of the form $V_i(\sigma) \equiv \beta(\sigma)\gamma\Delta r_i + \big[1-\beta(\sigma)\big]\Delta s_i$, with $\beta(\sigma) \in [0,1]$. This implies that, for any two strategy profiles σ and σ', and any legislator $i \in \{i \in N : \gamma r_i > s_i\}$, we have $V_i(\sigma) \geq V_i(\sigma')$ if and only if $\beta(\sigma) \geq \beta(\sigma')$.

Suppose first that $\big|L\big| < q$. Let $\{\sigma^m\}$ be a sequence in $\mathcal{E}(R,1)$ that satisfies $\lim_{m\to\infty} \beta(\sigma^m) = \inf_{\sigma\in\mathcal{E}(R,1)} \beta(\sigma)$, so that $\lim_{m\to\infty} V_i(\sigma^m) = \inf_{\sigma\in\mathcal{E}(R,1)} V_i(\sigma)$ for all $i \in W \equiv \{i \in N : \gamma r_i \geq s_i\}$. Fix $m \in \mathbb{N}$. Every proposal that may successfully be made by the last proposer in the first period under σ^m (both on and off the path) must be accepted by some decisive player i in W. That is, i's continuation payoff from accepting the proposal, say U_i^a, must be at least as large as her payoff from rejecting it; i.e., $U_i^a \geq (1-\delta)\gamma\Delta r_i + \delta V_i(\sigma^r)$, where $\sigma^r \in \mathcal{E}(R,1)$ is the equilibrium of $\Gamma(R,1)$ that is played from the next period on if i rejects the proposal in the first period. From the argument in the previous paragraph, we thus have $U_j^a \geq (1-\delta)\gamma\Delta r_j + \delta V_j(\sigma^r)$ for all $j \in W$. Similarly, every proposal that may successfully be made by the penultimate proposer in the first period under σ^m (both on and off the path) must also be accepted by some member i of W. Her payoff (and therefore the payoff of all members of W) from accepting must be at least as large as the payoff from rejecting which, as previously shown, must be at least $(1-\delta)\gamma\Delta r_i + \delta \inf \big\{V_i(\sigma)\colon \sigma \in \mathcal{E}(R,1)\big\}$. Applying the same argument recursively, we obtain that the acceptance of any proposal in the first period must give a payoff of at least $(1-\delta)\gamma\Delta r_i + \delta \inf \big\{V_i(\sigma)\colon \sigma \in \mathcal{E}(R,1)\big\}$ for all $i \in W$. Hence,

$$V_i(\sigma^m) \geq (1-\delta)\gamma\Delta r_i + \delta \inf \big\{V_i(\sigma)\colon \sigma \in \mathcal{E}(R,1)\big\} ,$$

for all $i \in W$. Taking the limit as $m \to \infty$ and recalling the definition of $\{\sigma^m\}$, we obtain $\gamma\Delta r_i = \inf \big\{V_i(\sigma)\colon \sigma \in \mathcal{E}(R,1)\big\}$ (since $\gamma\Delta r_i$ is maximum feasible payoff for a player $i \in W$ when R is good with probability one). This in turn implies that R must be implemented with probability one in every period of every equilibrium of $\Gamma(R,1)$.

The argument for the case where $\left|L\right| \geq q$ is analogous.

(ii) To prove the second part of the lemma, we proceed in three steps: first, we show that the infimum of every player i's equilibrium payoff in $\Gamma(S, \alpha_k)$ converges to Δs_i as $k \to \infty$; then, we show that for sufficiently large k, alternative S is implemented in every period of $\Gamma(S, \alpha_k)$; finally, we use the previous result to complete the proof of the lemma.

Let $\alpha_k \in A \setminus \{1\}$; and let $\mathcal{E}(S, \alpha_k)$ be the set of PBEs of $\Gamma(S, \alpha_k)$. Every period of $\Gamma(S, \alpha_k)$ begins with a belief α that alternative R is good; then, either S is implemented, in which case legislator i receives a payoff of Δs_i; or R is implemented, in which case i's expected payoff is $\alpha\gamma\Delta r_i$. Therefore, every strategy profile σ yields an expected payoff of the form

$$V_i^k(\sigma) \equiv \Delta\left[\beta_s^k(\sigma)s_i + \beta_1^k(\sigma)\gamma r_i + \sum_{\ell=k}^{\infty} \beta_\ell^k(\sigma)\alpha_\ell\gamma r_i\right]$$

to player i in $\Gamma(S, \alpha_k)$, where $\beta_s^k(\sigma) + \beta_1^k(\sigma) + \sum_{\ell=k}^{\infty} \beta_\ell^k(\sigma) = 1$ and $\beta_s^k(\sigma), \beta_1^k(\sigma), \beta_\ell^k(\sigma) \in [0, 1]$ for each $\ell = k, k{+}1, \ldots$. Moreover, as R must have been successfully tried at least once to be known to be good, $\beta_1^k(\cdot)$ is bounded above by $\alpha_k\gamma$. Coupled with the fact that $\alpha_\ell \leq \alpha_k$ for all $\ell \geq k$, this implies that $\lim_{k\to\infty} M_i^k \equiv \sup_{\sigma\in\mathcal{E}(S,\alpha_k)} \left[\beta_1^k(\sigma) + \sum_{\ell=k}^{\infty} \beta_\ell^k(\sigma)\alpha_\ell\right]\gamma r_i$ for each $i \in N$. This in turn implies that there is a null sequence $\{\varepsilon^k\}$ in $\mathbb{R}_+$ such that, for all $\sigma \in \mathcal{E}(S, \alpha_k)$, we have

$$\max_{i\in N} \left|V_i^k(\sigma) - \beta_s^k(\sigma)s_i\Delta\right| \leq \Delta \max_{i\in N} M_i^k < \varepsilon^k \ ,$$

for every $k \in \mathbb{N}$. Now for each $k \in \mathbb{N}$, let $\{\sigma^{k,m}\}$ be a sequence in $\mathcal{E}(S, \alpha_k)$ such that $\lim_{m\to\infty} \beta_s^k(\sigma^{k,m}) = \inf_{\sigma\in\mathcal{E}(S,\alpha_k)} \beta_s^k(\sigma)$. As $\sigma^{k,m}$ is an equilibrium of $\Gamma(S, \alpha_k)$, there must be at least one legislator, say i_k, such that

$$V_{i_k}^k(\sigma^{k,m}) \geq (1-\delta)\Delta s_{i_k} + \delta \inf_{\sigma\in\mathcal{E}(S,\alpha_k)} V_{i_k}^k(\sigma) \ ;$$

otherwise some player would have a profitable deviation in the first period of $\Gamma(S, \alpha_k)$. It follows that

$$\beta_s^k(\sigma^{k,m})\Delta s_{i_k} + \varepsilon^k \geq (1-\delta)\Delta s_{i_k} + \delta\left[\inf_{\sigma\in\mathcal{E}(S,\alpha_k)} \beta_s^k(\sigma)\Delta s_{i_k} - \varepsilon^k\right] .$$

Taking the limit as $m \to \infty$, we obtain $\inf_{\sigma\in\mathcal{E}(S,\alpha_k)} \beta_s^k(\sigma) \geq 1 - 2(\varepsilon^k/\Delta s_{i_k})$. This implies that $\inf_{\sigma\in\mathcal{E}(S,\alpha_k)} \beta_s^k(\sigma)$ converges to one as $k \to \infty$ and, therefore, that there exists a null sequence $\{\eta^k\}$ such that $\lim_{k\to\infty} \max_{i\in N} \left| \inf_{\sigma\in\mathcal{E}(S,\alpha_k)} V_i^k(\sigma) - \Delta s_i \right| < \eta^k$, for all $k \in \mathbb{N}$, thus completing the first step of the argument.

We now turn to the second step of the proof. Observe first that, as $(1-\delta)(s_i - \alpha_k\gamma r_i)\Delta - \delta\eta^k$ converges to $(1-\delta)\Delta s_i > 0$ as $k \to \infty$, there is a sufficiently large $K \in \mathbb{N}$ such that $(1-\delta)(s_i - \alpha_k\gamma r_i)\Delta - \delta\eta^k > 0$, for all $k \geq K$. Let $k \geq K$, and suppose that $\Gamma(S,\alpha_k)$ has an equilibrium in which alternative R is implemented with positive probability. Consider the first period of $\Gamma(S,\alpha_k)$ in which R may be implemented. Every decisive voter i's benefit from rejecting any proposal to change S to R is bounded below by $(1-\delta)(s_i-\alpha_k\gamma r_i)\Delta+\delta\big[(\Delta s_i-\eta^k)-\Delta s_i\big] > 0$, where the bracketed term represents the difference between the lower and upper bounds on i's continuation payoffs from rejecting R and accepting it, respectively. (Recall that each legislator i's maximum payoff is Δs_i when the belief is smaller than or equal to $\hat{\alpha}_n$.) Hence, then every proposal to amend S to R is rejected in any equilibrium of $\Gamma(S,\alpha_k)$. We thus have $V_i^k(\sigma) = \Delta s_i$, for all $i \in N$ and all $\sigma \in \mathcal{E}(S,\alpha_k)$.

If $\alpha_K > \hat{\alpha}_n$, then Lemma 1(ii) follows immediately from the previous paragraph; so suppose that $\alpha_K \leq \hat{\alpha}_n$. To complete the proof of the result, consider the first period of $\Gamma(S,\alpha_K)$. If alternative R is implemented, then the expected payoff to each legislator i is $\big[1-\delta(1-\gamma\Delta)\big]\alpha_K\gamma\Delta r_i+\delta(1-\alpha_K\gamma\Delta)\Delta s_i < \Delta s_i$, where the inequality follows from $\alpha_K \leq \hat{\alpha}_n$ and the definition of the legislators' optimal cutoffs in Subsection 3.2; if alternative S is instead implemented, then her expected payoff will be a convex combination of $\big[1-\delta(1-\gamma\Delta)\big]\alpha_K\gamma\Delta r_i+\delta(1-\alpha_K\gamma\Delta)\Delta s_i$ (if R is implemented with positive probability in a future period) and Δs_i, with a positive coefficient on the latter. Therefore, all legislators are strictly better off implementing R: they all reject proposals to amend S to R (when decisive). Hence, $V_i^k(\sigma) = \Delta s_i$, for all $i \in N$, $k \geq K$ and $\sigma \in \mathcal{E}(S,\alpha_k)$. Applying the same argument recursively from belief α_{K-1} to belief $\hat{\alpha}_n$ we obtain that, for all $\alpha_k \leq \hat{\alpha}_n$, $\Gamma(S,\alpha_k)$ has a unique equilibrium outcome: Alternative S is implemented in every period. By the same logic, the same is also true in game $\Gamma(S,\alpha_k)$, $\alpha_k \leq \hat{\alpha}_n$. In such a game, every decisive voter receives her largest possible payoff Δs_i if she accepts a proposal to change the status quo R to S, since the latter will then never be amended. Any such a proposal must therefore be successful and, as S is the ideal policy of all players, some proposer must successfully propose it in equilibrium.

Finally, S being the ideal alternative of all the players, it is easy to construct an equilibrium in which all players always propose alternative S (conditional on being recognized to propose), accept any proposal to change status quo R to alternative S, and reject any proposal to amend status quo S. □

We now return to the proof of the main proposition. Suppose first that $\big|L\big| < q$. Having characterized the unique PBE outcome of $\Gamma(a, 1)$ and $\Gamma(a, \alpha_k)$, for all $a \in \{R, S\}$ and all $\alpha_k \leq \hat{\alpha}_n$, we begin with an inductive argument. Take any $k \in \mathbb{N}$ such that (i) for each $a \in \{R, S\}$, alternative S is implemented in every period in any PBE of $\Gamma(a, \alpha_{k+1})$, and (ii) $\alpha_k \leq \hat{\alpha}_q$. (From Lemma A1(ii), we already know that this is the case if $\alpha_k \leq \hat{\alpha}_n$.) In the first period of $\Gamma(S, \alpha_k)$, the expected payoff to legislator i in any PBE must be a convex combination of $f_i(\alpha_k) \equiv \big[1 - \delta(1 - \gamma\Delta)\big]\alpha_k\gamma\Delta r_i + \delta(1 - \alpha_k\gamma\Delta)s_i\Delta$ and $s_i\Delta$. It follows from the definition of $\hat{\alpha}_i$, $s_i\Delta > f_i(\alpha_k)$ if and only if $\alpha_k \leq \hat{\alpha}_i$. By the same logic as in the proof of Lemma A1(i), this implies that any proposal to change status quo S to the risky alternative R must be rejected by the members of the winning coalition $\{1, \ldots, q\}$. Hence, $s_i\Delta$ is legislator i's unique equilibrium payoff in $\Gamma(S, \alpha_k)$. It follows that legislator i's payoff in the continuation game $\Gamma(R, \alpha_k)$ is $f_i(\alpha_k)$ if R is implemented in the first period, and Δs_i otherwise. This implies that every member of the winning coalition $\{1, \ldots, q\}$ must accept any proposal to amend R to S (when decisive) and, therefore, at least one proposer must successfully propose S in the first period in equilibrium. We have thus established that for all $\alpha_k \leq \hat{\alpha}_q$, S is implemented in every period of $\Gamma(R, \alpha_k)$ in any PBE. The unique equilibrium outcome of $\Gamma(R, \hat{\alpha}_q)$ is therefore the stoppling rule with cutoff $\hat{\alpha}_q$.

In Section 3, we defined $\overline{V}_i(\alpha_k)$ as the expected payoff to legislator i induced by the stopping rule with cutoff α_k in $\Gamma(S, \alpha_0)$. For every $0 \leq \ell \leq k$, we can similarly define the expected payoff to legislator i induced by this stopping rule in $\Gamma(S, \alpha_\ell)$ as

$$\overline{U}_i(\alpha_k \mid \alpha_\ell) \equiv \big[1 - \delta^{k-\ell}(1 - \gamma\Delta)^{k-\ell}\big]\alpha_\ell\gamma\Delta r_i + \delta^{k-\ell}\big[1 - \alpha_\ell + (1 - \gamma\Delta)^{k-\ell}\alpha_\ell\big]s_i\Delta\ .$$

Differentiating the right side of the above equation with respect to k reveals that $\overline{U}_i(\alpha_k|\alpha_\ell)$

is single-peaked in k and, therefore, in the cutoff α_k: it decreases with α_k if $\alpha_k < \hat{\alpha}_i$, and increases with α_k if $\alpha_k > \hat{\alpha}_i$. Now take any belief $\alpha_\ell > \hat{\alpha}_q$ such that the unique equilibrium outcome of $\Gamma(R, \alpha_{\ell+1})$ is the stoppling rule with cutoff $\hat{\alpha}_q$. (From the previous paragraph, this is the case if $\alpha_{\ell+1} = \hat{\alpha}_q$.) The expected payoff to legislator i in any PBE of $\Gamma(R, \alpha_\ell)$ is a convex combination between $\overline{U}_i(\hat{\alpha}_q \mid \alpha_\ell)$ and $s_i\Delta$. Moreover, it follows from the single-peakedness of $\overline{U}_i(\cdot \mid \alpha_\ell)$ that $\overline{U}_i(\hat{\alpha}_q \mid \alpha_\ell) > s_i\Delta = \overline{U}_i(\alpha_\ell \mid \alpha_\ell)$, for all $i \geq q$ — recall that $\hat{\alpha}_i \leq \hat{\alpha}_q < \alpha_\ell$ for all $i \geq q$. Every member of the blocking coalition $\{q, \ldots, n\}$ therefore rejects any proposal to amend the status quo R to S in the first period of $\Gamma(R, \alpha_\ell)$. This shows in particular that the unique PBE outcome of $\Gamma(R, \alpha_1)$ is the stopping rule with cutoff $\hat{\alpha}_q$.

Finally, consider the first period of the game — i.e., the first period of $\Gamma(S, \alpha_0)$. It follows from the previous paragraph that, in any equilibrium, legislator i's payoff must be a convex combination between $\overline{U}_i(\hat{\alpha}_q) = \overline{U}_i(\hat{\alpha}_q \mid \alpha_0)$ and $s_i\Delta = \overline{U}_i(\alpha_0 \mid \alpha_0)$. Suppose first that $\overline{U}_{n-q+1}(\hat{\alpha}_q) < s_{n-q+1}\Delta$, so that $\overline{U}_i(\hat{\alpha}_q) < s_i\Delta$ for every member i of the blocking coalition $\{1, \ldots, n-q+1\}$. Every member of this coalition must therefore reject any proposal to amend the status quo S to R in the first period. This in turn implies that the stopping rule with cutoff α_0 is the unique PBE outcome. Suppose now that $\overline{U}_{n-q+1}(\hat{\alpha}_q) \geq s_{n-q+1}\Delta$. Denoting the supremum of the PBE payoffs of each legislator i by U_i^{sup}, we thus have $\overline{U}_i(\hat{\alpha}_q) \geq U_i^{\text{sup}} \geq s_i\Delta$ for every member i of the winning coalition $C \equiv \{n-q+1, \ldots, n\}$. This implies that, in any PBE, the payoff of each legislator $i \in C$ from accepting a proposal to amend status quo S to R (when decisive) in the first period of the game — i.e. $\overline{U}_i(\hat{\alpha}_q)$ — must therefore exceed her payoff from rejecting it — $(1-\delta)\Delta s_i + \delta U_i^{\text{sup}}$. Hence, at least one member of the legislature must successfully propose alternative R in the first period. This in turn implies that the stopping rule with cutoff $\hat{\alpha}_q$ is the unique PBE outcome, completing the proof of Proposition 1.

B Proof of Proposition 2

B.1 Proof of Part (i) of Proposition 2

Fix $\tau^{\max} > 0$. To prove the first part of Proposition 2, we will first define the threshold $\overline{\Delta}$ (Subsection B.1.1). Then, for every $\Delta < \overline{\Delta}$, we will construct a stationary Markov strategy profile σ^{Δ} that supports the optimal stopping rule (Subsection B.1.2). Finally, we will demonstrate that, for all $\Delta < \overline{\Delta}$, σ^{Δ} is a renegotiation-proof equilibrium (Subsection B.1.3).

B.1.1 Definition of $\overline{\Delta}$

To begin we must establish some notation. For each $i \in N$, let coalition C^i be defined by $C^i = N \setminus \{n\}$ if $i = 1$, and $C^i = N \setminus \{i-1\}$ otherwise. Note that, as $q < n$, each coalition C^i is winning. Let $x^i \in X$ be defined by

$$x^i_j \equiv \begin{cases} 1/(n-1) & \text{if } j \in C^i \ , \\ 0 & \text{otherwise,} \end{cases}$$

and let

$$\bar{x}_i \equiv \frac{1}{n-1} \sum_{j \colon C^j \ni i} p_j \ ,$$

where $p_j \in (0,1)$ is the probability (induced by the protocol) that legislator j proposes first in any period. To shorten the notation slightly we will henceforth refer to the upper bound on the tax rate $\tau^{\max}$ more concisely as $\hat{\tau}$. Next, let $k^* \in \mathbb{N}$ be implicitly defined by $\alpha_{k^*} \equiv \alpha^*$ and, for every $i, j \in N$, let the function $W^j_i \colon A \to \mathbb{R}_+$ be defined by

$$W^j_i(\alpha) \equiv \begin{cases} w_i(R, \hat{\tau}, x^j \mid 1) & \text{if } \alpha = 1, \\ \\ w_i(S, \hat{\tau}, x^j \mid \alpha) & \text{if } \alpha = \alpha_k \text{ with } k \geq k^*, \\ \\ \big[1 - \delta^{k^*-k}(1-\gamma\Delta)^{k^*-k}\big] w_i(R, \hat{\tau}, x^j \mid \alpha) & \text{if } \alpha = \alpha_k \text{ with } k < k^*, \\ +\delta^{k^*-k}\big[1 - \alpha_k + (1-\gamma\Delta)^{k^*-k}\alpha_k\big] w_i(S, \hat{\tau}, x^j \mid \alpha) & \end{cases}$$

for all $\alpha \in A$. In words: for every α, $W_i^j(\alpha)$ is legislator i's average discounted payoff is implemented along with (constant) redistributive policy $(\hat{\tau}, x^j)$ when the belief is α. Finally, let $W_i^0\colon A \to \mathbb{R}_+$ be defined by $W_i^0(\alpha) \equiv \sum_{j\in N} p_j W_i^j(\alpha)$, for all $\alpha \in A$. The interpretation of $W_i^0(\alpha)$ is analogous to $W_i^j(\alpha)$'s, but each redistributive policy $(\hat{\tau}, x^j)$ is implemented with probability p_j. Observe that, for all $\alpha \in A$ and $j \in N$, the payoff vectors $\big(W_i^j(\alpha)\big)_{i\in N}$ and $\big(W_i^0(\alpha)\big)_{i\in N}$ belong to the Pareto frontier. This observation will play an important role in the equilibrium construction below.

The definition of the threshold $\overline{\Delta}$ hinges on the following lemma:

Lemma B1. *Suppose $\tau^{\max} > 0$ and, for all $i, j \in N$, let W_i^j and W_i^0 be defined as above. There exists $\overline{\Delta} > 0$ such that the following inequalities hold for all $\Delta < \overline{\Delta}$, all $i, j \in N$ with $i \in C^j$, and all $k \in \mathbb{N}$:*

$$\begin{aligned} W_i^j(1) &> (1-\delta)\gamma\Delta\bar{r} + \delta W_i^0(1)\ , \\ W_i^j(\alpha_k) &> (1-\delta)\bar{s}\Delta + \delta W_i^0(\alpha_k)\ ,\ \textit{and} \\ W_i^j(\alpha_k) &> (1-\delta)\alpha_k\gamma\Delta\bar{r} + \delta\alpha_k\gamma\Delta W_i^0(1) + \delta(1-\alpha_k\gamma\Delta)W_i^0(\alpha_{k+1})\ . \end{aligned}$$

Proof. Let $i, j \in N$ with $i \in C^j$, and all $k \in \mathbb{N}$. By definition of W_i, we have

$$\begin{aligned} \frac{W_i^j(1) - (1-\delta)\gamma\Delta\bar{r} - \delta W_i^0(1)}{\Delta} &= (1-\delta)\gamma(1-\hat{\tau})(r_i - \bar{r}) \\ &\quad + \gamma\hat{\tau}\bar{r}\big[\delta(x_i^j - \bar{x}_i) - (1-\delta)(1-x_i^j)\big]\ . \end{aligned}$$

As $x_i^j - \bar{x}_i > 0$, there exists $\hat{\Delta}_{i,j}^1 > 0$ such that $W_i^j(1) - (1-\delta)\gamma\bar{r} - \delta W_i^0(1) > 0$ whenever $\Delta < \hat{\Delta}_{i,j}^1$. By the same logic, if $k \geq k^*$, then there exists $\hat{\Delta}_{i,j}^2 > 0$ such that $W_i^j(\alpha_k) - (1-\delta)\bar{s}\Delta - \delta W_i^0(\alpha_k) = \bar{s}\Delta\big[\delta(x_i^j - \bar{x}_i) - (1-\delta)(1-x_i^j)\big] > 0$ whenever $\Delta < \hat{\Delta}_{i,j}^2$. Now suppose that $k < k^*$. Observe that

$$\begin{aligned} \frac{W_i^j(\alpha_k) - W_i^0(\alpha_k)}{\Delta} =& \Big\{\big[1 - \delta^{k^*-k}(1-\gamma\Delta)^{k^*-k}\big]\alpha_k\gamma\bar{r} + \delta^{k^*-k}\big[1 - \alpha_k + (1-\gamma\Delta)^{k^*-k}\alpha_k\big]\bar{s}\Big\} \\ &\times (x_i^j - \bar{x}_i)\ , \end{aligned}$$

where the first bracketed term on the right-hand side represents the expected social welfare (divided by Δ) under the optimal stopping rule. As $\alpha_k > \alpha^*$, this term is greater than or equal to $\bar{s}$. Hence, $W_i^j(\alpha_k) - W_i^0(\alpha_k) \geq \bar{s}\Delta(x_i^j - \bar{x}_i) > 0$, and

$$\frac{W_i^j(\alpha_k) - (1-\delta)\bar{s}\Delta - \delta W_i^0(\alpha_k)}{\Delta} \geq \frac{1-\delta}{\Delta} W_i^j(\alpha_k) - (1-\delta)\bar{s} + \frac{\delta}{\Delta}\bar{s}\hat{\tau}(x_i^j - \bar{x}_i) \ .$$

An application of l'Hôpital's rule shows that $(1-\delta)/\Delta \to \rho$ as $\Delta \to 0$. As $W_i^j(\cdot)$ and $W_i^0(\cdot)$ are bounded, there exists $\hat{\Delta}_{i,j}^3 > 0$ such that $W_i^j(\alpha_k) > (1-\delta)\bar{s}\Delta - \delta W_i^0(\alpha_k) > 0$ whenever $\Delta < \hat{\Delta}_{i,j}^3$.

Consider now the last inequality in the lemma. Let $\Psi(\alpha_k) \equiv W_i^j(\alpha_k) - (1-\delta)\alpha_k\gamma\Delta\bar{r} - \delta\alpha_k\gamma\Delta W_i^0(1) - \delta(1-\alpha_k\gamma\Delta)W_i^0(\alpha_{k+1})$. Suppose first that $k \geq k^*$. It is readily checked that

$$\lim_{\Delta\to 0}\frac{\Psi(\alpha_k)}{\Delta} = (1-\hat{\tau})s_i + \hat{\tau}x_i^j\bar{s} - \big[(1-\hat{\tau})s_i + \hat{\tau}\bar{x}_i\bar{s}\big] = \hat{\tau}(x_i^j - \bar{x}_i)\bar{s} > 0 \ .$$

Therefore, there exists $\hat{\Delta}_{i,j}^4 > 0$ such that $W_i^j(\alpha_k) > (1-\delta)\alpha_k\gamma\bar{r} + \delta\alpha_k\gamma\Delta W_i^0(1) + \delta(1-\alpha_k\gamma\Delta)W_i^0(\alpha_{k+1})$ whenever $\Delta < \hat{\Delta}_{i,j}^4$. Finally, suppose that $k < k^*$. By definition of W_i^j, we have

$$\begin{aligned}\Psi(\alpha_k) &= -(1-\delta)\alpha_k\gamma\Delta\bar{r}(1-x_i^j) + \delta\alpha_k\gamma\Delta\big[W_i^j(1) - W_i^0(1)\big] \\ &\qquad\qquad\qquad\qquad + \delta(1-\alpha_k\gamma\Delta)\big[W_i^j(\alpha_{k+1}) - W_i^0(\alpha_{k+1})\big] \\ &\geq -(1-\delta)\alpha_k\gamma\Delta\hat{\tau}(1-x_i^j)\bar{r} + \delta(x_i^j - \bar{x}_i)\left[\alpha_k\gamma\Delta\hat{\tau}\bar{r} + (1-\alpha_k\gamma\Delta)\hat{\tau}\bar{s}\right] \\ &> -(1-\delta)\alpha_k\gamma\Delta\hat{\tau}(1-x_i^j)\bar{r} + \delta\hat{\tau}(x_i^j - \bar{x}_i)\bar{s} \ ,\end{aligned}$$

where the first inequality follows from $W_i^j(\alpha_{k+1}) - W_i^0(\alpha_{k+1}) \geq \bar{s}\hat{\tau}(x_i^j - \bar{x}_i)$ (as established above), and the second follows from our assumption that $\gamma\bar{r} > \bar{s}$. Therefore, there exists $\hat{\Delta}_{i,j}^5 > 0$ such that $W_i^j(\alpha_k) > (1-\delta)\alpha_k\gamma\Delta\bar{r} + \delta\alpha_k\gamma\Delta W_i^0(1) + \delta(1-\alpha_k\gamma\Delta)W_i^0(\alpha_{k+1})$ whenever $\Delta < \hat{\Delta}_{i,j}^5$. Setting $\bar{\Delta} \equiv \min\{\hat{\Delta}_{i,j}^\ell \colon i,j \in N \ \& \ \ell = 1,\ldots,5\}$, we obtain the lemma. $\square$

Let $\overline{\Delta}$ be defined as in the above lemma. Henceforth, we assume that $\Delta < \bar{\Delta}$.

B.1.2 Definition of Stationary Markov Strategy Profile σ^Δ

This subsection describes the behavior prescribed by strategy profile σ^Δ to each legislator $i \in N$, for all $\Delta < \bar{\Delta}$. Observe that, in each proposal stage, i's behavior only depends on the current status quo and belief and, in each voting stage, her behavior only depends on the current status quo, the belief, and the list of remaining proposers in the current period. Hence, σ^{Δ} is stationary Markov.

• **Proposal stages.** Consider first proposer i's behavior in a period where the order of proposers is $\pi = (\pi_1, \ldots, \pi_n)$ with $\pi_\ell = i$ for some ℓ; and the first $\ell - 1$ proposers have failed to amend the status quo. There are three cases:

Case P1: The belief is α_k, where $k < k^*$.

Proposer i offers $(R, \hat{\tau}, x^i)$ (which, in cases where the status quo is $(R, \hat{\tau}, x^i)$ means that she passes).

Case P2: The belief is α_k, where $k > k^*$.

Proposer i offers $(S, \hat{\tau}, x^i)$ (which, in cases where the status quo is $(S, \hat{\tau}, x^i)$ means that she passes).

Case P3: The belief is α^*.

Case 3.1: If the status quo is a policy $(a, \tau, x) \neq (R, \hat{\tau}, x^j)$ for all $j \in N$, then proposer i offers $(S, \hat{\tau}, x^i)$.

Case 3.2: If the status quo is $(R, \hat{\tau}, x^j)$ for some $j \in N$, then proposer i offers $(S, \hat{\tau}, x^j)$ if $i \in C^j$, and $(S, \hat{\tau}, x^i)$ otherwise.

• **Voting stages.** Consider now voter i's behavior in a period where the order of proposers is $\pi = (\pi_1, \ldots, \pi_n)$. There are several cases:

Case V1: The status quo is $(R, \hat{\tau}, x^j)$ for some $j \in N$; the belief is α_k, where $k < k^*$; and a proposer π_ℓ has just proposed policy $(a, \tau, y) \neq (R, \hat{\tau}, x^j)$.

If voter i is a member of C^j, then she rejects the proposal; otherwise, she accepts the proposal if and only if

$$W_i^j(\alpha_k) > (1-\delta) w_i(a, \tau, y \mid \alpha_k) + \delta \begin{cases} \alpha_k \gamma \Delta W_i^0(1) + (1 - \alpha_k \gamma \Delta) W_i^0(\alpha_{k+1}) & \text{if } a = R, \\ W_i^0(\alpha_k) & \text{if } a = S. \end{cases}$$

Case V2: The status quo is (a, τ, x), where $(a, \tau, x) \neq (R, \hat{\tau}, x^j)$ for all $j \in N$; the belief is α_k, where $k < k^*$.

Case V2.1: Proposer π_n has just proposed policy $(R, \hat{\tau}, x^j)$ for some $j \in N$.

If voter i is a member of C^j, then she accepts the proposal; otherwise, she accepts the proposal if and only if

$$W_i^j(\alpha_k) > (1-\delta)w_i(a, \tau, x \mid \alpha_k) + \delta \begin{cases} \alpha_k \gamma \Delta W_i^0(1) + (1 - \alpha_k \gamma \Delta) W_i^0(\alpha_{k+1}) & \text{if } a = R, \\ W_i^0(\alpha_k) & \text{if } a = S. \end{cases}$$

Case V2.2: Proposer π_n has just proposed policy (b, τ', y), where $(b, \tau', y) \neq (R, \hat{\tau}, x^j)$ for all $j \in N$.

Voter i accepts the proposal if and only if

$$(1 - \delta)w_i(a, \tau, x \mid \alpha_k) + \delta \begin{cases} \alpha_k \gamma \Delta W_i^0(1) + (1 - \alpha_k \gamma \Delta) W_i^0(\alpha_{k+1}) & \text{if } a = R, \\ W_i^0(\alpha_k) & \text{if } a = S. \end{cases}$$

$$< (1 - \delta)w_i(b, \tau', y \mid \alpha_k) + \delta \begin{cases} \alpha_k \gamma \Delta W_i^0(1) + (1 - \alpha_k \gamma \Delta) W_i^0(\alpha_{k+1}) & \text{if } b = R, \\ W_i^0(\alpha_k) & \text{if } b = S. \end{cases}$$

Case V2.3: Proposer π_ℓ, $\ell < m$, has just proposed policy $(R, \hat{\tau}, x^j)$ for some $j \in N$.

If voter i is a member of C^j, then she accepts the proposal; otherwise, she accepts the proposal if and only if $W_i^j(\alpha_k) \geq W_i^{\pi_{\ell+1}}(\alpha_k)$.

Case V2.4: Proposer π_ℓ, $\ell < m$, has just proposed policy (b, τ', y), where $(b, \tau', y) \neq (R, \hat{\tau}, x^j)$ for all $j \in N$.

Voter i accepts the proposal if and only if

$$W_i^{\pi_{\ell+1}}(\alpha_k) < (1-\delta)w_i(b, \tau', y \mid \alpha_k) + \delta \begin{cases} \alpha_k \gamma \Delta W_i^0(1) + (1 - \alpha_k \gamma \Delta) W_i^0(\alpha_{k+1}) & \text{if } b = R, \\ W_i^0(\alpha_k) & \text{if } b = S. \end{cases}$$

Case V3: The status quo is $(S, \hat{\tau}, x^j)$ for some $j \in N$; the belief is α_k, where $k > k^*$; and a proposer π_ℓ has just proposed policy $(a, \tau, y) \neq (S, \hat{\tau}, x^j)$.

If voter i is a member of C^j, then she rejects the proposal; otherwise, she accepts the proposal if and only if

$$W_i^j(\alpha_k) > (1-\delta)w_i(a,\tau,y \mid \alpha_k)+\delta \begin{cases} \alpha_k\gamma\Delta W_i^0(1) + (1-\alpha_k\gamma\Delta)W_i^0(\alpha_{k+1}) & \text{if } a = R, \\ W_i^0(\alpha_k) & \text{if } a = S. \end{cases}$$

Case V4: The status quo is (a, τ, x), where $(a, \tau, x) \neq (S, \hat{\tau}, x^j)$ for all $j \in N$; and the belief is α_k, where $k > k^*$.

Case V4.1: Proposer π_n has just proposed policy $(S, \hat{\tau}, x^j)$ for some $j \in N$.

If voter i is a member of C^j, then she accepts the proposal; otherwise, she accepts the proposal if and only if

$$W_i^j(\alpha_k) \geq (1-\delta)w_i(a,\tau,x \mid \alpha_k)+\delta \begin{cases} \alpha_k\gamma\Delta W_i^0(1) + (1-\alpha_k\gamma\Delta)W_i^0(\alpha_{k+1}) & \text{if } a = R, \\ W_i^0(\alpha_k) & \text{if } a = S. \end{cases}$$

Case V4.2: Proposer π_n has just proposed policy (b, τ', y), where $(b, \tau', y) \neq (S, \hat{\tau}, x^j)$ for all $j \in N$.

Voter i accepts the proposal if and only if

$$\begin{aligned} &(1-\delta)w_i(a,\tau,x \mid \alpha_k) + \delta \begin{cases} \alpha_k\gamma\Delta W_i^0(1) + (1-\alpha_k\gamma\Delta)W_i^0(\alpha_{k+1}) & \text{if } a = R, \\ W_i^0(\alpha_k) & \text{if } a = S. \end{cases} \\ &\leq (1-\delta)w_i(b,\tau',y \mid \alpha_k) + \delta \begin{cases} \alpha_k\gamma\Delta W_i^0(1) + (1-\alpha_k\gamma\Delta)W_i^0(\alpha_{k+1}) & \text{if } b = R, \\ W_i^0(\alpha_k) & \text{if } b = S. \end{cases} \end{aligned}$$

Case V4.3: Proposer π_ℓ, $\ell < m$, has just proposed policy $(S, \hat{\tau}, x^j)$ for some $j \in N$.

If voter i is a member of C^j, then she accepts the proposal; otherwise, she accepts the proposal if and only if $W_i^j(\alpha_k) \geq W_i^{\pi_{\ell+1}}(\alpha_k)$.

Case V4.4: Proposer π_ℓ, $\ell < m$, has just proposed policy (b, τ', y), where $(b, \tau', y) \neq (S, \hat{\tau}, x^j)$ for all $j \in N$.

Voter i accepts the proposal if and only if

$$W_i^{\pi_{\ell+1}}(\alpha_k) < (1-\delta)w_i(b,\tau',y \mid \alpha_k)+\delta \begin{cases} \alpha_k\gamma\Delta W_i^0(1) + (1-\alpha_k\gamma\Delta)W_i^0(\alpha_{k+1}) & \text{if } b = R, \\ W_i^0(\alpha_k) & \text{if } b = S. \end{cases}$$

Case V5: The status quo is $(R, \hat{\tau}, x^j)$ for some $j \in N$; and the belief is α^*.

Case V5.1: Proposer π_n has just proposed policy $\left(S, \hat{\tau}, x^{j'}\right)$ for some $j' \in N$.

If voter i is a member of $C^{j'}$, then she accepts the proposal; otherwise, she accepts the proposal if and only if

$$W_i^{j'}(\alpha^*) \geq (1-\delta) w_i(R, \hat{\tau}, x^j \mid \alpha^*) + \delta \alpha^* \gamma \Delta W_i^0(1) + \delta(1 - \alpha^* \gamma \Delta) W_i^0(\alpha_{k^*+1} .$$

Case V5.2: Proposer π_n has just proposed policy (b, τ', y), where $(b, \tau', y) \neq (S, \hat{\tau}, x^j)$ for all $j \in N$.

Voter i accepts the proposal if and only if

$$\begin{aligned} &(1-\delta) w_i(R, \hat{\tau}, x^j \mid \alpha^*) + \delta \alpha^* \gamma \Delta W_i^0(1) + \delta(1 - \alpha^* \gamma \Delta) W_i^0(\alpha_{k^*+1}) \\ &\leq (1-\delta) w_i(b, \tau', y \mid \alpha^*) + \delta \begin{cases} \alpha^* \gamma \Delta W_i^0(1) + (1 - \alpha^* \gamma \Delta) W_i^0(\alpha_{k^*+1}) & \text{if } b = R, \\ W_i^0(\alpha^*) & \text{if } b = S. \end{cases} \end{aligned}$$

Case V5.3: Proposer π_ℓ, $\ell < m$, has just proposed policy $\left(S, \hat{\tau}, x^{j'}\right)$ for some $j' \in N$.

If voter i is a member of C^j and $j' = j$, then she accepts the proposal; if she is a member of C^j, $j' \neq j$ and there is $\ell' > \ell$ such that $\pi_{\ell'} \in C^j$, then she rejects the proposal; otherwise, she accepts the proposal if and only if $W_i^{j'}(\alpha^*) > W_i^{\pi_{\ell+1}}(\alpha^*)$.

Case V5.4: Proposer π_ℓ, $\ell < m$, has just proposed policy (b, τ', y), where $(b, \tau', y) \neq (R, \hat{\tau}, x^j)$ for all $j \in N$.

Voter i accepts the proposal if and only if

$$W_i^{\pi_{\ell+1}}(\alpha^*) < (1-\delta) w_i(b, \tau', y \mid \alpha^*) + \delta \begin{cases} \alpha^* \gamma \Delta W_i^0(1) + (1 - \alpha^* \gamma \Delta) W_i^(\alpha_{k^*+1}) & \text{if } b = R, \\ W_i^0(\alpha^*) & \text{if } b = S. \end{cases}$$

Case V6: The status quo is (a, τ, x), with $(a, \tau, x) \neq (R, \hat{\tau}, x^j)$ for all $j \in N$; and the belief is α^*.

Voter i behaves as in Cases 3 and 4 (with $k = k^*$).

Case V7: The belief is equal to one. In this case, apply Cases V1 and V2 replacing α_k and α_{k+1} by 1.

B.1.3 Verification that σ^{Δ} Is a Renegotiation-proof equilibrium

Optimal stopping rule. Before proceeding to the verification that σ^{Δ} is a renegotiation-proof equilibrium, it is worth noting that it sustains the optimal stopping rule. The game starts with status quo $(S, 0, x^0)$ and belief α_0. The first proposer, say j, is prescribed to propose policy $(R, \hat{\tau}, x^j)$ (see Case P1 above), which is accepted by all the members of decisive coalition C^j (see Case V2.3). This policy is then implemented again in every future period that begins with a belief greater than α^*, as any proposal to amend it is voted down by the members of C^j (see Case V1). If the belief becomes α^*, then all members of C^j reject any proposal until one of them proposes policy $(S, \hat{\tau}, x^j)$, which they all accept (see Cases P3.2 and V5.3). (As C^j is a decisive coalition, $C_j \cap \{\pi_1, \ldots, \pi_m\} \neq \emptyset$ and, therefore, at least one of its members is a proposer.) Policy $(S, \hat{\tau}, x^j)$ is then never amended, as any proposal to change it is voted down by the members of C^j (see Case V3). Hence, the optimal stopping rule is implemented on the path. Any deviation from this path leads the next proposer j' to successfully propose policy $\left(R, \hat{\tau}, x^{j'}\right)$ if the belief is greater than α^*, or policy $\left(S, \hat{\tau}, x^{j'}\right)$ if the belief smaller than or equal to α^*. The induced path again supports the optimal stopping rule.

Continuation values and renegotiation-proofness. Let $V_i(b, \tau, y|\alpha)$ be player i's average discounted value (induced by σ^{Δ}) from implementing policy (b, τ, y) when the belief is equal to α. Observe first that if a policy $(R, \hat{\tau}, x^j)$, with $j \in N$, is implemented in a period that begins with belief α_k, $k < k^*$, then it is also implemented in any future period beginning with a belief greater than α^* (Cases V1 and V7 above). If the belief becomes equal to α^*, then $(R, \hat{\tau}, x^j)$ is amended to policy $(S, \hat{\tau}, x^j)$ (see Cases P3.2, V5.1 and V5.3), which is then implemented in every future period (see Cases V6 and V3). This implies that $V_i(R, \hat{\tau}, x^j|\alpha_k) = W_i^j(\alpha_k)$ for all $i, j \in N$ and $k < k^*$. Similar arguments establish that $V_i(R, \hat{\tau}, x^j|1) = W_i^j(1)$ and $V_i(S, \hat{\tau}, x^j|\alpha_k) = W_i^j(\alpha_k)$, for all $i, j \in N$ and $k \geq k^*$. By construction of σ^{Δ}, these are all the possible continuation values induced by σ^{Δ} at the start of any continuation game. As they all belong to the Pareto frontier (Subsection B.1.1), this implies that if σ^{Δ} is an equilibrium, then it must be renegotiation proof.

Now suppose that a policy (b, τ, y), where $(b, \tau, y) \neq (R, \hat{\tau}, x^j)$ for all $j \in N$, is implemented in a period that starts with a belief α_k, $k < k^*$. Player i receives $(1-\delta)w_i(R, \tau, y \mid \alpha_k)$ in that period. If b is alternative R, then there are two possible cases: it succeeds with probability $\alpha_k \gamma \Delta$, in which case the next period's first proposer j successfully offers policy $(R, \hat{\tau}, x^j)$ (see Cases P1, V7 and V1); and it fails with probability $1 - \alpha_k \gamma \Delta$ in which case, the next period's first proposer j successfully offers policy $(R, \hat{\tau}, x^j)$ if $k < k^* - 1$ (see Case V2.3), or $(S, \hat{\tau}, x^j)$ if $k = k^* - 1$ (see Cases V6 and V4.3). Therefore,

$$
\begin{aligned}
V_i(R, \tau, y|\alpha_k) &= (1-\delta)w_i(R, \tau, y \mid \alpha_k) + \delta \alpha_k \gamma \Delta \sum_{j \in N} p_j V_i(R, \hat{\tau}, x^j|1) \\
&\quad + \delta(1 - \alpha_k \gamma \Delta) \sum_{j \in N} p_j \begin{cases} V_i(R, \hat{\tau}, x^j|\alpha_{k+1}) & \text{if } k < k^* - 1 \\ V_i(S, \hat{\tau}, x^j|\alpha^*) & \text{if } k = k^* - 1 \end{cases} \\
&= (1-\delta)w_i(b, \tau, y \mid \alpha_k) + \delta \alpha_k \gamma \Delta \sum_{j \in N} p_j W_i^j(1) \\
&\quad + \delta(1 - \alpha_k \gamma \Delta) \sum_{j \in N} p_j \begin{cases} W_i^j(\alpha_{k+1}) & \text{if } k < k^* - 1 \\ W_i^j(\alpha^*) & \text{if } k = k^* - 1 \end{cases} \\
&= (1-\delta)w_i(R, \tau, y \mid \alpha_k) + \delta\big[\alpha_k \gamma \Delta W_i^0(1) + (1-\alpha_k)\gamma \Delta W_i^0(\alpha_{k+1})\big] \ .
\end{aligned}
$$

If b is alternative S then, in the next period, the first proposer j successfully offers policy $(R, \hat{\tau}, x^j)$ (see Cases P1 and V2.3); so that

$$
\begin{aligned}
V_i(S, y|\alpha_k) &= (1-\delta)w_i(S, \tau, y \mid \alpha_k) + \delta \sum_{j \in M} p_j V_i(s, \hat{\tau}, x^j|\alpha_k) \\
&= (1-\delta)w_i(S, \tau, y \mid \alpha_k) + \delta \sum_{j \in N} p_j W_i^j(1) \\
&= (1-\delta)w_i(S, \tau, y \mid \alpha_k) + \delta W_i^0(1) \ .
\end{aligned}
$$

Using parallel arguments, one can show that player i's continuation value from implementing a policy (b, τ, y), where $(b, \tau, y) \neq (S, \hat{\tau}, x^j)$ for all $j \in N$, in a period with belief α_k, $k \geq k^*$, is given by

$$
V_i(b, \tau, y|\alpha_k) = (1-\delta)w_i(b, \tau, y \mid \alpha_k) + \delta \begin{cases} W_i^0(\alpha_{k+1}) & \text{if } b = r \ , \\ W_i^0(\alpha_k) & \text{if } b = s \ , \end{cases}
$$

and that her continuation value from implementing a policy (b, τ, y), where $(b, \tau, y) \neq (R, \hat{\tau}, x^j)$ for all $j \in N$, in a period where the belief is equal to one is given by

$$V_i(b, \tau, y|1) = (1-\delta)w_i(b, \tau, y \mid 1) + \delta W_i^0(1) \ .$$

It follows directly from this characterization of continuation values and from Lemma **??** that $(R, \hat{\tau}, x^j)$ with $i \in C^j$ is player i's ideal policy when the belief is greater than α^* — i.e., $V_i(R, \hat{\tau}, x^j|\alpha) \geq V_i(b, \tau, y|\alpha)$ for all $i, j \in N$ such that $i \in C^j$, $(b, \tau, y) \in \{R, S\} \times [0, \hat{\tau}] \times X$, and $\alpha > \alpha^*$ — and that $(S, \hat{\tau}, x^j)$ with $i \in C^j$ is her ideal policy when the belief is smaller than or equal to α^* — i.e., $V_i(S, \hat{\tau}, x^j|\alpha) \geq V_i(b, \tau, y|\alpha)$ for all $i, j \in N$ such that $i \in C^j$, $(b, \tau, y) \in \{R, S\} \times [0, \hat{\tau}] \times X$, and $\alpha \leq \alpha^*$.

Voting stages. To verify that σ^Δ is an equilibrium, we will first check that all possible deviations in voting stages are unprofitable. To do so, we will consider in turn the various cases in the definition of voting strategies.

In Case V1, (decisive) voter i receives a payoff of $V_i(R, \hat{\tau}, x^j|\alpha_k) = W_i^j(\alpha_k)$ if she rejects the proposal (a, τ, y) to amend status quo $(R, \hat{\tau}, x^j)$ (as any future attempt to amend it in this period will be unsuccessful), and a payoff of

$$V_i(a, \tau, y|\alpha_k) = (1-\delta)w_i(a, \tau, y \mid \alpha_k) + \delta \begin{cases} \alpha_k \gamma \Delta W_i^0(1) + (1 - \alpha_k \gamma \Delta) W_i^0(\alpha_{k+1}) & \text{if } a = R \ , \\ W_i^0(\alpha_k) & \text{if } a = S \ , \end{cases}$$

if she accepts it. Hence, she cannot profitably deviate from σ^Δ if she is not a member of C^j. Moreover, it follows from Lemma **??** and the above equality that $V_i(a, \tau, y|\alpha_k) < W_i^j(\alpha_k) = V_i(R, \hat{\tau}, x^j|\alpha_k)$ for all $i \in C^j$, so that voter i cannot profitably deviate from rejecting (a, τ, y) if she is a member of C^j.

It follows immediately from our characterization of continuation values above that, in Cases V2.1 and V2.2, σ^Δ prescribes voter i to accept the last proposer's offer if and only if her continuation value from implementing this offer exceeds her continuation value from implementing the status quo. Therefore, deviations are also unprofitable in these cases. In case V2.3, (decisive) player i anticipates that if she rejects the ℓth proposer's offer, $(R, \hat{\tau}, x^j)$, then the next proposer will successfully propose $(R, \hat{\tau}, x^{\pi_{\ell+1}})$. It is therefore optimal for her to accept $(R, \hat{\tau}, x^j)$ if and only if $V_i(R, \hat{\tau}, x^j|\alpha_k) = W_i^j(\alpha_k) \geq W_i^{\pi_{\ell+1}}(\alpha_k) = V_i\big(R, \hat{\tau}, x^{\pi_{\ell+1}}|\alpha_k\big)$, as prescribed by σ^Δ to every $i \notin C^j$. Moreover, by definition of W_i^j, $W_i^j(\alpha_k) \geq W_i^{j'}(\alpha_k)$ for all $i, j, j' \in N$ such that $i \in C^j$. Therefore, it is always optimal for voter i to accept $(R, \hat{\tau}, x^j)$ if $i \in C^j$. The same argument applies to Case 2.4 except that in this case, the ℓth proposal offer is some $(b, \tau, y) \neq (R, \hat{\tau}, x^j)$ for all $j \in N$, so that the value from accepting it is equal to

$$V_i(b, \tau, y) = (1-\delta)w_i(b, \tau, y \mid \alpha_k) + \delta \begin{cases} \alpha_k\gamma\Delta W_i^0(1) + (1-\alpha_k\gamma\Delta)W_i^0(\alpha_{k+1}) & \text{if } b = R, \\ W_i^0(\alpha_k) & \text{if } b = S. \end{cases}$$

The arguments to show that there is no profitable deviation from σ^Δ in all the other cases, but Case V5.3, are analogous: in each of these cases, σ^Δ prescribes decisive voter i the action that maximizes her continuation value. In Case V5.3, (decisive) voter i anticipates that if she rejects the ℓth proposer's offer, $\big(R, \hat{\tau}, x^{j'}\big)$, then the next proposer will successfully propose $(R, \hat{\tau}, x^{\pi_{\ell+1}})$, and she will consequently receive $V_i\big(R, \hat{\tau}, x^{\pi_{\ell+1}}|\alpha^*\big) = W_i^{\pi_{\ell+1}}(\alpha^*)$. If she is a member of coalition C^j and $j' = j$, then it is optimal for her to accept the offer (as prescribed by σ^Δ): by definition of W_i, $V_i\big(R, \hat{\tau}, x^{j'}|\alpha^*\big) = V_i(R, \hat{\tau}, x^j|\alpha^*) = W_i^j(\alpha^*) \geq W_i^{\pi_{\ell+1}}(\alpha^*)$. This implies that every member of coalition C^j knows that if the status quo is $(R, \hat{\tau}, x^j)$ in a period where the belief is α^*, then the first proposer in C^j will successfully offer policy $(S, \hat{\tau}, x^j)$. It therefore follows from Lemma **??** and our the characterization of continuation values above that every member i of C^j obtains her highest possible continuation value, $V_i(S, \hat{\tau}, x^j|\alpha^*) = W_i^j(\alpha^*)$, by rejecting any offer until a proposer in C^j successfully offers $(S, \hat{\tau}, x^j)$ (as prescribed by σ^Δ). Finally, if i is not a member of C^j, or if (off the path) all the proposers in C^j have failed to amend the status quo, then σ^Δ optimally prescribes her, as in the previous cases, to choose the action that maximizes her continuation value.

Proposal stages. Consider now player i's behavior in a period where she is the ℓth proposer and the first $\ell - 1$ proposers have failed to amend the status quo. We begin with cases where the belief α is greater than α^*. If the status quo is $(R, \hat{\tau}, x^j)$ for some $j \in N$, then any proposal to amend it is voting down by decisive coalition C^j (see Cases V1 and

V7). Therefore, proposer i's payoff will be $V_i(R, \hat{\tau}, x^j|\alpha)$, irrespective of the action she takes. If the status quo is a policy $(a, \tau, x) \neq (R, \hat{\tau}, x^j)$ for all $j \in N$ and she proposes $(R, \hat{\tau}, x^i)$ (as prescribed by σ^Δ), then her proposal is accepted by all the members of C^i (see Cases V2.1 and V2.3) and she obtains a payoff of $V_i(R, \hat{\tau}, x^i|\alpha) = W_i^i(\alpha)$. As we established above, this is the highest payoff she can get if $\alpha > \alpha^*$. Therefore, any deviation from σ^Δ must be unprofitable.

Suppose now that the belief α is smaller than α^*. If the status quo is $(S, \hat{\tau}, x^j)$ for some $j \in N$, then any proposal to amend it is voting down by decisive coalition C^j (see Case V3). Therefore, proposer i's payoff will be $V_i(S, \hat{\tau}, x^j|\alpha)$, irrespective of the action she takes. If the status quo is a policy $(a, \tau, x) \neq (S, \hat{\tau}, x^j)$ for all $j \in N$ and she proposes $(S, \hat{\tau}, x^i)$ (as prescribed by σ^Δ), then her proposal is accepted by all the members of C^i (see Cases V4.1 and V4.3) and she obtains a payoff of $V_i(S, \hat{\tau}, x^i|\alpha) = W_i^i(\alpha)$. As we established above, this is the highest payoff she can get if $\alpha \leq \alpha^*$. Therefore, any deviation from σ^Δ must again be unprofitable.

Finally, suppose that the belief is equal to α^*. If the status quo is a policy $(a, \tau, x) \neq (R, \hat{\tau}, x^j)$ for all $j \in N$, then the proof that proposer i cannot deviate from proposing $(S, \hat{\tau}, x^i)$ is the same as in the previous paragraph. If the status quo is a policy $(R, \hat{\tau}, x^j)$ for some $j \in N$, then there are several possible cases:

(i) If i is a member of C^j and she proposes $(R, \hat{\tau}, x^j)$, then her proposal is accepted by all the the members of decisive coalition C^j. As this policy is one of her ideal policies when the belief is α^*, she cannot profitably deviate from the behavior prescribed in Case P3.2.

(ii) If i is not a member of C^j and all proposers in C^j have already (unsuccessfully) proposed, then all the members of decisive coalition C^i accept proposal $(S, \hat{\tau}, x^i)$ (see Case V5.1 and the last sub-case in Case V5.3). As we established above, this is the policy that maximizes her continuation value when the belief is α^*. It is therefore impossible for her to profitably deviate from proposing it, as prescribed in Case P3.1.

(iii) If i is not a member of C^j and some proposers in C^j have not yet proposed, then any proposal that differs from $(S, \hat{\tau}, x^j)$ is rejected by the members of C^j (second sub-case in Case V5.3) and, by the end of the period, some proposer in C^j will successfully propose $(S, \hat{\tau}, x^j)$. This implies that, irrespective of the proposal she makes, proposer i's payoff will be $V_i(S, \hat{\tau}, x^j | \alpha^*)$. Therefore, any deviation is unprofitable.

This proves that σ^Δ is a renegotiation-proof stationary equilibrium, thus completing the proof of part (i) of Proposition 2.

B.2 Proof of Part (ii) of Proposition 2

Suppose now that $\tau^{\max} = 1$. In this case, the Pareto frontier of the set of feasible payoff vectors is the simplex $P(\Delta) \equiv \big\{(w_1, \ldots, w_n) \in \mathbb{R}^n_+ \colon \sum_{i \in N} w_i = V^*(\alpha_0)\big\}$ for all $\Delta > 0$. Fix $\varepsilon > 0$; let $\overline{V}^*(\alpha_0) \equiv \lim_{\Delta \to 0} V^*(\alpha_0)$; let $P^* \equiv \big\{(w_1, \ldots, w_n) \in \mathbb{R}^n_+ \colon \sum_{i \in N} w_i = \overline{V}^*(\alpha_0)\big\}$; and let $W_\varepsilon \equiv \big\{(w_1, \ldots, w_n) \in \mathbb{R}^n_+ \colon \sum_{i \in N} w_i \leq \overline{V}^*(\alpha_0) - \varepsilon\big\}$. As $\varepsilon > 0$, it is readily checked that there exists $\bar{\lambda} \in (0, 1)$ such that every vector in W_ε is strictly Pareto dominated by some vector in $P^*_{\bar{\lambda}} \equiv \big\{(w_1, \ldots, w_n) \in P^* \colon w_i \geq \bar{\lambda}\overline{V}^*(\alpha_0), \forall i \in N\big\}$; that is, for every $w \in W_\varepsilon$, there is $w' \in P^*_{\bar{\lambda}}$ such that $w'_i > w_i$, for all $i \in N$.

In the previous subsection, we provided an equilibrium construction which, in the case where $\tau^{\max} = 1$, gives each legislator i an expected payoff of

$$\bar{x}_i V^*(\alpha_0) = V^*(\alpha_0) \frac{1}{n-1} \sum_{j \colon C^j \ni i} p_j \ ,$$

for arbitrarily small values of Δ. A similar construction yields any payoff vector of the form $\big(\lambda_1 V^*(\alpha_0), \ldots, \lambda_n V^*(\alpha_0)\big)$, with $\lambda = (\lambda_1, \ldots, \lambda_n) \in \big\{(\lambda'_1, \ldots, \lambda'_n) \in [\bar{\lambda}, 1] \colon \sum_{i=1}^n \lambda_i = 1\big\}$, if one replaces the revenue distribution x^i proposed by legislator i (if recognized) in the first period by $y^i \in X$, defined by

$$y^i_j(\lambda) \equiv \begin{cases} \lambda_j + \eta \left(1 - \sum_{\ell \colon C^\ell \ni j} p_\ell\right) & \text{if } j \in C^i \ , \\ \lambda_j - \eta \sum_{\ell \colon C^\ell \ni j} p_\ell & \text{if } j \in C^i \ , \end{cases}$$

where $\eta > 0$ is chosen in such a way that $y^i_j(\lambda) \geq 0$ for all $i, j \in N$. That is, for every λ in the compact set $\big\{(\lambda'_1, \ldots, \lambda'_n) \in [\bar{\lambda}, 1] \colon \sum_{i=1}^n \lambda_i = 1\big\}$, there exists $\overline{\Delta}(\lambda) > 0$, such that the payoff vector $\big(\lambda_1 V^*(\alpha_0), \ldots, \lambda_n V^*(\alpha_0)\big)$ is supported by a renegotiation-proof equilibrium whenever $\Delta < \overline{\Delta}(\lambda)$. Moreover, inspection of the proof of Lemma B1 reveals that the threshold $\overline{\Delta}(\lambda)$ can be taken to be continuous in λ. It follows that there is a uniform threshold $\overline{\Delta}_0 > 0$ such that any payoff vector in $\big\{(w_1, \ldots, w_n) \in P(\Delta) \colon w_i \geq \bar{\lambda}\overline{V}^*(\alpha_0), \forall i \in N\big\}$ can be supported by a renegotiation-proof equilibrium whenever $\Delta < \overline{\Delta}_0$. Coupled with the previous paragraph, this observation yields the second part of Proposition 2.

C Proof of Proposition 4

As in Section A, we use the notation $\Gamma(p, \alpha)$ to represent the game that begins with an initial status quo p and in which legislators initially hold belief α.

Lemma C1. *Let $q = n$, $\tau \in [0, 1]$ and $x \in X$. For all $\Delta > 0$:*
(i) $\Gamma(R, \tau, x \mid 1)$ has a renegotiation-proof equilibrium $\sigma^1_{\tau,x}$ and, in any such an equilibrium, each player i's payoff is $\gamma\Delta\big[(1-\tau)r_i + \tau x_i \bar{r}\big]$; and
(ii) $\Gamma(S, \tau, x \mid 1)$ has a renegotiation-proof equilibrium.

Proof. Consider an equilibrium (or a subgame perfect equilibrium) of $\Gamma(R, \tau, x \mid 1)$. Player i can always choose to reject any proposal at every history, and so receive $(1-\delta)\gamma\Delta\big[(1-\tau)r_i + \tau x_i \bar{r}\big]$ in every period. Hence, her equilibrium payoff must be at least $\gamma\Delta\big[(1-\tau)r_i + \tau x_i \bar{r}\big]$. As this is true for all players and the payoff vector $\big(\gamma\Delta\big[(1-\tau)r_j + \tau x_j \bar{r}\big]\big)_{j \in N}$ belongs to the Pareto frontier, $\gamma\Delta\big[(1-\tau)r_i + \tau x_i \bar{r}\big]$ is also i's maximum equilibrium payoff.

Consider a game of the form $\Gamma(a, \tau', y \mid 1)$, with $a \in \{R, S\}$, $\tau' \in [0, 1]$ and $y \in X$; and let $\widetilde{\Gamma}(a, \tau', y \mid 1)$ be a variant on this game where, as in standard model of Baron and Ferejohn (1989), the first successful proposal is never amended. Routine arguments show that $\widetilde{\Gamma}(a, \tau', y \mid 1)$ has a stationary equilibrium in which: the first period's last proposer successfully offers a Pareto efficient policy if the status quo (a, τ', y) is not itself efficient; and (a, τ', y) is never amended otherwise — see Baron and Ferejohn (1989) for more details. It is easy to see that a strategy profile for $\Gamma(R, \tau, x \mid 1)$ (or $\Gamma(S, \tau, x \mid 1)$) that prescribes the same actions as that stationary equilibrium in the first period of each continuation game $\Gamma(a, \tau', y \mid 1)$ is an equilibrium. Moreover, as the policy implemented in each period is Pareto efficient, the equilibrium thus obtained is renegotiation-proof. □

Lemma C2. *Suppose $q = n$ and $\tau^{\max} = 1$. Then, there exists $\widehat{\Delta}_1 > 0$ such that the following holds for all $\Delta < \widehat{\Delta}_1$, $\tau \in [0,1]$, $x \in X$ and $\alpha_k \leq \alpha^*$:*

(i) For every coalition $C \subset N$ comprising $n-1$ legislators, $\Gamma(R,\tau,x \mid \alpha^*)$ has a renegotiation-proof equilibrium, in which the payoff to each legislator $i \in C$ is equal to $P_i(\tau,x|\alpha^*) \equiv \gamma\Delta\big[(1-\tau)r_i + \tau x_i \bar{r}\big]\alpha^*$; and

(ii) the set of renegotiation-proof equilibrium payoff vectors for the game $\Gamma(R,\tau,x \mid \alpha^*)$ is $\big\{(w_1,\ldots,w_n) \in \mathbb{R}^n_+ \colon \sum_{i=1}^n w_i = \bar{s}\Delta$ and $w_i \geq P_i(\tau,x|\alpha^*), \forall i \in N\big\}$.

Proof. We begin with the definition of the threshold $\widehat{\Delta}_1$. An application of l'Hôpital's rule gives

$$\lim_{\Delta\to 0} \alpha^* = \frac{\rho\bar{s}}{\gamma\big[(\rho+\gamma)\bar{r} - \bar{s}\big]} \ ,$$

so that

$$\lim_{\Delta\to 0} \Big[\bar{s} - \frac{1}{\Delta}\sum_{i\in N} P_i(\tau,x|\alpha^*)\Big] = \bar{s} - \gamma\bar{r}\lim_{\Delta\to 0}\alpha^* = \frac{\bar{s}(\gamma\bar{r}-\bar{s})}{\gamma\big[(\rho+\gamma)\bar{r}-\bar{s}\big]} > 0 \ ,$$

for all $\tau \in [0,1]$ and $x \in X$. Therefore, there exists $\widehat{\Delta}_1 > 0$ such that, whenever $\Delta < \widehat{\Delta}_1$,

$$\delta(1-\alpha^*\gamma\Delta)(n-2)\Big[\bar{s} - \frac{1}{\Delta}\sum_{i\in N} P_i(\tau,x|\alpha^*)\Big] > \big[1-\delta(1-\alpha^*\gamma\Delta)\big]\bar{s} \ ,$$

for all $\tau \in [0,1]$ and $x \in X$. Henceforth, we assume that $\Delta < \widehat{\Delta}_1$.

Part (i). The proof is constructive. We begin with an intuitive description of the equilibrium. The safe alternative S is implemented in each period (both on and off the path). As the belief that R is good is less than or equal to α^*, this implies that payoff vectors are Pareto optimal and, therefore, that the putative equilibrium is renegotiation-proof. Once S has been implemented, all proposers pass in all future periods, irrespective of the tax rate and distribution of revenues. If S has not yet been implemented, then behavior is determined by a set of n "phases," each corresponding to one legislator in N. If the status quo is of the form (R,τ',y) and the belief is $\alpha_k \leq \alpha^*$ then, in phase i, every proposer successfully offers a policy that gives a payoff of $R_i(\tau',y|\alpha_k) \equiv \bar{s}\Delta - \sum_{j\neq i} P_j(\tau',y|\alpha_k)$ to legislator i and $P_j(\tau',y|\alpha_k)$ to each legislator $j \neq i$. The idea is that i receives her "reward payoff" and the others their "punishment payoffs." If a proposer, say i, deviates from, then every legislator (other than i) rejects her proposal and the game transitions to phase k, where k is the identity of the first legislator who rejected the proposal. If voter i rejects a proposal which she should have accepted, then the game moves to phase i.

We now turn to the formal definition of the equilibrium, σ^*, for $\Gamma(R, \tau, x \mid \alpha^*)$, with $\tau \in [0,1]$ and $x \in X$. Let $C \subset N$ be a coalition comprising $n-1$ legislators. The game begins in phase $\ddot{\imath} \notin C$. Then, σ^* prescribes the following behavior:

• **Proposal stages with uncertainty.** Consider first proposer i's behavior in a period where the order of proposers is $\pi = (\pi_1, \ldots, \pi_n)$ with $\pi_\ell = i$ for some ℓ, the first $\ell - 1$ proposers have failed to amend the status quo, the belief is $\alpha_k \leq \alpha^*$, and the game is in phase $j \in N$. There are two cases:

Case P1: The status quo is (R, τ', y), where $\tau' \in [0,1]$ and $y \in X$.

Proposer i's strategy prescribes her to offer $(S, 1, x^j)$, where $x^j \in X$ is defined by: $x^j_\ell \equiv P_\ell(\tau', y|\alpha_k)/(\bar{s}\Delta)$ for all $\ell \neq j$. The game remains in phase j, irrespective of her move.

Case P2: The status quo is (S, τ', y), where $\tau' \in [0,1]$ and $y \in X$.

Proposer i's strategy prescribes her to pass — i.e., to offer (S, τ', y). The game remains in phase j, irrespective of her move.

• **Voting stages with uncertainty.** Consider now voter i's behavior in a period where the order of proposers is $\pi = (\pi_1, \ldots, \pi_n)$, the belief is $\alpha_k \leq \alpha^*$, and the game is in phase $j \in N$. There are several cases:

Case V1: The status quo is (R, τ', y), where $\tau' \in [0,1]$ and $y \in X$, and policy $(S, 1, x^j)$ has just been proposed.

Legislator i's strategy prescribes her to accept this proposal. If all voters accept the proposal,then the game remains in phase j; otherwise, it transitions to phase $\hat{\jmath} + 1$ where $\hat{\jmath}$ is the first voter who rejected the proposal. (We set $\hat{\jmath} + 1 = 1$ if $\hat{\jmath} = n$.)

Case V2: The status quo is (R, τ', y), where $\tau' \in [0,1]$ and $y \in X$, and a proposer π_ℓ has just offered policy $p \neq (S, 1, x^j)$.

Case V2.1: Proposal p is of the form (S, τ'', z) for some $\tau'' \in [0,1]$ and $z \in X$.

The strategy of legislator $i \neq \pi_\ell$ prescribes her to accept this proposal if and only if

$$\big[1 - \delta(1-\gamma\Delta)\big] w_i(R, \tau', y \mid \alpha_k) + \delta(1 - \alpha_k\gamma\Delta) R_i(\tau', y|\alpha_{k+1}) < w_i(S, \tau'', z \mid \alpha_k);$$

and the strategy of legislator π_ℓ prescribes her to accept this proposal if and only if

$$\big[1 - \delta(1-\gamma\Delta)\big] w_i(R, \tau', y \mid \alpha_k) + \delta(1 - \alpha_k\gamma\Delta) P_i(\tau', y|\alpha_{k+1}) < w_i(S, \tau'', z \mid \alpha_k).$$

If (S, τ'', z) is rejected by some $i \neq \pi_\ell$, then the game transitions to phase $\hat{\imath}$, where $\hat{\imath}$ is the first voter (other than π_ℓ) who rejected it. Otherwise, the game transitions to phase $\hat{\jmath}$, where $\hat{\jmath} \equiv \min\{i \in N : i \neq \pi_\ell\}$.

Case V2.2: Proposal p is of the form (R, τ'', z) for some $\tau'' \in [0,1]$ and $z \in X$.

The strategy of legislator $i \neq \pi_\ell$ prescribes her to accept this proposal if and only if

$$\begin{aligned} &\big[1 - \delta(1-\gamma\Delta)\big] w_i(R, \tau', y \mid \alpha_k) + \delta(1 - \alpha_k\gamma\Delta) R_i(\tau', y|\alpha_{k+1}) \\ &< \big[1 - \delta(1-\gamma\Delta)\big] w_i(R, \tau'', z \mid \alpha_k) + \delta(1 - \alpha_k\gamma\Delta) \begin{cases} R_i(\tau'', z|\alpha_{k+1}) & \text{if } i = \hat{\jmath}\ , \\ P_i(\tau'', z|\alpha_{k+1}) & \text{if } i \neq \hat{\jmath}\ , \end{cases} \end{aligned}$$

where $\hat{\jmath}$ is defined as in V2.1; and the strategy of legislator π_ℓ prescribes her to accept this proposal if and only if

$$\big[1 - \delta(1-\gamma\Delta)\big] w_i(R, \tau', y \mid \alpha_k) + \delta(1 - \alpha_k\gamma\Delta) P_i(\tau', y|\alpha_{k+1}) < P_i(\tau'', z|\alpha_k)\ .$$

If (R, τ'', z) is rejected by some $i \neq \pi_\ell$, then the game transitions to phase $\hat{\imath}$, where $\hat{\imath}$ is the first voter (other than π_ℓ) who rejected it. Otherwise, the game transitions to phase $\hat{\jmath}$.

Case V3: The status quo is (S, τ', y), where $\tau' \in [0,1]$ and $y \in X$, and a policy $p \neq (S, \tau', y)$ has just been proposed.

Case V3.1: Proposal p is of the form (S, τ'', z) for some $\tau'' \in [0,1]$ and $z \in X$.

Legislator i's strategy prescribes her to accept this proposal if and only if $(1-\tau'')s_i + \tau'' z_i \bar{s} > (1-\tau')s_i + \tau' y_i \bar{s}$. The game remains in the same phase, irrespective of the legislators' voting behavior.

Case V3.2: Proposal p is of the form (R, τ'', z) for some $\tau'' \in [0,1]$ and $z \in X$.

Legislator i's strategy prescribes her to accept this proposal if and only if

$$w_i(S, \tau', y \mid \alpha_k) < \big[1 - \delta(1-\gamma\Delta)\big] w_i(R, \tau'', z \mid \alpha_k) + \delta(1-\alpha_k\gamma\Delta)x_i^j \bar{s}.$$

The game remains in the same phase, irrespective of the legislators' voting behavior.

• **Proposal/voting stages without uncertainty.** In any continuation game where the belief is equal to one, σ^* prescribes the same actions as one of the equilibria obtained in Lemma C1.

To see that σ^* is an equilibrium, consider first legislator i's behavior in case V1. If another legislator has already rejected proposal $(S, 1, x^j)$, then her action has no impact on her payoff and is, therefore, trivially optimal. Suppose that all previous voters have accepted the proposal. If she also accepts (as prescribed by σ^*), then she will receive a payoff of

$$x_i^j \bar{s}\Delta = \begin{cases} R_i(\tau', y|\alpha_k) & \text{if } i = j \ , \\ P_i(\tau', y|\alpha_k) & \text{otherwise;} \end{cases}$$

if she rejects, then she will receive

$$\big[1 - \delta(1-\gamma\Delta)\big] w_i(R, \tau' y \mid \alpha_k) + \delta(1-\alpha_k\gamma\Delta)P_i(\tau', y|\alpha_{k+1}) = P_i(\tau', y|\alpha_k) \ .$$

As $R_i(\tau', y|\alpha_k) - P_i(\tau', y|\alpha_k) = \big[\bar{s} - (1/\Delta)\sum_{j\in N} P_j(\tau', y|\alpha_k)\big]\Delta > 0$ when $\Delta < \widehat{\Delta}_1$, deviating from σ^* is unprofitable for i in this case.

Consider now case V2.1. By the same argument as above, player i cannot profitably deviate if another legislator, other than the proposer π_ℓ, has already rejected the proposal. Suppose that, when it is legislator i's turn to vote, all previous voters have accepted $p \neq (S, 1, x^j)$. If i also accepts p, then it will never be amended (see Case P2) and i's payoff will therefore be $w_i(S, \tau'', z \mid \alpha_k)$. If i rejects p, then (R, τ', y) is implemented in the current period and, if R is a success, will never be amended. If R is a failure (which happens with probability $(1-\alpha_k\gamma\Delta)$) and $i \neq \pi_\ell$ [resp. $i = \pi_\ell$], then the game moves to phase i [resp. $\hat{\jmath}$]. Hence, legislator i's payoff from rejecting p is

$$\big[1-\delta(1-\gamma\Delta)\big]w_i(R,\tau',y \mid \alpha_k) + \delta(1-\alpha_k\gamma\Delta)\begin{cases} R_i(\tau',y|\alpha_{k+1}) & \text{if } i \neq \pi_\ell \ , \\ P_i(\tau',y|\alpha_{k+1}) & \text{otherwise.}\end{cases}$$

Hence, she cannot profitably deviate from σ^*. Note for future reference that, as $q = n$, a necessary condition for the proposal to be accepted is

$$\big[1-\delta(1-\gamma\Delta)\big]\alpha_k\gamma\Delta\bar{r} + \delta(1-\alpha_k\gamma\Delta)\Big[P_{\pi_\ell}(\tau',y|\alpha_{k+1}) + \sum_{i\neq\pi_\ell} R_i(\tau',y|\alpha_{k+1})\Big] < \bar{s}\Delta$$

or, equivalently,

$$\big[1-\delta(1-\gamma\Delta)\big]\alpha_k\gamma\bar{r} + \delta(1-\alpha_k\gamma\Delta)(n-2)\big[\bar{s} - \frac{1}{\Delta}\sum_{i\in N} P_i(\tau',y|\alpha_{k+1})\big] < \big[1-\delta(1-\alpha_k\gamma\Delta)\big]\bar{s} \ .$$

As $\Delta < \widehat{\Delta}_1$ and $\big[1-\delta(1-\gamma\Delta)\big]\alpha_k\gamma\bar{r} > 0$, this inequality never holds and the proposal is consequently rejected. By the same logic, there are no profitable deviations and the proposal is always rejected in case V2.2.

In case V.3, σ^* prescribes each legislator to vote for the policy that would maximize her continuation value (in the vent that she is still decisive when she has to vote). Deviations are therefore unprofitable. Moreover, as the status quo payoffs belong to the Pareto frontier (and $q = n$), all proposals must be unsuccessful.

It follows from the analysis of cases V1 and V.2 that, in case P1, $(S, 1, x^j)$ is the only proposal that can successfully be made. If proposer i makes this proposal, then she either receives $R_i(\tau',y|\alpha_k)$ or $P_i(\tau',y|\alpha_k)$ — the former if $i = j$, and the latter otherwise. If she deviates by making any other proposal, then she receives $P_i(\tau',y|\alpha_k)$. As $P_i(\tau',y|\alpha_k) < R_i(\tau',y|\alpha_k)$, deviations are unprofitable. Finally, in case P2, any proposal would be rejected and, as the status quo is absorbing, proposers cannot improve on passing.

We have thus constructed an equilibrium for $\Gamma(R,\tau,x \mid \alpha^*)$, in which player $\ddot{\imath} \notin C$ receives $R_{\ddot{\imath}}(\tau,x|\alpha^*)$ and each $i \in C$ receives $P_i(\tau,x|\alpha^*)$.

Part (ii). If legislator i rejects every proposal to amend the status quo in $\Gamma(R,\tau,x \mid \alpha^*)$, then she receives a payoff of $P_i(\tau,x|\alpha^*)$. Therefore, her payoff must be at least $P_i(\tau,x|\alpha^*)$

in any equilibrium. Moreover, it is easy to see that one can use a similar construction to that used in part (i) to obtain any equilibrium vector in $\{(w_1,\ldots,w_n)\in\mathbb{R}^n_+ :\ \sum_{i=1}^n w_i = \bar{s}\Delta \text{ and } w_i \geq P_i(\tau,x|\alpha^*), \forall i\in N\}$: if every legislator i is willing to implement a policy that yields her "punishment payoff" $P_i(\tau,x|\alpha^*)$ when $\Delta<\widehat{\Delta}_1$, then she is also willing to implement a policy that gives her any higher payoff. This implies that $\{(w_1,\ldots,w_n)\in\mathbb{R}^n_+ :\ \sum_{i=1}^n w_i = \bar{s}\Delta \text{ and } w_i \geq P_i(\tau,x|\alpha^*), \forall i\in N\}$ is the Pareto frontier of the set of equilibrium payoff vectors and, consequently, the set of renegotiation-proof equilibrium payoff vectors for $\Gamma(R,\tau,x\mid\alpha^*)$. □

Lemma C3. *Suppose $q=n$ and $\tau^{\max}=1$. Then, there exists $\widehat{\Delta}_1>0$ such that, for all $\Delta<\widehat{\Delta}_1$, $\tau\in[0,1]$ and $x\in X$, the set of renegotiation-proof equilibrium payoff vectors of $\Gamma(R,\tau,x\mid\alpha_1)$ is $\{(w_1,\ldots,w_n)\in\mathbb{R}^n_+ :\ \sum_{i=1}^n w_i = V^*(\alpha_1) \text{ and } w_i\geq P_i(\tau,x|\alpha_1),\forall i\in N\}$.*

Proof. Let $k^*\in\mathbb{N}$ be implicitly defined by $\alpha_{k^*}\equiv\alpha^*$. We begin by characterizing the set of renegotiation-proof equilibrium payoff vectors of $\Gamma(R,\tau,x\mid\alpha_{k^*-1})$. Given the (renegotiation-proof) equilibrium continuation values obtained in Lemmas C1 and C2, it is routine to show that $\Gamma(S,\tau,x\mid\alpha_{k^*-1})$ possesses a renegotiation-proof equilibrium. Pick an arbitrary payoff vector $w\in\{(w_1,\ldots,w_n)\in\mathbb{R}^n_+ :\ \sum_{i=1}^n w_i = V^*(\alpha_{k^*-1}) \text{ and } w_i\geq P_i(\tau,x|\alpha_{k^*-1}),\forall i\in N\}$, and let $y\in X$ be defined by: $y_i\equiv w_i/V^*(\alpha_{k^*-1})$ for each $i\in N$. As $\Delta<\widehat{\Delta}_1$, we can use the equilibrium characterization of Lemma C2 to construct a renegotiation-proof equilibrium σ for $\Gamma(R,\tau,x\mid\alpha_{k^*-1})$ that supports w. The construction is in the same vein as that used in the proof and we only provide an intuitive description of σ:

• The strategy prescribes the first period's proposer π_1 to offer policy $(R,1,y)$ and all legislators to accept this proposal. If any other proposer is called upon to make a proposal (off the path), then she passes.

• For each $i\in N$, let σ^i be a renegotiation-proof equilibrium of $\Gamma(R,\tau,x\mid\alpha^*)$ that gives legislator i a payoff of $P_i(\tau,x|\alpha^*)$ — we know from Lemma C2 that such an equilibrium exists. If π_1 proposes some policy $p\neq(R,1,y)$, then the proposal is rejected, the other proposers pass and equilibrium σ^1 is played in the continuation game $\Gamma(R,\tau,x\mid\alpha^*)$. The reason why p is rejected in equilibrium is that the payoff vector from the policy sequence that consists of (R, τ, x) followed by the policies induces by σ^1 is in the Pareto frontier. This guarantees that at least one legislator is better-off under this policy sequence than under the policy sequence induced by accepting p. By construction, $w_{\pi_1} \geq P_{\pi_1}(\tau, x|\alpha_{k^*-1}) = \big[1-\delta(1-\gamma\Delta)\big]w_{\pi_1}(R, \tau, x \mid \alpha_{k^*-1}) + \delta(1-\alpha_{k^*-1}\gamma\Delta)P_{\pi_1}(\tau, x|\alpha^*)$, so that π_1 is better off successfully proposing $(R, 1, y)$ than proposing p. A similar argument shows why passing is optimal for each proposer π_ℓ, $\ell \geq 2$.

• The proposal $(R, 1, y)$ is unanimously accepted for the following reason. If it is accepted, then σ prescribes to play the renegotiation-proof equilibrium that gives legislator i a payoff of $y_i\bar{s} = y_i V^*(\alpha^*) = w_i$ in the continuation game $\Gamma(R, \tau, x \mid \alpha^*)$ (Lemma C2); if it is rejected, then σ prescribes to play σ^j in $\Gamma(R, \tau, x \mid \alpha^*)$, where j is the first legislator who rejected $(R, 1, y)$. Given that the other legislators accept $(R, 1, y)$, player i would thus receive $\big[1-\delta(1-\gamma\Delta)\big]w_i(R, \tau, x \mid \alpha_{k^*-1})+\delta(1-\alpha_{k^*-1}\gamma\Delta)P_i(\tau, x|\alpha^*) = P_i(\tau, x|\alpha_{k^*-1}) \leq w_i$ if she rejected it.

As w was chosen arbitrarily from $\big\{(w_1, \dots, w_n) \in \mathbb{R}^n_+ \colon \sum_{i=1}^n w_i = V^*(\alpha_{k^*-1})$ and $w_i \geq P_i(\tau, x|\alpha_{k^*-1}), \forall i \in N\big\}$, this proves that the latter set is the set of renegotiation-proof equilibrium payoff vectors of $\Gamma(R, \tau, x \mid \alpha_{k^*-1})$. Applying the same logic recursively from $k = k^* - 2$ to $k = 1$, we obtain the lemma. $\square$

By assumption, $\rho\bar{s} < \alpha_0\gamma\big[(\rho+\gamma)\bar{r} - \bar{s}\big]$. This implies that there exists $\widehat{\Delta}_2 > 0$ such that $\delta(n-2)\big[V^*(\alpha_0) - \bar{s}\Delta\big] > (1-\delta)V^*(\alpha_0)$ for all $\Delta < \widehat{\Delta}_2$. Henceforth, we assume that $\Delta < \min\big\{\widehat{\Delta}_1, \widehat{\Delta}_2\big\}$. To complete the proof of Proposition 3, therefore, it suffices to show that, for every vector w in the simplex $W^* \equiv \big\{(w_1, \dots, w_n) \in \mathbb{R}^n_+ \colon \sum_{i=1}^n w_i = V^*(\alpha_0)$ and $w_i \geq \Delta s_i, \forall i \in N\big\}$, there is a renegotiation-proof equilibrium that supports w. In one such equilibrium, policy $(R, 1, x)$, where $x_i \equiv w_i/V^*(\alpha_0)$ is implemented in every period unless the belief becomes equal to α^*, in which case policy $(S, 1, x)$ is implemented in all future periods. The construction of this equilibrium parallels closely that used in the proof of Lemma C2, and will be omitted. The two main differences are that: (i) it focuses on what happens at the initial belief α_0, using the equilibrium continuation values obtained in Lemmas C1 and C3 to describe behavior at other beliefs; and (ii) the key condition ensuring that all proposals $p \neq (R, 1, x)$ in the first period are unsuccessful now relies on $\delta(n-2)\big[V^*(\alpha_0) - \bar{s}\Delta\big] > (1-\delta)\bar{s}\Delta$. This is because, in equilibrium, the first legislator i (different from the proposer) who rejects p is promised a continuation value of $V^*(\alpha_0) - \sum_{j\neq i} s_j\Delta$, while the proposer $\hat{\imath}$ is promised $s_{\hat{\imath}}\Delta$. To be accepted, p must therefore give a payoff greater than or equal to $(1-\delta)\Delta s_i + \delta\big[V^*(\alpha_0) - \sum_{j\neq i} s_j\Delta\big]$ to each legislator $i \neq \hat{\imath}$, and a payoff greater than or equal to $s_{\hat{\imath}}\Delta$ to legislator $\hat{\imath}$. Summing across the legislators, feasibility requires

$$(1-\delta)\bar{s}\Delta + \delta\Big[s_{\hat{\imath}} + (n-1)V^*(\alpha_0) - \sum_{i\neq\hat{\imath}}\sum_{j\neq i} s_j\Delta\Big] \leq V^*(\alpha_0)$$

or, equivalently,

$$(1-\delta)\bar{s}\Delta + \delta(n-2)\big[V^*(\alpha_0) - \bar{s}\Delta\big] \leq (1-\delta)V^*(\alpha_0)\ .$$

As $(1-\delta)\bar{s}\Delta > 0$, this is impossible if $\delta(n-2)\big[V^*(\alpha_0) - \bar{s}\Delta\big] > (1-\delta)V^*(\alpha_0)$ — which holds since $\Delta < \widehat{\Delta}$.

References

Acemoglu, D., Golosoc, M. and Tsyvinski, A. (2008), 'Political Economy of Mechansims', Econometrica 76(3), 619-641.

Acemoglu, D., Naidu, S., Restrepo, P. and Robinson, J. A. (2015), Democracy, Redistribution, and Inequality, Vol. 2B, North Holland, chapter 21.

Anesi, V. and Seidmann, D. J. (2015), 'Bargaining in Standing Committees with an Endogenous Default', The Review of Economic Studies 82(3), 825-867.

Baron, D. P. and J.A. Ferejohn (1989), 'Bargaining in Legislatures', The American Political Science Review 83, 1181-1206.

Baron, D. P. (1996), 'A Dynamic Theory of Collective Goods Programs', The American Political Science Review 90(2), 316-330.

Battaglini, M. and Palfrey, T. R. (2012), 'The Dynamics of Distributive Politics', Economic Theory 49(3), 739-777.

Bergin, J. and MacLeod, W. B. (1993), 'Continuous Time Repeated Games', International Economic Review 34(1), 21-37.

Besley, T. and Coate, S. (1998), 'Sources of Inefficiency in a Representative Democracy: A Dynamic Analysis', American Economic Review 88(1), 139-156.

Bowen, T. R. and Baron, D. (2014), 'Dynamic Coalitions', working paper .

Bowen, T. R. and Zahran, Z. (2012), 'On Dynamic Compromise', Games and Economic Behavior 76, 391-419.

Bouton, L., Llorente-Sauer, A. and Malherbe, F. (forthcoming), 'Get Rid of Unanimity: The Superiority of Majority Rules with Veto Power ', Journal of Political Economy.

Cai, H. and Treisman, D. (2009), 'Political Decentralization and Policy Experimentation', Quarterly Journal of Political Science 4(1), 35-58.

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

Callander, S. and Hummel, P. (2014), 'Preemtive Policy Experimentation', Econometrica 82(4), 1509-1528.

Dewatripont, M. and Roland, G. R. (1992), 'Economic Reform and Dynamic Political Constraints', Review of Economic Studies 59, 703-730.

Diermeier, D., Egorov, G. and Sonin, K. (2015), 'Political Economy of Redistribution', Econometrica 85 (3), 851-870.

Dziuda, W. and Loeper, A. (2016), 'Dynamic Collective Choice with Endogenous Status Quo', Journal of Political Economy 124(4).

Dziuda, W. and Loeper, A. (2017), 'Dynamic Pivotal Politics', working paper.

Fernandez, R. and Rodrik, D. (1991), 'Resistance to Reform: Status Quo Bias in the Presence of Individual-Specific Uncertainty', American Economic Review 81(5), 1146-1155.

Hörner, J. and Samuelson, L. (2013), 'Incentives for experimenting agents', Rand Journal of Economics 44(4), 632-663.

Kalandrakis, T. (2004), 'A Three-Player Dynamic Majoritarian Bargaining Game', Journal of Economic Theory 116(2), 294-322.

Kalandrakis, T. (2010), 'Minimum winning coalitions and endogenous status quo', International Journal of Game Theory 39, 617-643.

Keller, G., Rady, S. and Cripps, M. (2005), 'Strategic Experimentation with Exponential Bandits', Econometrica 73(1), 39-68.

Majumdar, S. and Mukand, S. W. (2004), 'Policy Gambles', The American Economc Review 94(4), 1207-1222.

Meltzer, A. H. and Richard, S. F. (1981), ‘A Rational Theory of the Size of Government’, Journal of Political Economy 89(5), 914-927.

Messner, M. and Polborn, M. K. (2012), ‘The option to wait in collective decisions and optimal majority rules’, Journal of Public Economics 96(5-6), 524-540.

Murto, P. and Välimäki, J. (2011), ‘Learning and Information Aggregation in an Exit Game’, Review of Economic Studies 78, 1426-1461.

Nunnari, S. (2014), ‘Dynamic Legislative Bargaining with Veto Power: Theory and Experiments’, Working Paper .

Richter, M. (2014), ‘Fully absorbing dynamic compromise’, Journal of Economic Theory 152, 92-104.

Strulovici, B. (2010), ‘Learning while Voting: Determinants of Collective Experimentation’, Econometrica 78(3), 933-971.

Tornell, A. (1998), ‘Reform from Within’, NBER Working Paper No. 6497 .

Volden, C., Ting, M. M. and Carpenter, D. P. (2008), ‘A Formal Model of Learning and Policy Diffusion’, The American Political Science Review 102(3), 319-332.

Zapal, J. (2016), ‘Markovian Equilibria in Dynamic Spatial Legislative Bargaining: Existence with Three Players’, Games and Economic Behavior 98, 235-242.