Strategic Renegotiation in Repeated Games

Mikhail Safronov
University of Cambridge

Bruno Strulovici*
Northwestern University

September 17, 2016

Abstract

Cooperative approaches to study renegotiation in repeated games have assumed that Pareto-ranked equilibria could not coexist within the same renegotiation-proof set. With strategic renegotiation, however, a proposal to move to a Pareto-superior equilibrium can be deterred by a different continuation equilibrium which harms the proposer and rewards the rejector. This paper studies strategic renegotiation in repeated games, defining renegotiation-proof outcomes by a simple equilibrium refinement. We provide distinct necessary and sufficient conditions for renegotiation-proofness which converge to each other as renegotiation frictions become negligible and which are straightforward to characterize graphically. The analysis suggests a novel mechanism to explain the persistence of equilibrium inefficiencies, such as miscoordination and status quo bias, even when information is complete, communication is frictionless, and players are arbitrarily patient and can credibly agree on efficient outcomes.

1 Introduction

The punishment equilibria used to sustain cooperation in repeated games are often Pareto inefficient. This puts into question their viability and, hence, the viability of cooperative outcomes based on such punishments, when players are free to renegotiate the continuation of the game. Incorporating renegotiation satisfactorily in repeated games has been a longstanding challenge.

To address this question, economists have introduced various concepts of renegotiation-proofness based on the following idea: roughly speaking, an equilibrium is *not* renegotiation-proof if it entails a continuation play that is Pareto dominated by some "credible" equilibrium (Pearce (1987), Bernheim and Ray (1989), Farrell and Maskin (1989), Abreu and Pearce (1991), and Asheim (1991)).[1]

*We are grateful for comments from Larry Samuelson, Ilya Segal, Tadashi Sekiguchi, Takuo Sugaya, Joel Watson, and seminar participants at the 2015 SAET meeting, the 2015 World Congress of the Econometric Society, and the 2016 Asian Meeting of the Econometric Society. Strulovici gratefully acknowledges financial support from an NSF CAREER Award (Grant No. 1151410) and a fellowship from the Alfred P. Sloan Foundation

[1]The first discussion along these lines is due to Farrell (1983), which is subsumed by Farrell and Maskin (1989). Other approaches to renegotiation include DeMarzo (1988), Benoît and Krishna (1993), and Bergin and MacLeod (1993). All these papers follow axiomatic approaches.

These concepts mainly differ regarding what "credible" means and yield contrasted results: while Pearce (1987) argued,[2] as in the first paragraph, that maximal cooperation may not be sustained due to the lack of a credible and severe enough punishment, Farrell and Maskin (1989) found that most renegotiation-proof outcomes, as players become arbitrarily patient, had to be Pareto efficient.[3]

Owing to their cooperative (i.e., non-strategic) nature, these concepts have left unexplored an aspect of renegotiation which arises naturally when one considers an explicit protocol of renegotiation: what happens when a player rejects another player's proposal? Suppose that during the punishment phase of a two-player repeated game, the continuation payoffs are (X_1, X_2) and player 1 proposes a Pareto-improving equilibrium with payoffs (Y_1, Y_2). Clearly, such a Pareto-improvement need not be accepted if, by rejecting 1's proposal, player 2 gets rewarded by a higher continuation payoff $Z_2 > Y_2$. Moreover, if 1's continuation payoff Z_1 after 2 has rejected the offer is less than X_1, then it is suboptimal for 1 to propose the Pareto improvement in the first place. With strategic renegotiation, a Pareto dominated equilibrium may thus withstand renegotiation as long as any off-path proposal may be deterred in this fashion. Punishing a player who deviates (here, in proposals) and rewarding other players is standard in repeated game analysis. It also seems plausible: for example, if an agent tries to bribe another one to obtain some advantage (a Pareto improvement for these agents, other things equal), the agent who rejects and exposes the bribe may be rewarded and the corruptor punished as a result.

This paper models strategic renegotiation in repeated games by adding a simple stage at the end of each period: after actions and payoffs have been chosen and observed in period t, one of the players may be selected with fixed probability to propose a continuation *plan*. A plan for period $t+1$ is described recursively as follows: it prescribes players' actions, proposals, and acceptance decisions in period $t+1$, as well as the continuation plan for period $t+2$ as a function of the actions, proposals, and acceptance decisions observed in period $t+1$.

To give traction to renegotiation, we introduce a simple equilibrium refinement. An equilibrium of the enlarged game is *renegotiation-proof* if, as long as no off-path proposal has been accepted, any accepted proposal is played. Intuitively, this refinement may be viewed as a social norm under which accepted proposals are taken seriously until, possibly, players abandon this norm by accepting

[2] See also Abreu and Pearce (1991) and Abreu, Pearce and Stacchetti (1993).

[3] Farrell and Maskin, like Bernheim and Ray, introduce weak and strong concepts of renegotiation-proofness. The strong notion is arguably the more satisfactory one as it allows external comparisons: For example, the infinite repetition of any fixed Nash equilibrium of the stage game always constitutes a weakly renegotiation-proof equilibrium, but typically fails to be strongly-renegotiation proof. The strong concept is particularly demanding, as a strong renegotiation-proof equilibrium must win payoff comparisons even against those equilibria which are themselves eliminated by that concept.

some off-path proposal.[4]

Our main object of study is the set of payoffs generated by renegotiation-proof equilibria. Starting with two players, we characterize this set when players become arbitrarily patient and renegotiation frictions—modeled as the probability that no one gets to make a proposal within any fixed time window—become negligible. This set is nonempty and straightforward to describe graphically: it is the intersection of the set of individually rational payoffs (as in the Folk Theorem) and of the positive orthant whose boundaries go through the endpoints of the Pareto frontier. More generally, we compute distinct necessary and sufficient conditions for renegotiation-proofness which are strictly nested for any finite level of renegotiation friction and converge to each other as renegotiation frictions become negligible.

Our construction implies *path dependence* for the set of proposals which are acceptable in equilibrium. For example, the cooperative proposal (Y_1, Y_2) mentioned above may be acceptable at the beginning of the game, but not after a deviation. Several authors have expressed discomfort with static notions of renegotiation-proofness, precisely because they ruled out path dependence (Abreu and Pearce (1991) and Asheim (1991)). Path dependence arises naturally when renegotiation is considered to be part of the equilibrium of a larger game, rather than a restriction on the set of equilibria of the underlying repeated game.[5]

The analysis provides a novel explanation for the existence of inefficient equilibria, such as coordination failures and political inertia, when agents are i) arbitrarily patient, ii) freely able to communicate, iii) and able to credibly agree on future strategies. As noted, strategic negotiation requires that continuation play be prescribed not only after off-path actions but also after off-path proposals, such as suggestions to move to a Pareto-improving equilibrium. In the context of an oppressive regime, for instance, the equilibrium may specify that any "subversive" (i.e., regime-threatening) proposal triggers a punishment for the proposer and rewards the players who reject the proposal.

With three or more players, new conceptual issues emerge. In particular, one must specify the continuation equilibrium when only a subset of players accept the proposal made by a given player.[6] We explore several specifications, whose predictions for the set of renegotiation-proof

[4]We also consider a milder refinement requiring only that "credible" accepted proposals be played, where credibility is defined in Section 5. With two players, the refinements turn out to yield the same necessary and sufficient conditions.

[5]It also suggests that Farrell and Maskin's and Bernheim and Ray's concepts of weak renegotiation proofness and internal consistency may rule out reasonable renegotiation-proof equilibria by preventing path dependence of acceptable continuation equilibria. In particular, while those concepts require that no Pareto ranked equilibria coexist within a given norm, this paper suggests that Pareto-dominated equilibria may withstand renegotiation as long as the social norm specifies clear punishments and rewards for the proposers and rejectors of Pareto improving equilibria.

[6]Other issues arise when a player makes a proposal to a subset of players. While such a proposal is formally

payoffs range from the Folk Theorem to Pareto efficiency. The simplest one requires that the continuation equilibrium when a proposal fails to receive unanimous approval be independent of the rejectors' identities. In this case, our necessary and sufficient conditions for renegotiation-proof payoffs both take the form of upper-orthants, a useful qualitative property to model renegotiation in repeated games. For each assumption that we considered, the sets characterizing necessary and sufficient conditions become arbitrarily close to each other as renegotiation frictions vanish. The analysis of this more general environment is contained in Section 7.

Several papers have studied negotiation by players who engage in several rounds of cheap talk before choosing their actions in a one-shot game, and asked whether this pre-play communication could help select efficient equilibrium in the one-shot game. Farrell (1987) considers an entry game in which firms simultaneously announce their intention of whether to enter the market. With pre-play communication, firms achieve a higher payoff than they do in the symmetric one-shot equilibrium, but do not achieve perfect coordination. In Rabin (1994), players simultaneously propose Nash equilibria of the one-shot game and an equilibrium is played if both players propose it. With sufficiently many rounds of communication, each player is guaranteed to get at least her worst payoff in the "Pareto meet," which is the set of Pareto-efficient Nash equilibria in the one-shot game. Players need not achieve a Pareto-efficient outcome, however. In inefficient equilibria, players keep proposing their preferred outcome with high probability and may thus fail to reach an agreement.

The papers closest to ours are Santos (2000) and Miller and Watson (2013). In the alternative-offer model studied by Santos (2000), players bargain over Nash equilibria of a one-shot game, and play whichever equilibrium is agree upon. In that paper, each player is guaranteed to get a payoff in the Pareto meet, but players may still end up playing a Pareto-inefficient equilibrium. More recently, Miller and Watson (2013) study equilibrium selection in a repeated game with an extensive bargaining protocol and unbounded transfers in each period. Their goals and analysis are quite different from this paper's. In particular, they are interested in understanding how axiomatic restrictions on disagreements affect bargaining outcomes. This, together with the presence of a transfer stage separate from the action stage, distinguishes their analysis and results from ours. The relation between these papers and ours is explained in detail in Sections 4.5 and 6.

2 Setting

We first consider the case of two players, indexed by $i \in \{1, 2\}$, engaged in a repeated game. Player i's stage-game action, a_i, lies in a finite set denoted $\mathcal{A}_i$. The vector $\mathbf{a} = (a_1, a_2)$ of actions

similar to a global proposal which requires only the approval of a subset of players, it has a specific structure which we do not investigate in this paper and hope to explore in future work.

determines the players' payoffs for the current period, $\mathbf{u}(\mathbf{a}) = (u_1(\mathbf{a}), u_2(\mathbf{a}))$. A distribution α_i over $\mathcal{A}_i$ is a *mixed action* for i, and $\alpha = (\alpha_1, \alpha_2)$ denotes the vector of mixed actions for both players. Players put a weight $\varepsilon \in (0,1)$ on the current period, which corresponds to a common discount factor $\delta = 1 - \varepsilon$.

Each period consists of the following stages:

1) Players observe the realization z of a public randomization device taking values in $[0,1]$;

2) They simultaneously choose mixed actions $\alpha_i \in \Delta(\mathcal{A}_i)$, $i \in \{1,2\}$. Mixing probabilities are not observable.[7] Conditional on the realization z of the public randomization device, players choose their mixed actions independently from each other;

3) The vector $\mathbf{a}$ of actions is observed and the period's payoffs are realized;

4) With probability $p < 1$, one of the players is chosen to propose a new *plan* for continuation of the game.[8] Each player has the same probability of $\frac{p}{2}$ being chosen.[9] The chosen player may conceal his proposal opportunity by remaining silent, or mix between making a proposal or staying silent. The object of a proposal is an infinite-horizon plan m from the set $\mathcal{M}$ of all possible plans, which will be described shortly;

5) If i made a proposal, player $-i$ decides whether to accept it, possibly mixing between acceptance and rejection. The resulting decision, D_{-i}, is set to 1 if $-i$ accepts the proposal and 0 if he rejects it;

[7]In accordance with current practice, we allow players to use privately mixed strategies. This feature distinguishes our analysis from some of the earlier work on renegotiation. For example, Farrell and Maskin (1989) assume that players can observe each other's mixing strategies, rather than just the realized actions. This distinction can affect the set of weakly renegotiation-proof equilibria (Farrell and Maskin's concept), as we show in Appendix H. Intuitively, when players observe each other's mixed strategy, there is without loss a single continuation payoff vector, conditional on players' mixed strategies. When mixtures are unobservable, however, there must be a continuation vector for each possible outcome of the mixture, chosen so as to make each player indifferent across all actions in the support of his mixed strategy. Moreover, all of these vectors must belong to the renegotiation-proof set. This is problematic because some of these continuations may have Pareto-ranked payoffs, violating weak renegotiation-proofness. Bernheim and Ray (1989) rule out mixing altogether, focusing their analysis on pure-strategy equilibria.

[8]It is possible to define proposals in terms of a different message space. For example, players could simply propose continuation payoffs. Although many message spaces are possible, any accepted message must correspond to some equilibrium of the dynamic game, and the message space must be rich enough to include all equilibria which the players may want to consider as continuations of the game. The advantage of presenting messages directly as plans is to make this connection explicit including, in particular, not only players' continuation payoffs but also whether the equilibrium implementing the payoffs satisfies some renegotiation-proof refinement or whether it belongs to some social norm.

[9]Our results extend to the case of asymmetric probabilities. The sufficient conditions are unchanged, but necessary conditions entail a payoff lower bound on each player, which increases with that player's proposal probability, consistent with the intuition that a higher proposal probability means an increased bargaining power. This extension is covered in Appendix G.

The public history of a period consists of the realisation z of the randomization device, the action vector $\mathbf{a}$, the proposal (which we will later denote as μ_i) or absence thereof and, if applicable, the acceptance decision D_{-i}. In addition, each player privately observes the mixing probability used for each of his decisions.

A *plan* in period t describes players' strategy for the infinite repetition of the stage-game described above, *from period* $t+1$ *onwards.* These decisions (actions, proposals, and acceptance mixtures) are history-dependent. The setting being time invariant, it is convenient to define recursively the set $\mathcal{M}$ of plans. A plan $m \in \mathcal{M}$ in period t is described by the following elements:

a) For each realization z of the public randomization device, a pair $\alpha = \alpha[m](z)$ of mixed actions that players should play in period $t+1$;

b1) For each player i, a distribution $\bar{\mu}_i = \bar{\mu}_i[m](z, \mathbf{a}) \in \Delta(\mathcal{M} \cup \emptyset)$ over proposals, where the outcome $\emptyset$ means that i abstains from making a proposal (unbeknownst to player $-i$). We assume that distributions have a finite support over plans.[10] The proposal distribution can depend on the realization z of the public randomization device and on the pair $\mathbf{a}$ of observed actions. Because $p < 1$, not observing any proposal from either player is always consistent with "on-path" behavior. The realized proposal is denoted μ_i;

b2) A probability $q_{-i} = q_{-i}[m](z, \mathbf{a}, \mu_i)$ that $-i$ accepts i's proposal (whenever $\mu_i \neq \emptyset$), which may depend on z, $\mathbf{a}$, and μ_i;

b3) If no one made a proposal, the acceptance stage is skipped. To economize on notation, we assume that some player i is, even in that case, conventionally selected (randomly or deterministically) as the proposer and let $\mu_i = \emptyset$ and $D_{-i} = 0$. (So, $-i$'s conventional response is to systematically "reject" a non proposal.)

c) A continuation plan $m_{+1} = m_{+1}[m](z, \mathbf{a}, i, \mu_i, D_{-i}) \in \mathcal{M}$ for period $t+2$ onwards, as a function of z, $\mathbf{a}$, i, μ_i, D_{-i}, where i indicates the identity of the last proposer.[11]

This protocol allows plans in which all proposals are ignored regardless of whether they are accepted (babbling). In the next section, we introduce a refinement requiring that some accepted proposals be played. The protocol also allows plans for which any rejected proposal results in the same "default" continuation, which is the "No-Fault Disagreement" Axiom studied by Santos (2000) and Miller and Watson (2013) and, in a simultaneous-offer setting, by Farrell (1987), Rabin (1994), Arvan, Cabral, and Santos (1999).[12]

[10] We will in fact impose a uniform upper bound on this support, as explained below.

[11] Clearly, this plan must be independent of i whenever $\mu_i = \emptyset$, so that the convention chosen for the proposer in the absence of any actual proposal is indeed irrelevant. This restriction is applied throughout.

[12] The protocol also allows counter-intuitive plans for which an accepted proposal is followed by a continuation plan which has nothing to do with the initial proposal. Appendix G explains why one could without loss restrict attention to plans which are "truthful," i.e. such that accepted on-path proposals are always played. This discussion is postponed to avoid cluttering the analysis.

While the above definition seems natural, it turns out to be too permissive for the set of plans to be well-defined: there does not exist a set of plans so large as to contain all the possible continuation prescriptions allowed above.[13] The appendix provides restrictions on plan prescriptions guaranteeing that the set of plans is well-defined and flexible enough to include all the plans discussed in this paper, so this difficulty can be ignored in a first reading.[14]

3 Concepts

The previous section has introduced an enlarged infinite-horizon game which we term *repeated game with renegotiation.* Any strategy profile of this game can be identified with a plan. Indeed, a plan prescribes history-dependent distributions over decisions (actions, proposals, or acceptance) at each stage of each period of the game. Accordingly, the subgame perfect equilibria (SPEs) of the repeated game with renegotiation can be identified as a subset $\mathcal{S}$ of $\mathcal{M}$. Unless stated otherwise, in this paper "SPE" refers to an equilibrium of the repeated game with renegotiation.

Our main concept is a simple equilibrium refinement, which confers its strength to renegotiation.

DEFINITION 1 *An SPE m is* renegotiation-proof *if, as long as no off-path proposal has been accepted, any proposal that is accepted is played.*

Equivalently, players implement all accepted proposals until, possibly, some off-path proposal is accepted. Unlike the previous literature on renegotiation-proofness, this concept involves no *set* of equilibria; it reduces renegotiation-proofness to a simple equilibrium refinement. However, there is an equivalent formulation of the concept that follows the more familiar of approach of viewing a

[13] For example, in the above construction, a plan must specify an acceptance decision for each possible proposal. Therefore, each plan m must specify—among other things—a function which maps each element of $\mathcal{M}$ (the proposal) to a binary decision (acceptance). This implies that the set $\mathcal{M}$ of plans must contain, in order to include all possible prescriptions, its power set $2^{\mathcal{M}}$. Such a set does not exist, since any set has a strictly lower cardinality than its power set, by Cantor's Power Set Theorem (see, e.g., Mendelson (1997)).

[14] These restrictions are of three kinds: First, the support of the proposal distribution $\bar{\mu}_i[m](z, \mathbf{a})$ and the set $\mathcal{M}^{+1}[m]$ of possible continuation plans have cardinalities which are uniformly bounded over m, z, and $\mathbf{a}$. Second, all plans prescribe to accept on-paths proposals and reject off-path ones. Third, the continuation plan is chosen in $\mathcal{M}^{+1}[m]$ according to a choice rule which depends only on the following information: i) whether the proposal was on path, ii) whether it was accepted, and iii) for each player and continuation plan in $\mathcal{M}^{+1}[m]$, the pairwise ranking of the proposal's payoff relative that continuation plan's payoff. There exists a set of plans obeying these restrictions which has the same cardinality, $\beth_2$, as the set of real-valued functions over $\mathbb{R}$. Having fixed this set, one may construct the extensive formulation of each plan from its recursive formulation. To avoid cluttering the analysis, we defer the details to Appendix F.

norm as a set of equilibria, and which will facilitate the comparison of the concept with existing notions of renegotiation-proofness.

Definition 2 *A set* $\mathcal{N} \subset \mathcal{S}$ *of equilibria forms a* **norm** *if, for any* $m \in \mathcal{N}$ *such that* $\mu_i \in \bar{\mu}_i[m](z, \boldsymbol{a})$ *or* $D_{-i} = 0$*,* $m_{+1}[m](z, \boldsymbol{a}, i, \mu_i, D_{-i}) \in \mathcal{N}$*.*[15]

The definition implies that, if players start with an equilibrium in the norm, then all on-path proposals (whether they are accepted or rejected), as well as rejected off-path proposals, have their continuations in the norm. In particular, deviations in actions are punished within the norm, as long as no off-path proposal to leave the norm has been accepted. One may view $\mathcal{N}$ as a social norm: it describes the set of continuations which players perceive as consistent with "business as usual." A norm can only be abandoned if some player makes an off-path proposal outside the norm that is accepted by the other player. The following notion of stability requires that such proposals be taken seriously by the players.

Definition 3 *A norm* $\mathcal{N}$ *is* **stable** *if, in any period starting with a continuation equilibrium in* $\mathcal{N}$*, whenever* i *proposes an equilibrium* $\mu \in \mathcal{S}$ *and* $-i$ *accepts it,* μ *is implemented.*

Since all on-path continuations of equilibria in $\mathcal{N}$ must all belong to $\mathcal{N}$—by definition of a norm—stability implies that any Pareto-improving proposal lying outside the norm is rejected with probability 1; for if it were accepted, stability would require that the proposal be implemented. Stability thus requires that no player ever has an incentive to make proposals outside of the norm—hence the terminology. Intuitively, stability is achieved by rewarding a player on the receiving end of a deviating proposal whenever he rejects it. Crucially, however, this continuation, which rewards the rejector and deters the proposer, must lie *within the norm.*

As anticipated, norm stability is equivalent to renegotiation-proofness in the following sense.

Proposition 1 *An equilibrium is renegotiation-proof if and only if it is part of a stable norm.*

The proof of this equivalence is straightforward: First, any SPE of a stable norm must be renegotiation-proof since all continuations of on-path proposals and rejected off-path proposals lie in the norm, and thus subject to the stability condition. For the reverse direction, take any renegotiation-proof SPE and consider the set consisting of this equilibrium together with all of its continuation equilibria at the beginning of periods following histories for which no off-path proposals has been accepted. This set forms a norm, by construction, which is stable, by renegotiation-proofness of the equilibrium.

[15]When viewing an equilibrium m as a plan, the notation $m_{+1}[m](z, \mathbf{a}, i, \mu_i, D_{-i})$ refers to the continuation equilibrium of m at the next period, following the observations $z, \mathbf{a}, i, \mu_i, D_{-i}$ in the current period.

These definitions may seem demanding: why should players treat all proposals as "credible" when they are accepted? A pragmatic answer is that the question is moot: Section 5 shows that restricting the refinement to a much narrower class of "credible" proposals yields the same necessary and sufficient conditions as the ones obtained when all proposals are covered by the refinement.[16]

A more conceptual answer is that all proposals are chosen from *equilibria* of the repeated game with renegotiation, which may have their own prescriptions regarding how proposals are handled, and do not have to satisfy the refinement of the initial norm. For instance, player 1 may offer to player 2 to move to another equilibrium of the underlying repeated game and treat all future proposals as babbling. If player 2 accepts this offer, it may be reasonable for the players to ignore future negotiation.

There is, in fact, an important conceptual reason *not* to impose the renegotiation-proofness refinement beyond the first accepted proposal: Players may agree, once, to implement a proposal lying outside the norm, as implied by the content of the refinement. However, requiring that further proposals which are off-path relative to the first one be played when accepted undermines the credibility of this refinement: since these further proposals are off-path, implementing them would imply *de facto* that the first proposal has not in fact been implemented and, hence, that the renegotiation-proofness refinement invoked when the players accepted the first proposal has been violated.

Payoffs

To delineate players' incentives at each node of the game, we introduce notation that distinguishes players' payoffs at different stages of each period. Given a subset $\mathcal{L}$ of SPEs, let $\mathcal{U}(\mathcal{L}) \subset \mathbb{R}^2$ (or just $\mathcal{U}$, when there is no confusion) denote the set of expected payoffs for the players across all possible SPEs in $\mathcal{L}$, computed before public randomization. $\mathcal{V}$ is defined identically but computed after the realization of the randomization device z. $\mathcal{U}$ is thus included in the convex hull of $\mathcal{V}$. Finally, $\mathcal{W}$ consists of continuations payoffs after actions and payoffs are observed and incurred in the current period, but before the proposal stage. Each element of $\mathcal{W}$ is a convex combination of three expected payoff vectors corresponding to the following events: player 1 gets to make a proposal, player 2 does, or no one does. Because elements of $\mathcal{W}$ define continuation payoffs excluding the current period, to make them commensurate with payoffs in $\mathcal{U}$, we evaluate them at the next period (i.e., ignoring the discount factor between the two periods). With this convention, payoffs in $\mathcal{W}$ are convex combinations of elements of $\mathcal{U}$.

Elements of $\mathcal{U}$, $\mathcal{V}$, and $\mathcal{W}$ are points of two-dimensional sets. For any point U, we let $\pi_i(U)$ denote the i^{th} component of U, i.e., i's continuation payoff at the relevant stage of the game.

[16] A proposal is *credible* relative to a norm if any ulterior deviation from the proposal, whether at the action or the proposal stage, triggers a reversal to the norm.

Definition 4 *A point A is* **q-renegotiation-proof** *if there exists $\bar{\varepsilon} \in (0, \frac{1}{q})$ such that for all $\varepsilon \leq \bar{\varepsilon}$ there exists a stable norm $\mathcal{N}$ for $p = q\varepsilon$ such that $A \in \mathcal{U}(\mathcal{N})$. Moreover, A is* **renegotiation-proof** *if it is q-renegotiation-proof for all q's large enough.*[17]

The coefficient q is inversely related to the amount of renegotiation frictions in the game:[18] when $q = 0$, players never get a chance to renegotiate and the game reduces to a standard repeated game. As with the Folk Theorem, any point in interior of the individually rational feasible set is 0-renegotiation-proof. Our main objective is to characterize the set of renegotiation-proof payoffs, i.e., the set of payoffs which sustain renegotiation, when opportunities to renegotiate become arbitrarily frequent.

4 Main Result

4.1 Statement

Let $\underline{v}_i$ denote i's minmax payoff in the stage game of the repeated game absent any renegotiation[19] and P_i denote the feasible payoff vector that gives i his maximal payoff among all payoff vectors above the minmax.[20] The weak individually-rational Pareto frontier—consisting of points which are not strictly Pareto dominated—is a piecewise linear curve joining P_1 and P_2.

Let $v_1 = \pi_1(P_2)$ and $v_2 = \pi_2(P_1)$.

Theorem 1 (Renegotiation-Proof Set)

- *(Sufficiency) If*
$$\pi_i(A) > v_i \qquad \text{for } i \in \{1, 2\} \tag{1}$$
or $A = P_1 = P_2$, then A is q-renegotiation-proof for all $q \geq 0$ and, hence, renegotiation-proof.

- *(Necessity) If A is q-renegotiation-proof, then*
$$\pi_i(A) \geq \underline{v}_i + \max\left\{0; \frac{q}{2+q}(\pi_i(P_{-i}) - \underline{v}_i)\right\} \tag{2}$$
for $i \in \{1, 2\}$. If A is renegotiation-proof, inequalities in (1) must hold for both players as weak inequalities.

[17] This definition concerns payoffs evaluated at the beginning of a period, i.e., before the realization of the public randomization device.

[18] One may embed the model into a continuous-time structure, with ε being the time interval between consecutive periods. For small ε (our case of interest), the coefficient q then corresponds to the rate at which players receive proposal opportunities per unit of time.

[19] As usual, player $-i$ is allowed to mix across actions to minmax i.

[20] If several such points exist, we choose the point among those with the lowest payoff for $-i$.

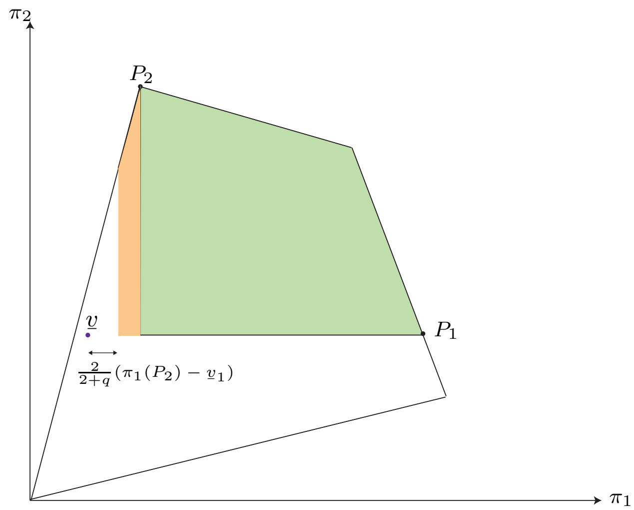

Figure 1: Necessary and sufficient conditions for fixed q

Condition (1) thus fully characterizes (up to its boundary) the set of renegotiation-proof payoffs. Sufficiency is established in Section 4.4 which constructs explicitly, for any A that satisfies (1), a renegotiation-proof equilibrium that implements A.[21] The necessary condition is derived in Appendix A.[22]

Figure 1 illustrates Theorem 1 for a fixed q: the green region represents the points known to be renegotiation-proof and the orange region represents the additional points which may be renegotiation-proof. When $q = 0$ (no renegotiation), the orange region extends all the way to the minmax point $\underline{v}$ and we recover the Folk Theorem. As renegotiation frictions become arbitrarily small ($q \to +\infty$), the orange region disappears as necessary and sufficient conditions become identical (up to their boundary).

One consequence of Theorem 1 is that the set of q-renegotiation-proof payoffs is nonempty for all values of $q \geq 0$ and so is the set of renegotiation-proof payoffs. In particular, our concept of renegotiation-proofness provides a well-defined counterpoint to the standard Folk Theorem when renegotiation is introduced to repeated games, allowing us to compare the impact of renegotiation across different strategic situations of the stage game, from perfectly aligned interests to extreme misalignments, and to establish for a large class of games the possibility of sustaining inefficient equilibria even when players are arbitrarily patient and can frictionlessly and credibly propose and

[21] The construction focuses on the case $P_1 \neq P_2$. If $P_1 = P_2$, players have perfectly aligned interests as they both want to implement P_1 and the construction is trivial.

[22] When $P_1 = P_2$, the necessary condition selects this point as the unique outcome as renegotiation frictions become negligible. If the weak Pareto frontier consists of a segment giving a constant payoff to one of the players—a degenerate case—the Pareto point maximizing the other player's payoff is renegotiation-proof. Our conditions do not pin down which of the other Pareto points are renegotiation-proof.

agree on Pareto improving equilibria.[23]

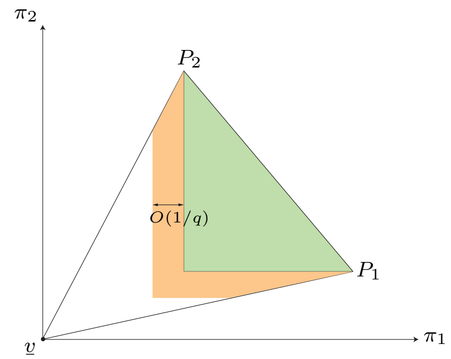

(a) Renegotiation destroys the Folk Theorem

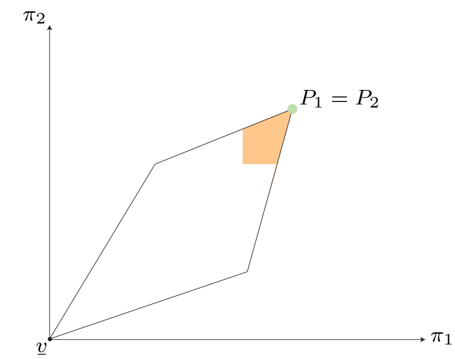

(b) Pareto frontier reduced to one point

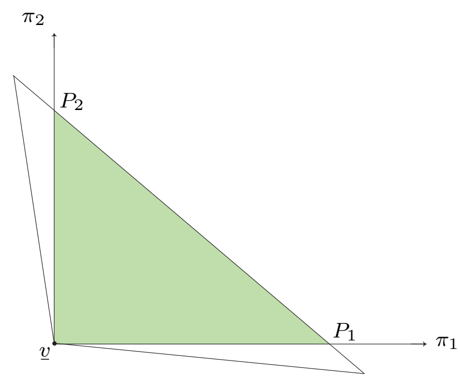

(c) Folk Theorem with extreme deterrence points

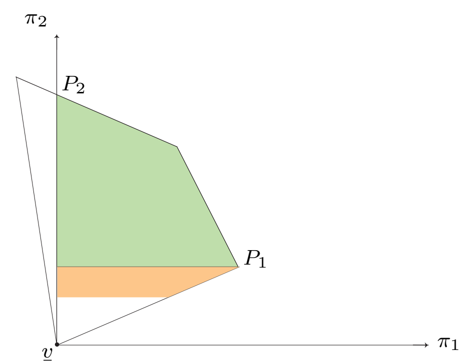

(d) Asymmetric case

Figure 2: Necessary and sufficient conditions for various configurations

4.2 Relation between players' interest alignment and renegotiation-proof outcomes

Figure 2 represents the set of renegotiation-proof payoffs for degrees of player alignment. In configuration (a), renegotiation constrains the set of implementable payoffs because the deterrence points P_1 and P_2 are too close to each other relative to the vector of minmax payoffs. Configuration (b) represents a perfectly cooperative game. The only renegotiation-proof outcome is the Pareto efficient point. In configuration (c), the punishment/reward points used to deter off-path proposals are sufficiently far apart and the Folk Theorem holds despite the presence of frictionless renegotiation.

[23] It should be noted that for fixed ε, there need not exist any stable norm—or, equivalently, any renegotiation-proof equilibrium—just as strongly renegotiation-proof equilibria (Farrell and Maskin (1989)) and externally consistent norms (Berneim and Ray (1989)) may fail to exist for fixed discount factors. Indeed, we have constructed a family of counter-examples for some fixed $\varepsilon > 0$ and all values of $q > 0$.

As the figure illustrates, the impact of renegotiation hinges on the alignment structure of the stage game. As the game becomes less cooperative (moving from (b) to (a) to (c) on the figure), there is more scope for disagreement among the players, which can be used to implement a larger set of feasible payoffs. Strategic renegotiation thus does not destroy the implementability of Pareto-efficient payoffs, but does not prevent Pareto-inefficient ones either, and the severity of the inefficiency which may be sustained increases as players' interests become more divergent.

4.3 Applications: a novel mechanism for miscoordination and inertia

Perhaps the most intriguing consequence of this analysis is the fact that players may be stuck with rules, constructed explicitly in Section 4.4, which prevent them from moving to Pareto efficient payoffs, even though the game has complete information, players are arbitrarily patient, and can perfectly communicate and commit to new equilibria.

Pareto inefficient equilibria may be desirable in some applications, in which the players who are explicitly modeled create externalities on other economic agents. A social planner in charge of designing rules between the players may be concerned that high payoffs for these players mean that they are colluding, polluting, shirking, or, more generally, adversely affecting society members who cannot influence these players in return.

Consider, for instance, a regulator wishing to prevent collusive pricing in an oligopolistic market. If the firms can be given self-enforcing rules that prevent collusion, such a design is of course cheaper for the regulator than explicitly monitoring the firms and administrating the punishments. Likewise, the manager of administrative office facing high costs of monitoring his employees may wish to create a social norm between them which implements high effort and under which an employee's proposal to shirk is rebuked by other employees and thwarted without requiring the manager's intervention. The designer's role is then simply to set the rules at the beginning of the game, specifying how players should interpret deviations in actions and proposals. Once this common understanding is reached, the designer completely withdraws from the game: the players enforce the rules themselves by punishing one another if one of them ever deviates from these rules.

Of course, proposal-deterring norms do not have to be designed by anyone: players may simply be trapped in an equilibrium with this feature—perhaps the remain of an unmodeled evolution before which such a norm made sense. An example may be "acting tough" and discouraging any suggestion to "soften up" even when do so would in fact lead to a Pareto improvement.

Another insight of the analysis—related to Section 4.2 on players' alignment—is to emphasize the potential value, from a designer's perspective, of *creating* actions which benefit only one player but not the other, in order to deter collusive proposals more easily. Instead of taking the stage game as given, that is, the designer (regulator, manager, etc) may instil some amount of potential disagreement between the players.

Trapped in an improvement-deterring norm, players may be tempted to reach out to each other anyway and urge each other to forget about the current norm. Note, however, that such an attempt *precisely* constitutes the kind of proposal that this paper analyzes. Unless players can somehow create neologisms that are perceived by other players as immune from the consequences of the current rules, the attempt will be interpreted as a proposal by the receiver and the proposer punished as a result. Worse: *it is in the receiver's interest to interpret any attempt to reach out as a deviating proposal*, since by design rejecting such a proposal provides a reward to the rejector, above and beyond the improvement entailed by the proposal.

In other applications, the designer's motivations need not reflect any concern for *any* agent's welfare—whether explicitly modeled or passive— other than himself. We briefly describe one application of each kind below, where Pareto-inefficient equilibria from the players' perspective are beneficial or harmful from a broader social perspective.

Cournot competition.

Consider two symmetric firms which, in equilibrium, produce together more than the monopolistic output. These firms could achieve a higher profit by each producing half of the monopolistic output. However, proposals to move away from the current equilibrium may be subject to a norm treating any such proposal as corrupt behavior. The firm on the receiving end of such a proposal could reject it, triggering a new equilibrium in which, say, the rejector produces the Stackelberg leader's output in each period and the proposer produces the Stackelberg follower's output. These outputs constitute an equilibrium, which gives the proposer a lower payoff than the competitive equilibrium and his competitor a higher payoff than the half of the monopoly's profit.[24]

Political inertia and dictatorship.

Consider an authoritarian regime facing the risk of a revolution. In this regime, citizens may be exploited through high taxes, expropriation, and other channels. Faced with this situation, various citizen factions may attempt to persuade others to start a revolution (a proposal that lies outside of the current norm). If all citizens agree on the revolution, the authoritarian regime is toppled and citizens all become better off. However, the regime may impose a norm that thwarts this threat by rewarding anyone who reveals the plot and punishing its instigator. Rewards and punishments are all administered by the citizens, without the dictator needing to get involved or even monitor them.[25] This provides a novel, completely endogenous explanation for the stability of dictatorships,

[24]The punishment for the proposer, i.e., the Stackelberg equilibrium, is inefficient. However, it suffices to incentivize a rejection to the proposal and thus deter a Pareto-improving proposal.

[25]While the application obviously involves more than two players, the gist of the many-player analysis is identical to the two-player one, as shown in the many-player extension of Section 7.

which allows coordination among citizens but exposes the limits of attempts to coordinate when the norm in place anticipates such attempts.

The dynamic nature of social norms and the importance of neologisms

The previous examples hint at the dynamic nature of social norms, particularly with regard to how "innovative" proposals are perceived. For instance, starting from a Pareto-dominated equilibrium, a proposal to increase cooperation and increase both players' payoffs may be viewed as a good idea and implemented. To be sustained, however, this cooperation may require the threat of punishments in which the kind cooperation originally proposed is no longer acceptable.

In other applications, this norm dynamic looks as if each player had an endogenous "reputation," not about an intrinsic type—such type does not exist in our model—but about the kind of actions or proposals that she is allowed to pursue. If a player proposes to disrupt the current equilibrium (say, by implementing a higher cooperation, or a revolution), she risks losing her reputation if the other player rejects her proposal. Following a rejection, the proposer may be treated as if she were "soft" or "crazy" (for making such a proposal) and shunned or punished as a result.

Players facing an improvement-deterring norm may be tempted to reach out to each other anyway and argue that they should forget the current norm and coordinate on a Pareto-improving equilibrium. As noted earlier, however, such a message amounts to exactly the kind of proposal that our model aims to capture. If it is clear to all players (firms, citizens) that reaching out in this way constitutes a proposal punishable within the norm, then it is definitely in the receiver's interest to reject the proposal, which deters the proposer from making his proposal in the first place. When "proposals" belong to a general message space (a possibility mentioned in Section 2), a comprehensive norm should prescribe the deterrence response to any possible message. Realistically, however, this opens the possibility of environments with "weak" norms, which are not immune to neologisms to which no pre-specified continuation has been attributed.

4.4 Proof (sufficiency): construction of a renegotiation-proof equilibrium

Outline. We construct, for any point A of the feasible set that satisfies (1) and ε small enough, an equilibrium m which implements A and is renegotiation-proof for all $q \geq 0$. The construction is based on points A_1 and A_2 such that A_i gives i his worst possible payoff among all continuations of m.[26] When i's continuation payoff is at an ε-independent distance above his payoff from A_i, it is easy to incentivize him to follow any prescribed action, since any deviation provides a maximal gain of order ε and can be punished by implementing A_i. One challenge is to choose A_i so that i is adequately incentivized near A_i. The second important points of the construction are D_1 and

[26]Unless stated otherwise, points refer to payoffs viewed from the beginning of the current period, i.e., to elements of the set $\mathcal{U}$.

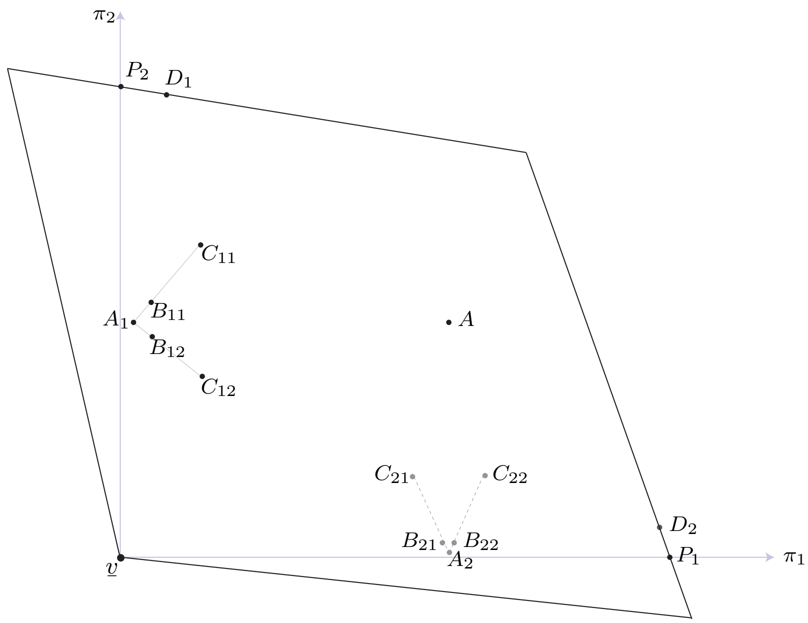

Figure 3: Construction of a renegotiation-proof equilibrium

D_2, which serve to deter off-path proposals. These points are chosen to be Pareto efficient, and set so that any relevant off-path proposal by i may be deterred by having $-i$ reject the proposal and have D_i be implemented instead. D_i must therefore be chosen so that $-i$ is sufficiently rewarded, and i punished, for any proposal that i may entertain.

For each player i, there are two configurations to consider, depending on whether i's minmax payoff $\underline{v}_i$ lies above or below $\pi_i(P_{-i})$. We first consider the case in which both players are in the former configuration.

Notation: throughout the analysis, for any payoff vector X implemented by some continuation of m, we denote by X^m the corresponding continuation.

Case 1: $\underline{v}_1 = \pi_1(P_2)$ **and** $\underline{v}_2 = \pi_2(P_1)$

Consider any point A satisfying (1). For ε small enough, the points A_1 and A_2 with coordinates

$$\pi_1(A_1) = \underline{v}_1 + \varepsilon^{\frac{1}{2}}; \qquad \pi_2(A_1) = \pi_2(A)$$

and

$$\pi_1(A_2) = \pi_1(A); \qquad \pi_2(A_2) = \underline{v}_2 + \varepsilon^{\frac{1}{2}}$$

are individually rational and such that $\pi_1(A_1) < \pi_1(A)$ and $\pi_2(A_2) < \pi_2(A)$.

The equilibrium A_1^m implementing A_1 is constructed as follows (A_2^m has a similar construction):

1) Action stage: player 2 minmaxes player 1, possibly mixing between several actions $\{a_{2j}\}_j$. Player 1 best responds by a pure action $a_{1,minmax}$ achieving his minmax payoff.

1a) If no deviation in action is observed, the continuation payoff vector $B_{1j} \in \mathcal{W}$ is a function of 2's realized action, a_{2j}, and is chosen so that i) 2 is indifferent between all actions a_{2j} used to minmax 1, ii) 1's continuation payoff is independent of j (so the vectors $\{B_{1j}\}_j$ all lie on the same vertical line as shown on Figure 3), and iii) the promise-keeping condition is satisfied for both players. In particular,

$$\pi_1(A_1) = \varepsilon \underline{v}_1 + (1-\varepsilon)\pi_1(B_{1j}) \tag{3}$$

for all indices j corresponding to some action a_{2j} in 2's minmaxing distribution. In particular, the points B_{1j} all lie within an ε-proportional distance of A_1.

1b) If 2 deviates in action (i.e., chooses an action outside of the mixture used to minmax 1), the continuation payoffs jump to the point A_2, mentioned above, which gives 2 her lowest possible payoff.[27] For small ε, this punishment suffices to incentivize 2 because any deviation gain is of order ε whereas $\pi_2(A_2)$ is arbitrarily close to 2's minmax payoff, causing 2 an ε-independent loss.

1c) If 1 deviates in action, disregard this. Such a deviation is suboptimal since 1 was prescribed to best respond to being minmaxed by 2.

2) Proposal stage: the equilibrium B_{1j}^m implementing B_{1j} is as follows: if either 2 gets a chance to make a proposal, or no player does, the continuation payoffs return to A_1^m. 2 is prescribed to remain silent. If 1 gets a chance to make a proposal, he proposes a continuation C_{1j}^m whose corresponding payoff vector C_{1j} lies on the line going through A_1 and B_{1j} and is chosen so as to satisfy the promise-keeping condition

$$\pi_1(B_{1j}) = \left(1 - \frac{p}{2}\right)\pi_1(A_1) + \frac{p}{2}\pi_1(C_{1j}) \tag{4}$$

Player 2 is prescribed to accept proposal C_{1j}^m. The points $\{C_{1j}\}_j$ give the same payoff to 1, independently of j. Their implementation is described in 3) below.

2a) If 1 proposes any plan other than C_{1j}^m that improves his payoff, he is punished by an equilibrium D_1^m such that i) $\pi_1(D_1) < \pi_1(C_{1j})$ and ii) 2 prefers $\pi_2(D_1)$ to her payoff under 1's proposal. Precisely, D_1 is defined as the point of the Pareto frontier that gives 1 a payoff of

$$\frac{\pi_1(A_1) + \pi_1(C_{1j})}{2} \tag{5}$$

As explained shortly, 1's payoff at C_{1j} is of order $\sqrt{\varepsilon}$ above what 1 gets at A_1 or B_{1j}. If 1 proposes a plan that makes him worse off than C_{1j}^m, 2 accepts it if only if improves her payoff. Of course, such a proposal never arises in equilibrium.

2b) If 2 deviates by making a proposal or rejecting 1's offer to move to C_{1j}, the equilibrium jumps to A_2^m, which punishes 2's deviation (in the former case, it is optimal for 1 to rejects 2's proposal and trigger A_2^m).

[27] More precisely, it jumps to the point B_{21}, which is the analogue of the point B_{11}, following the implementation of A_2.

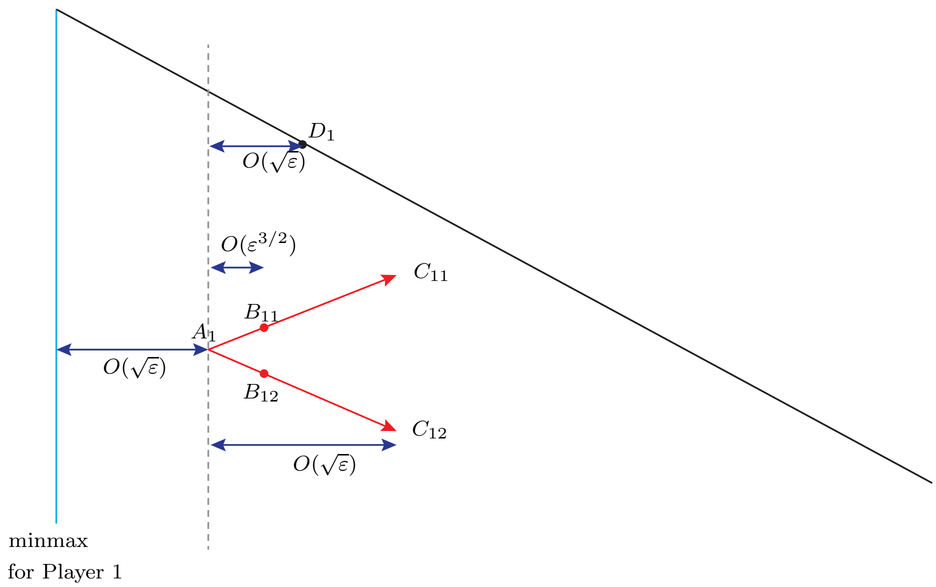

Figure 4: Construction details

3) Next periods: the equilibrium C_{1j}^m is easily implemented because it gives 1 a payoff of order $\sqrt{\varepsilon}$ above what A_1 and B_{1j} give him. A deviation in action by 1 brings a gain of order ε and is punished by a drop of order $\sqrt{\varepsilon}$ in 1's continuation payoff, and is thus suboptimal, for ε small enough. C_{1j}^m can be implemented by a deterministic sequence of actions keeping players' continuation payoffs within a distance $K\varepsilon$ from C_{1j}. The rules implementing this sequence are simple: play a deterministic action profile keeping continuation payoffs ε-close to C_{1j} and do not allow any proposal. If 1 deviates in actions, jump to one of the continuations B_{1j}^m; if he deviates in proposals, jump to D_1^m if 2 rejects this offer. A similar rule is applied for player 2, who has even more to lose from a deviation.

4) The payoff D_1 also gives 1 a payoff of $\sqrt{\varepsilon}$ above A_1 and B_{1j}. D_1^m can therefore be implemented similarly to C_{1j}^m. Again, any proposal is ignored.

The construction is represented on Figure 3. The magnitudes of payoff differences between the points involved in the construction are indicated on Figure 4.

We verify the claim that all C_{1j}'s lie at a $\sqrt{\varepsilon}$-proportional distance to the right of A_1. From (3) and (4), we have

$$\pi_1(A_1) = \varepsilon \underline{v}_1 + (1-\varepsilon)\pi_1(B_j) = \varepsilon \underline{v}_1 + (1-\varepsilon)\left[\left(1-\frac{q\varepsilon}{2}\right)\pi_1(A_1) + \frac{q\varepsilon}{2}\pi_1(C_{1j})\right]$$

Ignoring the terms of order ε^2 and higher, this implies that

$$\pi_1(A_1) = \varepsilon \underline{v}_1 + \left(1-\left(1+\frac{q}{2}\right)\varepsilon\right)\pi_1(A_1) + \frac{q\varepsilon}{2}\pi_1(C_{1j}).$$

Subtracting $\pi_1(A_1)$ from both sides and dividing by ε yields

$$\varepsilon^{\frac{1}{2}} = \pi_1(A_1) - \underline{v}_1 = \frac{q}{2}\left(\pi_1(C_{1j}) - \pi_1(A_1)\right), \tag{6}$$

which shows the claim.

The direction of each vector $\overrightarrow{A_1C_{1j}}$, which is also $\overrightarrow{A_1B_{1j}}$'s direction, depends only on 2's action, a_{2j}; it does not change when ε goes to 0. This shows that, for ε small enough, C_{1j} is a feasible payoff and $\pi_2(C_{1j})$ exceeds $\pi_2(A_2)$ by an ε-independent value.

As noted, the system of actions and proposals implementing A_i^m's, B_{ij}^m's and C_{ij}^m's and D_i^m's is incentive compatible in actions and in proposals. To conclude the construction, observe that A gives each player i a payoff higher than A_i, by an amount that is independent of ε. One may therefore implement A by a deterministic sequence of actions, chosen so that the continuation payoffs stay within a distance $K\varepsilon$ of A.[28] Deviations in actions are punished by moving to B_{11}^m or B_{21}^m, depending on the deviator's identity. Deviations in proposals are similarly punished by moving to D_1^m or D_2^m.

To verify that the equilibrium is renegotiation-proof, notice that whenever 1 gets to make a proposal (at any of continuations considered in the construction), his payoff is at least $\pi_1(D_1)$. Since D_1 is on the Pareto frontier, any proposal giving 1 strictly more than $\pi_1(D_1)$ must give 2 less than $\pi_2(D_1)$. This means that D_1^m can serve as a punishment in case 1 makes such a proposal.

Remaining cases: $\underline{v}_1 < \pi_1(P_2)$ **and/or** $\underline{v}_2 < \pi_2(P_1)$

The construction is almost identical in other cases. The only difficulty is that the difference $\pi_1(A_1) - \underline{v}_1$ is now bounded below away from zero, whereas it was previously of order $\sqrt{\varepsilon}$. This may lead to situations in which the points C_{1j} constructed above are no longer feasible and/or give 2 a payoff lower than $\pi_2(A_2)$. The difficulty is easily addressed by adding, for each j, a point E_{1j} lying on the segment $[A_1B_{1j}]$—and thus also on the line (A_1C_{1j})—such that if player 2 gets a chance to make a proposal, or if nobody does, players' continuation payoffs jump to E_{1j}. The promise keeping condition (4) becomes

$$\pi_1(B_{1j}) = \left(1 - \frac{p}{2}\right)\pi_1(E_{1j}) + \frac{p}{2}\pi_1(C_{1j}) \tag{7}$$

Choosing E_{1j} close enough to B_{1j} ensures that C_{1j} lies within a distance $\sqrt{\varepsilon}$ of B_{1j} and, hence, of A_1. This guarantees that C_{1j} is feasible and does not drop below $\pi_2(A_2)$, so that the rest of the argument for the first case can be applied. To implement E_{1j}^m, we use public randomization to implement it as a probabilistic mixture of A_1^m and C_{1j}^m.

[28] It is possible to show that A, A_1, and A_2 can all be implemented so that players' continuation payoffs eventually converge to a Pareto-efficient point. Under this "redemptive" implementation, if players switch to a Pareto-inefficient element following a deviation, they will eventually forgive and forget past deviations.

4.5 Relation to the existing literature

When renegotiation is viewed as a strategic interaction, renegotiation-proof equilibria may contain Pareto-ranked continuations. This happens when Pareto-improving proposals are dissuaded by punishing the proposer and rewarding the rejector beyond the proposal. This idea also underlies the results of Santos (2000) who considers players bargaining over which equilibrium to play in a one-shot game, as well as Miller and Watson's (2013) Theorem 1, which shows that renegotiation has no restrictive power when it must only obey their "Internal Agreement Consistency" Axiom. That theorem and ours differ in two important ways. First, their argument requires unbounded transfers: to punish a proposer, say player 1, one requires him to make a very high transfer to 2 in the next period. If the weight of a single period is ε, the transfer must be of order $\frac{1}{\varepsilon}$, hence the necessity of unbounded transfers as ε goes to zero. These large transfers permit 1's continuation value to jump immediately from some punishment payoff v_1^0 to a higher continuation value v_1, which is easy to implement. Second, the transfer stage takes place, in each period, before the action stage (and, in particular, is distinct from it). If 1 deviates by making a lower transfer than prescribed, it suffices to have him minmaxed by the other player and reset the continuation value to v_1^0 for the next period in order to punish this deviation.

When stage-game payoffs are bounded, as in our setting, the continuation value of a player cannot jump by an ε-independent amount. The equilibrium construction must thus keep track of continuation values and make sure that these continuation values are implementable at all stages and following all deviations. In the absence of a separate transfer stage, moreover, if player 1 deviates in action when implementing v_1^0, his continuation value must fall below v_1^0. Implementing this lower value may be difficult or even impossible. In fact, it is this impossibility which creates new restrictions on the set of renegotiation-proof payoffs and destroys the Folk Theorem obtained in Miller and Watson's Theorem 1.

Both Santos (2000) and Miller and Watson (2013) consider a further restriction, which is that the continuation of the game, in case of a disagreement, be independent of the identity of the proposer and of the nature of the proposals. This restriction guarantees a higher level of efficiency. The consequences for our model of such a refinement are studied in Section 6.

4.6 Comparative statics: bargaining frictions and discounting

In standard repeated games with public randomization, it is well known that the set of implementable payoffs gets larger as players become more patient. This property does not hold with renegotiation. For example, suppose that the stage game has an inefficient Nash equilibrium that violates the necessary conditions obtained by Theorem 1 for $q = \frac{1}{2}$. For small ε, Theorem 1 implies that this Nash equilibrium payoff, and an open neighborhood around it, is not renegotiation-proof.

However as ε goes to 1, there is an equilibrium in which players follow this Nash equilibrium in the first period (before possibly renegotiating to a Pareto superior continuation). Since the current-period weight is arbitrarily close to 1, players' payoffs are arbitrarily close to the inefficient Nash equilibrium's payoffs, which was impossible with a small enough value of ε.

Although discount-factor monotonicity is violated in the presence of renegotiation, a different kind of monotonicity arises here, with respect to negotiation frictions: the more opportunities players have to renegotiate their continuation equilibrium, the smaller the renegotiation-proof set. This result holds at all discount factor levels and is proved in Appendix D.

PROPOSITION 2 *For any fixed* $\varepsilon \in (0, 1)$*, the set of renegotiation-proof payoffs is decreasing in* q*.*

5 Equivalent notions of stability

This section considers two variations on our definition of renegotiation-proof equilibrium. The first one introduces the notion of a *credible* proposal and restricts our earlier refinement to credible proposals. The second concept is a set-theoretic definition of the renegotiation-proof set. We show that both concepts are equivalent to the one used in previous sections.

Since these variations are naturally presented in terms of norms rather than equilibrium refinements, we use throughout this section the language of stable norms of Definitions 2 and 3 rather than the renegotiation-proofness refinement of Definition 1, recalling that these definitions are equivalent, by Proposition 1.

5.1 Credible proposals

Stability requires that players implement any continuation equilibrium as long as it is proposed and accepted. When players are used to a given norm $\mathcal{N}$, one may wonder why players should take all proposals seriously, particularly when these proposals lie outside of the norm. It turns out that Theorem 1's necessary and sufficient conditions are identical if one restricts proposals to a much smaller subset of "credible" proposals.

DEFINITION 5 *Given a norm* $\mathcal{N}$*, an equilibrium is* $\mathcal{N}$*-credible (or just "credible", when there is no confusion) if any off-equilibrium play (action, proposal, or acceptance decision) triggers a continuation equilibrium in* $\mathcal{N}$ *for the corresponding stage.*

Starting with an equilibrium in some norm $\mathcal{N}$, a credible proposal is thus an SPE such that any deviation triggers a reversal to the norm. For example, if a norm includes a harsh punishment equilibrium for both players, the norm can sustain many credible equilibria, any deviation from which triggers the punishment equilibrium.

Definition 6 *A norm $\mathcal{N}$ is credibly stable if it satisfies the refinement of Definition 3 for all $\mathcal{N}$-credible proposals.*

Definition 6 is clearly more permissive than Definition 3, because it imposes the refinement over a smaller set of proposals. However, we have the following result.

Theorem 2 *The set of points sustained by credibly stable norms obeys the necessary and sufficient conditions of Theorem 1.*

The proof is straightforward: first, any stable norm is credibly stable since the latter must sustain fewer proposal challenges than the former. Our construction for the sufficiency condition thus still applies. Second, the proposals used in Appendix A to derive the necessary conditions of Theorem 1 are credible, as shown in this appendix. The necessary conditions are thus identical for stable and credibly stable norms.

5.2 Set-theoretic definition: closed vs. open norms

The norms that we defined earlier were *open* in the sense that they allowed players to actually depart from the norm in case an off-equilibrium proposal was made and accepted. This property is quite unlike the purely set-theoretic approach taken by earlier works which do not explicitly consider departures from their norms, whether on path or off path. These approaches are reconciled here: we show that norm stability can entirely be recast in terms of a purely set-theoretic definition.

We start by abandoning altogether the equilibrium refinement underlying norm stability in Section 3. Unlike this earlier section, players now do not actually take any proposal outside of the norm seriously, but instead consider their norm as the only possible outcomes, which in effect "closes" our definition of a norm:

Definition 7 *A subset $\mathcal{N}$ of $\mathcal{S}$ is a* **closed norm** *if for any $m \in \mathcal{N}$, $m_{+1}[m](z, \boldsymbol{a}, i, \mu_i, D_{-i}) \in \mathcal{N}$.*

The only difference with Definition 2 is that continuations now still belong to the norm even when off-path proposals are accepted. To offset this change, our earlier definition of stability is translated into the language of set-theoretic analysis. To keep in line with the previous section, we state the definition for credible proposals. Dropping credibility from the definition has no impact on the equivalence.

Definition 8 *A closed norm $\mathcal{N}$ is* **stable** *if it satisfies the following property: Consider any SPE of $\mathcal{N}$ and history at which i gets a chance to make a proposal, and let $\hat{U}_i$ denote i's continuation payoff. Then, for any credible proposal with payoff vector U which gives i a payoff $\pi_i(U) > \hat{U}_i$, there exists a payoff vector U' of $\mathcal{N}$ such that $\pi_{-i}(U') \geq \pi_{-i}(U)$ and $\pi_i(U') \leq \hat{U}_i$.*

THEOREM 3

1. For any closed norm $\mathcal{N}^c$, there exists an open norm $\mathcal{N}^o$ which has the same payoff set, and vice versa.

2. For any stable closed norm $\mathcal{N}^c$, there exists a stable open norm $\mathcal{N}^o$ which has the same payoff set, and vice versa.

6 Renegotiation-proof equilibria in the absence of proposer-specific punishments

One virtue of explicitly modeling the renegotiation process is to incorporate the logic of modern repeated games analysis into renegotiation: just as arbitrary continuation equilibria may follow from observed actions in a repeated game, here arbitrary continuations may follow rejected proposals. The paper has explored one consequence of this generality, which is that even good proposals may be deterred, and Pareto dominated equilibria be sustained as a result.

While in the applications discussed earlier this flexibility seemed reasonable or even desirable, in other environments it is natural to ask what equilibria may be sustained when proposers cannot be punished. Indeed, such a restriction is imposed in a number of models of explicit negotiation[29] and sometimes formalized as a "No-Fault Disagreement" (NFD) axiom. The axiom requires the continuation equilibrium following a rejected proposal to coincide with the default continuation in case no proposal was made. This section shows how our results are modified when this refinement is added.

In order to keep the language of the analysis as close as possible to the existing literature, this section adopts the "stable norms" terminology of Definitions 2 and 3 instead of the renegotiation-proofness refinement.[30]

DEFINITION 9 *A stable norm $\mathcal{N}$ is* forgiving *if for any SPE m in $\mathcal{N}$, for any i and μ_i, $m_{+1}[m](z, \boldsymbol{a}, i, \mu_i, 0){=}m_{+1}[m](z, \boldsymbol{a}, i, \emptyset, 0)$.*

Our concepts of renegotiation-proofness are modified as follows. A payoff vector A is said to be *forgivingly q-renegotiation-proof* if for all ε small enough, there is a forgiving stable norm containing an equilibrium which expected payoff is equal to A. A is forgivingly renegotiation-proof if it is forgivingly q-renegotiation-proof for all q's large enough.

The main result in this case is given by the novel necessary conditions, which are much more restrictive those of Theorem 1: the continuation payoffs must lie within a distance $O(\frac{1}{q})$ of the

[29] See Santos (2000) and Miller-Watson (2013). A similar idea appears in Farrell (1987), Rabin (1994), and Arvan, Cabral, Santos (1999) for the case of simultaneous announcements.

[30] The concepts of this section can be readily adapted to modify renegotiation-proof equilibria.

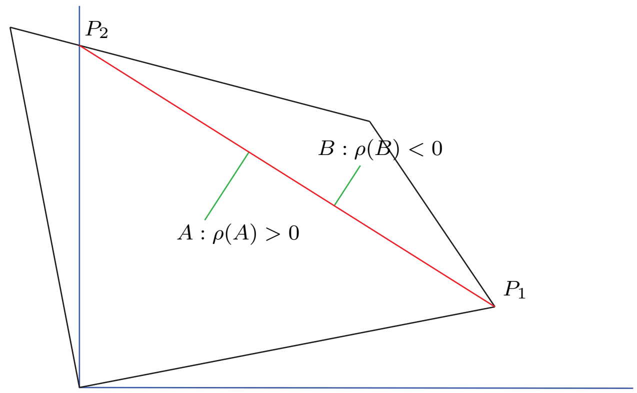

Figure 5: Signed distance from (P_1P_2)

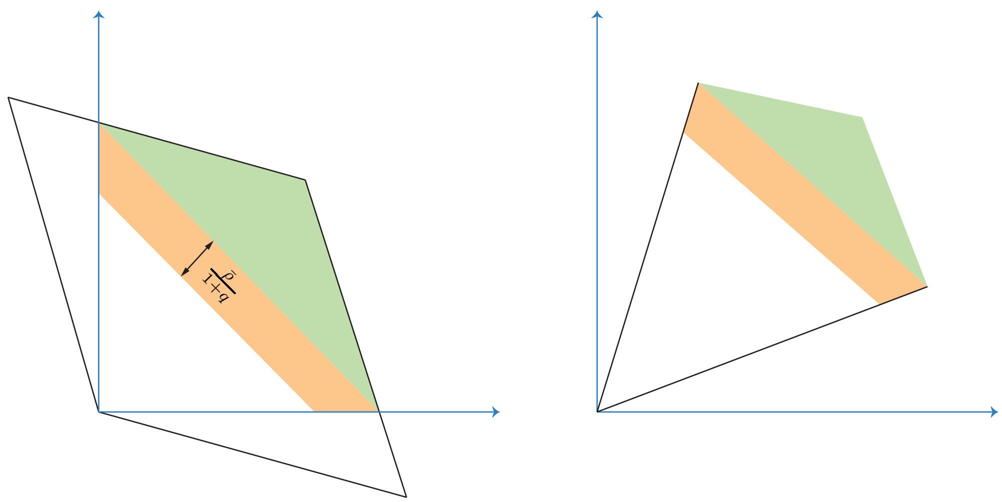

Figure 6: Renegotiation-proof payoffs without proposer-specific punishments

convex hull of the (individually-rational) Pareto frontier. More precisely, for each feasible payoff vector A, let $\rho(A)$ denote the signed distance from the line (P_1P_2), counted positively if A lies below (P_1P_2), and negatively otherwise, as indicated by Figure 5.

Let $\bar{\rho}$ denote the maximum value of ρ among all feasible payoff vectors.

Theorem 4 *If A is forgivingly q-renegotiation-proof, then* $\rho(A) \leq \frac{\bar{\rho}}{1+q}$.

One may also wonder whether all the feasible payoffs lying above the line (P_1P_2) can be achieved in this case. The next result provides a positive answer which is independent of negotiation frictions. To establish this result, we slightly modify the definition of stability, as follows: deviating proposal which is accepted needs to be implemented only if it improves the proposer's payoff by more than a constant $\eta > 0$, arbitrarily small but fixed, over his equilibrium payoff without the deviation.[31]

[31] Using the refinement in Theorem 4 affects the corresponding bound by a factor η.

Definition 10 *A norm $\mathcal{N}$ is η-**stable** if the following holds: consider any SPE of $\mathcal{N}$ and history at which i gets a chance to make a proposal and let $\hat{U}_i$ denote i's continuation payoff. Then, whenever i proposes a plan $\mu \in \mathcal{S}$ giving him at least $\hat{U}_i + \eta$, and $-i$ accepts it, μ is implemented.*

Theorem 5 *Assuming η-stability, any payoff vector A strictly above the segment (P_1P_2) is forgivingly renegotiation-proof.*

The role of η is to prevent off-path proposals whose payoffs lie near the boundary of the norm's payoff set, as detailed in the proof of the theorem.

7 Three or more players

The analysis so far has focused on two players, a common restriction to study renegotiation in repeated games.[32] Extending the analysis to more players raises new conceptual issues. Can proposals be targeted toward a subset of individuals? What happens if only a subset of the players accepts the proposal?

This section explores some of these issues, allowing for an arbitrary number, $n \geq 3$, of players. After a player has made a proposal, other players vote on accepting the proposal. The continuation payoff may *a priori* depend on the identity of the players who voted for the proposal. We consider several dependence structures which vary in their flexibility. The analysis mainly focuses on environments in which players vote simultaneously over the acceptance decision but also briefly considers the case of sequential acceptance decisions.

The setting is identical to the two-player case, except for the following aspects. At the proposal stage, each player i has the same probability $\frac{p}{n}$ $(p < 1)$ of being chosen to propose a new plan. The renegotiation friction parameter q is still defined by $p = q\varepsilon$. This player may choose to conceal his proposal opportunity. If i makes a proposal, other players decide on whether to accept it, resulting in a vector of acceptance votes $D_{-n} \in \{0,1\}^{n-1}$.[33] Given a plan m at the beginning of a period, the continuation plan at the next period, $m_{+1} = m_{+1}[m](z, \mathbf{a}, i, \mu_i, D_{-i}) \in \mathcal{M}$ is a function of z, $\mathbf{a}$, i, μ_i, and D_{-i}, where i indicates the identity of the last proposer.

In keeping with the previous two sections, we again focus on the concept of norm stability, which extends to an arbitrary number of players as follows.

Definition 11 *A subset $\mathcal{N}$ of $\mathcal{S}$ is a **norm** if, for any $m \in \mathcal{N}$ such that $\mu_i \in \bar{\mu}_i[m](z, \boldsymbol{a})$ or $D_{-i} \neq \{1\}^{n-1}$, $m_{+1}[m](z, \boldsymbol{a}, i, \mu_i, D_{-i}) \in \mathcal{N}$;*

[32] E.g., Farrell and Maskin (1989), Benoît and Krishna (1993), and Santos (2000). Abreu et al. (1993) focus instead on symmetric equilibria.

[33] As with the two-player case, if no proposal is made the identity of a proposer is arbitrarily chosen and the null proposal is assumed to be rejected by everyone else.

DEFINITION 12 *A norm $\mathcal{N}$ is* **stable** *if for any SPE of $\mathcal{N}$, whenever i proposes an equilibrium $\mu \in \mathcal{S}$ and all other players accept it, μ is implemented.*

Stability specifies the continuation payoff when *everyone* accepts the proposal. One must also define players' continuation payoffs following the rejection of the proposal. We consider three specifications. The most flexible one allows continuation payoffs to depend arbitrarily on the identity of the players who accepted and rejected the proposal. The other two specifications have binary continuation payoffs: one if all players agree on the proposal; the other if at least one player rejects it. These specifications differ with regard to the continuation payoff in case of a rejection: can the proposer be punished if his proposal is rejected, or does the game proceed as if no proposal had taken place, as already prescribed by the No-Fault Disagreement axiom studied in Section 6?

The last two specifications are captured by the following concepts.

DEFINITION 13 *A stable norm $\mathcal{N}$ is* **simple** *if, for any $m \in \mathcal{N}$ and $D_{-i}, D'_{-i} \neq \{1\}^{n-1}$,*

$$m_{+1}[m](z, \boldsymbol{a}, i, \mu_i, D_{-i}) = m_{+1}[m](z, \boldsymbol{a}, i, \mu_i, D'_{-i}).$$

DEFINITION 14 *A simple norm $\mathcal{N}$ is* forgiving *if, for any SPE $m \in \mathcal{N}$, $m_{+1}[m](z, \boldsymbol{a}, i, \mu_i, D_{-i} \neq \{1\}^{n-1})$=$m_{+1}[m](z, \boldsymbol{a}, i, \emptyset, \{0\}^{n-1})$, for any i, μ_i.*

The definitions of (forgivingly) q-renegotiation-proof payoffs and (forgivingly) renegotiation-proof payoffs are identical to those of the two-player case.

Throughout the analysis, we assume that the individually-rational payoff set has a full dimension an in Fudenberg and Maskin (1986).

7.1 Voter-dependent continuations

Suppose, first, that continuation payoffs can depend arbitrarily on the voting decision of each player—except if everyone agrees on a proposal, in which case stability dictates that the proposal is implemented. With this high degree of flexibility, norms may be constructed so that all negotiation proposals are dissuaded and the Folk Theorem obtains.

THEOREM 6 *For any feasible payoff vector π with $\pi_i > \underline{v_i}$ for all i, π is renegotiation-proof.*

To understand this result, we recall that in the underlying repeated game without negotiation, any strictly individually-rational payoff vector can be implemented for ε small enough by minmaxing any player i who deviates in actions, and switch to minmaxing any player $j \neq i$ who deviates when minmaxing player i. The same idea can be applied when negotiation is possible, by deterring it as follows: if a player, i, proposes another continuation, everyone else is prescribed to reject the proposal and to start minmaxing player i. If another player, j, deviates from the prescribed rejection by accepting i's proposal, and all other players reject it, then players are prescribed to minmax j instead of i. If two or more players accept i's proposal, it is implemented, which guarantees that the norm satisfies our stability refinement. This prescription guarantees that it is always suboptimal for a player to unilaterally accept a proposal and, consequently, that it is also suboptimal to make any proposal. Unless some additional restrictions are imposed on the continuation payoffs, allowing for the possibility of renegotiation with three or more players thus has no more predictive power on the set of equilibria and payoffs than the standard Folk Theorem.

7.2 Simple Norms

Suppose now that only two continuations may follow each proposal, depending on whether all players have agreed to it. As usual with voting games, we eliminate equilibria involving weakly dominated strategies.

ASSUMPTION 1 *A player votes in favor of the proposal if it gives him a strictly higher payoff than its continuation payoff in case of a rejection.*

Let $\mathcal{P}$ denote the Pareto frontier of the feasible payoffs in the stage game and, for each i, P_{-i} denote any individually-rational payoff vector of $\mathcal{P}$ which minimizes i's payoff.

The key question, for characterizing stable norms, is to determine each player i's worst possible punishment if he makes an unprescribed proposal. Suppose that i proposes an SPE of the renegotiated game, with corresponding payoff vector C, and let $\mathcal{V}$ denote the set of achievable payoff vectors in our candidate norm, $\mathcal{N}$. If $\mathcal{N}$ is stable, C will be implemented if all players accept i's proposal. If anyone rejects the proposal, norm simplicity implies that there is a single payoff vector *in the norm*, $D(C)$, which will be realized. If $D(C)$ gives $\pi_j(C)$ or more to at least one player $j \neq i$, that player will refuse the implementation of C, and the equilibrium implementing $D(C)$ will be played.

Following any proposal with a payoff C by player i, the worst punishment in $\mathcal{V}$ for player i minimizes i's utility over the set:[34]

$$\mathcal{D}(C, \mathcal{V}) = \{D(C) \in \mathcal{V} : \exists j \neq i : \pi_j(D(C)) \geq \pi_j(C)\}.$$

Let $\underline{\pi_i}(C, \mathcal{V})$ denote i's utility under this worst punishment.

Viewing $\underline{\pi_i}(C, \mathcal{V})$ as a function of C, one can then find the proposal with a continuation $C(\mathcal{V})$ which maximizes i's payoff at the worst punishment: $C(\mathcal{V}) = \arg\max_C \{\underline{\pi_i}(C, \mathcal{V})\}$, and the corresponding payoff, $\underline{\pi_i}(\mathcal{V})$, for i.

[34]For the existence of a *worst* punishment, the set $\mathcal{V}$ needs to be closed. Our construction will satisfy this condition.

One should think of the payoff $C(\mathcal{V})$ as follows. The most efficient way to prevent player i from making a non-prescribed proposal is by implementing his worst punishment. Anticipating this, if player i deviates from his prescription, he may as well choose the optimal proposal, which gives the payoff $C(\mathcal{V})$.

These observations lead to the following sequential construction. We start from the set $\mathcal{F}$ of strictly individually-rational payoffs in the stage game, i.e., what would be implementable without renegotiation. We consider the minimal payoffs $\underline{\pi_i}(\mathcal{F})$, $i \in \{1, ..., n\}$ which any player i could guarantee himself if having a chance to make a proposal and the payoffs sustained by the norm were in $\mathcal{F}$. We will build *two* decreasing sequences of sets, starting from $\mathcal{F}$, which will generate separate necessary and sufficient conditions for a payoff to be renegotiation-proof.

To derive sufficient conditions, the k^{th} set in the sequence, $\mathcal{F}_s^k$, is reduced by removing all the payoffs below $\underline{\pi_i}(\mathcal{F}_s^k)$, to form the $k+1$-th set in the sequence, starting with $\mathcal{F}_s^0 = \mathcal{F}$. We will show that this process converges to a stable set which defines sufficient conditions.

To derive necessary conditions, let $\pi_{min,i}(\mathcal{F}_n^k)$ denote the lowest expected payoff for player i at the beginning of a period, among all payoff vectors in $\mathcal{F}_n^k$. This value is lower than the continuation payoff $\underline{\pi_i}(\mathcal{F}_n^k)$ that i can guarantee himself when he gets a chance to make a proposal. We have

$$\pi_{min,i}(\mathcal{F}_n^k) \geq \varepsilon \underline{v_i} + (1-\varepsilon)[\frac{q\varepsilon}{n}\underline{\pi_i}(\mathcal{F}_n^k) + (1-\frac{q\varepsilon}{n})\pi_{min,i}(\mathcal{F}_n^k)]$$

Indeed, as in the two-player case, i gets at least $\underline{v_i}$ as his current payoff, and can guarantee himself $\underline{\pi_i}(\mathcal{F}_n^k)$ if he has a chance to make a proposal. As ε goes to 0, one can express the value $\pi_{min,i}(\mathcal{F}_n^k)$ as:

$$\pi_{min,i}(\mathcal{F}_n^k) \geq \frac{n\underline{v_i} + q\underline{\pi_i}(\mathcal{F}_n^k)}{n+q}. \tag{8}$$

At each step the set $\mathcal{F}_n^k$ is being reduced by removing the payoffs below (8). Iterations of this procedure converge to a steady set, as we show in the Appendix.

Proposition 3 *Both procedures converge to steady sets.*

We denote the limiting sets by $\mathcal{V}_s$ and $\mathcal{V}_n$. They both are positive orthants, whose vertices give lower bounds on players' payoffs (calculated at the beginning of period) under both procedures, and are denoted $\pi_{min,i}(\mathcal{V}_s)$ and $\pi_{min,i}(\mathcal{V}_n)$, for any player i. By construction, expression (8) holds as an equality for $\mathcal{V}_n$:

$$\pi_{min,i}(\mathcal{V}_n) = \frac{n\underline{v_i} + q\underline{\pi_i}(\mathcal{V}_n)}{n+q} \tag{9}$$

Similarly, we have $\pi_{min,i}(\mathcal{V}_s) = \underline{\pi_i}(\mathcal{V}_s)$.

We can now state the main result of this section. Let $\mathcal{R}$ denote the open positive orthant whose vertex is the vector $(\pi_i(P_{-i}))_{i=1}^n$. In the two-player case, this set characterized the sufficient conditions for renegotiation-proofness. With $n > 2$ players, we show that $\mathcal{R}$ still consists of renegotiation-proof vectors, thought it might not include all of them. The theorem is formulated for the case where Pareto frontier supports for each player a non-zero range of payoffs.

Theorem 7 *Any renegotiation-proof payoff lies in $\mathcal{V}_n$, and generically any payoff in the interior of $\mathcal{V}_s$ is renegotiation-proof. Moreover, any payoff in the interior of $\mathcal{R}$ is renegotiation-proof.*

In case of feasible set having the unique Pareto-efficient point, it is the only renegotiation-proof payoff. [35]

In the two-player case, the necessary and sufficient conditions became arbitrary tight as renegotiation frictions vanished. The same is true in this more general environment, as shown in Appendix E.4.

Proposition 4 *The sets $\mathcal{V}_s$ and $\mathcal{V}_n$ converge to each other as q goes to infinity.*

If the players are making responses to proposals sequentially, then one has the same result [36]. Each proposal has only two continuations. If the continuation in case of rejection benefits at least one responder, he rejects the proposal in any extended game. Otherwise, from backward induction, each player votes for the proposal if it gives higher payoff than rejection - same as in simultaneous voting.

7.2.1 Sequential voting in case of no restrictions

Sequential voting permits more than two continuation payoffs, depending on the sequence of acceptance decisions of the players. The resulting stable norm is qualitatively similar to the earlier analysis with only two continuations, and it is more permissive.

Proposition 5 *Suppose that each proposal is decided by sequential voting. Then, analogous constructions to the two-continuation case yield sufficient and necessary conditions characterized by upper orthants. Moreover, each of these sets is larger than the corresponding set obtained with only two continuations.*

Sequential voting with many continuations thus provides more predictive power than simultaneous voting, but less predictive power than the simultaneous-voting specification with only two continuations.

[35] In the non-generic case of several Pareto points, each giving the same payoff to one of players, in case of at least two other players having different payoffs on Pareto frontier, the renegotiation-proof payoffs always form a non-empty full-dimensional orthant. In case Pareto frontier gives the same payoff to all but one player, there is a best Pareto point, which is renegotiation-proof.

[36] We actually got the equivalence between sequential and simultaneous votes when assuming that player votes for the proposal if it gives him higher payoff compared to rejection.

7.3 Three or more players without proposer-specific punishment

Finally, consider the most restrictive case of a simple norm that is also forgiving, as defined in the two-player case. The necessary conditions resemble the two-player case. Consider the set of individually-rational Pareto-efficient payoffs $\mathcal{P}'$, and consider the convex hull of this set, $Co(\mathcal{P}')$. Then one has:

PROPOSITION 6 *If A is forgivingly q-renegotiation-proof, the distance from A to $Co(\mathcal{P}')$ is bounded above by a decreasing function of q, which converges to 0 as q becomes arbitrarily large.*

The proof closely mirrors the argument used for the two-player case and is only sketched here. Suppose that A is the point of the norm which has the largest distance from $Co(\mathcal{P}')$ and that A lies "too far" down away from $Co(\mathcal{P}')$. Whenever a player gets to make a proposal—which happens with probability proportional to q—he proposes a Pareto point (or close to it). Moreover, the continuation payoff A' which follows if the proposal is rejected cannot lie farther away from $Co(\mathcal{P}')$ than A does. Combining this puts a bound on A's distance to $Co(\mathcal{P}')$, which vanishes as q gets large.

We conclude this section with sufficient conditions, whose derivation is more involved and described in Appendix E.6.

THEOREM 8 *Assuming η-stability, any point A in the set $Co(\mathcal{P}')$ lying strictly above the minmax is forgivingly renegotiation-proof.*

8 Discussion

Understanding and tractably modeling renegotiation in repeated games has been a longstanding challenge. This paper's approach, based on strategic renegotiation and an equilibrium refinement, delivers a characterization of renegotiation-proof equilibria, which is straightforward to describe graphically and has several equivalent formulations. The concepts introduce here shed light on existing definitions of norms and informal notions of social norms, whether they are dynamic, encoded as part of a single equilibrium, or viewed as a set of possible equilibria.

With arbitrarily many players, our analysis suggests that under natural specifications the set of renegotiation-proof payoffs has useful analytical (convergence of necessary and sufficient conditions) and geometric properties (upper orthant characterization). Beyond these results, some important issues remain to be explored. In particular, what happens if a player can make a proposal to a subset of players? How such a proposal, if accepted, affects the strategies used by the players excluded from the proposal? This question seems challenging even when negotiations are public and actions are perfectly monitored, and should realistically be expanded to include private negotiations. Exclusive negotiations of this kind are common in economics, when agents are divided into relatively homogeneous groups within which negotiation is easier or when they are engaged in specific relationships like those arising in supplier chains. They may also arise in community enforcement models, in which matching parties may engage in local renegotiation to alleviate punishments (Ali, Miller, and Yang (2016)). Understanding how strategic renegotiation shapes equilibrium outcomes in environments with segmented groups seems a particularly interesting direction for future work.

References

Abreu, D., Pearce, D. (1991) "A perspective on renegotiation in repeated games," in R. Selten (ed.), *Game Equilibrium Models*, Springer Berlin Heidelberg.

Abreu, D., Pearce, D., and E. Stacchetti (1993) "Renegotiation and Symmetry in Repeated Games," *Journal of Economic Theory*, Vol. 60, pp. 217–240.

Ali, N., Miller, D., and D. Yang (2016) "Renegotiation-Proof Multilateral Enforcement," *Working Paper*, Pennsylvania State University and University of Michigan.

Arvan, L., Cabral, L., and V. Santos (1999) "Meaningful cheap talk must improve equilibrium payoffs," *Mathematical Social Sciences*, Vol. 37, pp. 97–106.

Asheim, G. (1991) "Extending Renegotiation-Proofness to Infinite Horizon Games," *Games and Economic Behavior*, Vol. 3, pp. 278–294.

Benoît, J.-P., Krishna, V. (1993) "Renegotiation in Finitely Repeated Games," *Econometrica*, Vol. 61, pp. 303–323.

Bergin, J., MacLeod, B. (1993) "Efficiency and Renegotiation in Repeated Games," *Journal of Economic Theory*, Vol. 61, pp. 42–73.

Bernheim, B.D., Ray, D. (1989) "Collective Dynamic Consistency in Repeated Games," *Games and Economic Behavior*, Vol. 1, pp. 295–326.

DeMarzo, P. (1988) "Coalitions and Sustainable Social Norms in Repeated Games," *mimeo*, Stanford University.

Farrell, J. (1983) "Credible Repeated Game Equilibria," *Unpublished Manuscript.*

Farrell, J. (1987) "Cheap Talk, Communication, and Entry." *Rand Journal of Economics*, Vol. 18, pp. 34-39.

Farrell, J., Maskin, E. (1989) "Renegotiation in Repeated Games," *Games and Economic Behavior*, Vol. 1, pp. 327–360.

Forster, T. (1995) *Set Theory with a Universal Set: Exploring an Untyped Universe*, Oxford University Press.

Fudenberg, D., Maskin, E. (1986) "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, Vol. 54, pp. 533–554.

Kuratowski, K., Mostowski, A. (1968) *Set theory*, Elsevier, Academic Press.

Mendelson, E (1997) *Introduction to Mathematical Logic*, 4th ed. Chapman & Hall, London.

Miller, D., Watson, J. (2013) "A Theory of Disagreement in Repeated Games with Bargaining," *Econometrica*, Vol. 81, pp. 2303–2350.

Pearce, D. (1987) "Renegotiation-Proof Equilibria: Collective Rationality and Intertemporal Cooperation," *Cowles Foundation Discussion Paper*, No. 855.

Rabin, M. (1994) "A Model of Pre-Game Communication," *Journal of Economic Theory*, Vol. 63, pp. 370-391.

Santos, V. (2000) "Alternating-Announcements Cheap Talk," *Journal of Economic Behavior & Organization*, Vol. 42, pp. 405–416.

A Proof of Theorem 1 (Necessary Conditions)

The interesting case is when $\underline{v}_i < \pi_i(P_{-i})$: otherwise, Theorem 1 predicts only that i's payoff must be individually rational. Without loss of generality, we only derive the necessary condition for player 1.

Suppose therefore that $\pi_1(P_2) > \underline{v}_1$ and, by contradiction, that there is a q-renegotiation-proof point A such that $\pi_1(A) < v_1 = \underline{v}_1 + \frac{q}{2+q}(\pi_1(P_2) - \underline{v}_1)$: one can construct, for any ε small enough, a renegotiation-proof equilibrium m that implements A.

Let C_1 denote 1's infimum payoff over all continuation equilibria of m following histories at which it is 1's turn to make a proposal and no off-path proposal has yet been accepted. Since the Pareto point P_2 is a possible proposal payoff,[37] and since it Pareto dominates all payoffs with $\pi_1 < \pi_1(P_2)$, C_1 must satisfy

$$\pi_1(P_2) \leq C_1.$$

We now contradict this inequality. Let $\mathcal{N}$ denote the set of continuation equilibria of m at the beginning of all periods following histories at which no off-path proposal has been accepted. This set forms a stable norm, by Proposition 1. Also let $A_1 = \inf_{V \in \mathcal{V}(\mathcal{N})} \pi_1(V)$, $B_1 = \inf_{W \in \mathcal{W}(\mathcal{N})} \pi_1(W)$, and $D_1 = \inf_{U \in \mathcal{U}(\mathcal{N})} \pi_1(U)$, and consider any sequence $\{V_k\} \in \mathcal{V}(\mathcal{N})$ such that $\pi_1(V_k) \to_{k \to +\infty} A_1$. For any V_k there is an action that implements it in the first period of the corresponding SPE. However, if player 1 deviates, he can guarantee himself an immediate payoff of at least $\underline{v}_1$, and the worst punishment for him after deviation gives him at least B_1. Therefore, $\pi_1(V_k) \geq \varepsilon \underline{v}_1 + (1-\varepsilon)B_1$. Since this inequality holds for all V_k we obtain, taking the limit:

$$A_1 \geq \varepsilon \underline{v}_1 + (1-\varepsilon)B_1 \tag{10}$$

Since any element of $\mathcal{U}(\mathcal{N})$ lies in the convex hull of $\mathcal{V}(\mathcal{N})$, and player 1 can always conceal his opportunity to propose, we have

$$C_1 \geq D_1 \geq A_1$$

Consider now a sequence $\{W_k\} \in \mathcal{W}(\mathcal{N})$ such that $\pi_1(W_k) \to B_1$. Any element W_k is a weighted average of an expected payoff vector EU_k^1 whenever 1 gets a chance to make a proposal, an expected payoff vector EU_k^2 when it is 2's turn to make a proposal, and a payoff vector U_k^0 in case no one gets to make a proposal:

$$W_k = \frac{p}{2}(EU_k^1) + \frac{p}{2}(EU_k^2) + (1-p)(U_k^0) \tag{11}$$

We note that EU_k^1 is a mixture of elements of $\mathcal{U}$ resulting from 1's mixture over proposals and 2's mixture over her acceptance decision. Similarly, EU_k^2 is a mixture of elements of $\mathcal{U}$.

Since all elements U_k's belong to $U(\mathcal{N})$, we have $\pi_1(EU_k^2) \geq A_1$ and $\pi_1(U_k^0) \geq A_1$. Equation (11) thus implies that

$$\pi_1(W_k) \geq (1 - \frac{p}{2})A_1 + \frac{p}{2}\pi_1(EU_k^1).$$

Recalling that C_1 denotes 1's infimum payoff when he gets to make a proposal, we get

$$\pi_1(W_k) \geq (1 - \frac{p}{2})A_1 + \frac{p}{2}C_1.$$

Taking limits,

$$B_1 \geq (1 - \frac{p}{2})A_1 + \frac{p}{2}C_1$$

or

$$B_1 \geq (1 - \frac{q\varepsilon}{2})A_1 + \frac{q\varepsilon}{2}C_1. \tag{12}$$

[37] By the Folk Theorem, P_2 can be implemented by an SPE of the repeated game without renegotiation. By treating all proposals as cheap talk, P_2 can thus also be implemented as an SPE of the game with renegotiation.

Combining (10) and (12), we conclude that

$$A_1 \geq \varepsilon \underline{v}_1 + (1-\varepsilon)[(1-\frac{q\varepsilon}{2})A_1 + \frac{q\varepsilon}{2}C_1]$$

or, ignoring terms of order ε^2 in right-hand side,

$$A_1 \geq \varepsilon \underline{v}_1 + (1-[1+\frac{q}{2}]\varepsilon)A_1 + \frac{q\varepsilon}{2}C_1.$$

Subtracting A_1 on both sides of the last equation and dividing by ε, we obtain

$$0 \geq \underline{v}_1 - [1+\frac{q}{2}]A_1 + \frac{q}{2}C_1 \tag{13}$$

From $A_1 \leq \pi_1(A)$, $C_1 \geq \pi_1(P_2)$, and $\pi_1(A) < v_1 = \underline{v}_1 + \frac{q}{2+q}(\pi_1(P_2) - \underline{v}_1)$, we get

$$0 < \underline{v}_1 - [1+\frac{q}{2}]A_1 + \frac{q}{2}C_1$$

which contradicts (13). This shows the necessary condition for player 1. An identical reasoning for player 2 shows the second necessary condition. This proves the result for $P_1 \neq P_2$. A similar reasoning applies when $P_1 = P_2$.

Credible proposals Section 5 introduced the concept of $\mathcal{N}$-credible proposals, and claimed that the necessity conditions were unaffected if the proposals involved in the definition of stability were restricted to being credible. To prove this claim, it suffices to verify that the proposal to move to P_2, used just above to derive the necessary condition, is $\mathcal{N}$-credible. The SPE implementing P_2 is constructed as follows: players are prescribed to play, in all periods, the pure action profile with payoff P_2, and to abstain from making any proposal. Any deviation, whether in action or in proposal, triggers the equilibrium implementing A—which is supposed to exist, by the contradiction hypothesis. Clearly, player 2 cannot benefit from deviating as she is getting her highest possible payoff in the game. Moreover, the difference $\pi_1(P_2) - \pi_1(A)$ is by assumption bounded below by $\frac{2}{2+q}(\pi_1(P_2) - \underline{v}_1)$, which is ε-independent. Therefore, 1 cannot benefit from deviating either: a deviation in action may create an immediate gain of order ε, but triggers a drop in continuation payoffs that is ε-independent and dominates the gain. A deviation in proposal yields the payoff vector A, which again is detrimental to 1.

B Concept equivalence

1. Any closed norm $\mathcal{N}^c$ is an open norm as well, so the first statement is trivially true. Now consider any open norm $\mathcal{N}^o$. To construct a payoff-equivalent closed norm $\mathcal{N}^c$, we modify each plan/equilibrium m of $\mathcal{N}^o$ as follows: m's rules on and off the equilibrium path are kept unchanged except when a player, say i, makes a proposal μ_i which is off the equilibrium path. In this case, because $\mathcal{N}^o$ is an open norm, the continuation equilibrium if $-i$ accepts the proposal need not lie in $\mathcal{N}^o$. Following such a proposal, players are instead prescribed to behave as if i had remained silent. The new rules define an equilibrium: when playing the original equilibrium m, i was not making the proposal μ_i anyway, so removing this option does not affect *equilibrium* behavior and payoffs. By construction, the set of modified equilibria form a closed norm $\mathcal{N}^c$, and because each equilibrium of $\mathcal{N}^o$ has been modified into a single payoff-equivalent equilibrium of $\mathcal{N}^c$, the norms are payoff equivalent.

2. We start with the observation that if two norms $\mathcal{N}^c$ and $\mathcal{N}^o$ have the same payoff sets, then any proposal that is credible according to either norm is credible according to the other norm.

We now consider any stable open norm $\mathcal{N}^o$ and construct the corresponding closed norm $\mathcal{N}^c$ as in Part 1. To show that $\mathcal{N}^c$ is stable, consider any SPE m of $\mathcal{N}^c$, history at which player i gets to propose, and credible proposal U such that $\pi_i(U)$ is strictly greater than i's continuation payoff $\hat{U}_i$. From the above observation, U is also credible for $\mathcal{N}^o$. If the proposal U gives player $-i$ a lower payoff that $\hat{U}$ does, then the payoff $U' = \hat{U}$ satisfies Definition 8. If the proposal U Pareto dominates $\hat{U}$, then for the equilibrium $\tilde{m}$ of $\mathcal{N}^o$ corresponding to m, and the same history, $-i$ must reject U with positive probability (for otherwise $\pi_i(U)$ would coincide with $\hat{U}_i$). Let U' denote the continuation payoff if $-i$ rejects U. By stability of $\mathcal{N}^o$, $-i$ knows that if he accepts U it will be implemented. Since it is weakly optimal for $-i$ to reject U, it must therefore be the case that $\pi_{-i}(U') \geq \pi_{-i}(U)$. Moreover, it must also be the case that $\pi_i(U') \leq \hat{U}_i$, for otherwise it would be strictly optimal for i to deviate by proposing U, and $\tilde{m}$ would not be an equilibrium. Using this U' in Definition 8, this implies that $\mathcal{N}^c$ is stable.

Next, consider any stable closed norm $\mathcal{N}^c$. To construct a payoff-equivalent stable open norm $\mathcal{N}^o$, we simultaneously modify all SPE's of $\mathcal{N}^c$. The modification proceeds in two steps, and is based on the recursive definition a plan. Recall that a plan is a prescription of actions, proposals and acceptance decisions for the next period (each depending on what happened in earlier stages), along with a continuation plan resulting from these stages to be applied in the period after next. In Step 1, we modify the prescriptions for time $t+1$, and still use plans of $\mathcal{N}^c$ as continuations plans. The purpose of this step is to make a prescription compatible with the requirement that if a Pareto-improving, credible proposal is made and accepted, then it has to be played. In Step 2, we replace these continuation plans of $\mathcal{N}^c$ by those built in Step 1, to guarantee that the rule applies at all periods, ensuring that credible proposals which are accepted are implemented, so that Definition 6 holds at all periods.

Consider any SPE m of $\mathcal{N}^c$. We modify m as follows. For the modified SPE $\tilde{m}$, the action stage and on-path proposals are prescribed exactly as in m.[38] Now consider a history at which i makes any proposal U which is not prescribed by m but which is $\mathcal{N}^c$-credible. If $-i$ accepts the proposal, we construct $\tilde{m}$ by prescribing that players implement this proposal.[39] If the proposal gives i a strictly higher payoff than his equilibrium continuation payoff $\hat{U}_i$, then by stability of $\mathcal{N}^c$, there must exist a payoff vector U' corresponding to some equilibrium m' of $\mathcal{N}^c$, which gives player $-i$ at least as much as U, and which gives player i at most $\hat{U}_i$. We prescribe playing the equilibrium corresponding to U' in case player $-i$ rejects the proposal. If U does not improve upon i's equilibrium continuation payoff, we prescribe playing the continuation equilibrium corresponding to any of i's equilibrium proposals in case $-i$ rejects U. Finally, if i makes a non-credible proposal, the proposal is ignored as if i had stayed silent.

We now verify that $\tilde{m}$ is an SPE that yields the same payoff as m. Since $\tilde{m}$ prescribes the same actions as m, players are incentivized to follow the prescription. If i gets a chance to make a proposal, any proposal prescribed by m (and hence $\tilde{m}$) yields the same continuation payoff as in m. If player i makes a credible, off-equilibrium proposal that improves upon his equilibrium payoff, then player $-i$ is incentivized to reject it, and i's continuation payoff is weakly lower than his equilibrium payoff. It is never optimal for i to make a credible proposal that is lower than his equilibrium payoff, regardless of $-i$'s acceptance decision. Finally, we replace all continuation plans by their modified versions.

There remains to verify that the set consisting of all modified equilibria forms a stable open norm, denoted $\mathcal{N}^o$, which is payoff equivalent to $\mathcal{N}^c$. First, we notice that continuation equilibria outside of $\mathcal{N}^o$ may only arise when a player makes an off equilibrium proposal (which, by construction, also has to be

[38] We need to make another modification to m whenever i proposes on path a continuation μ outside of the norm $\mathcal{N}^c$, which $-i$ is supposed to accept, and which is followed by a continuation μ' in the norm $\mathcal{N}^c$ (as it should, since the norm is closed). This sequence of moves is replaced by i directly proposing μ' and having it accepted by $-i$. The modified profile is also an equilibrium, as is easily checked. In fact, any SPE of the game can be turned into a payoff-equivalent "truthful" SPE of the game, i.e., one in which any proposal that is made and accepted *on the equilibrium path* is implemented. See Appendix G.

[39] At this point, we do not know yet that the proposal is $\mathcal{N}^o$-credible. We only know that it is $\mathcal{N}^c$-credible. However, the norm $\mathcal{N}^o$ that we are constructing will be payoff equivalent to $\mathcal{N}^c$ and hence have the same credible proposals.

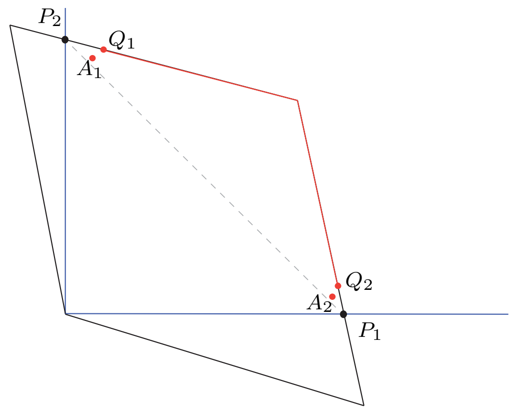

Figure 7: Norm construction with NFD (payoffs)

credible) which is accepted by the other player. Thus, $\mathcal{N}^o$ is an open norm. By construction, each element of $\mathcal{N}^o$ corresponds to exactly one element of $\mathcal{N}^c$, which yields the same expected payoff. Therefore, the norms are payoff equivalent. As observed earlier, this implies that they have the same set of credible proposals. This, in turn, implies that any Pareto-improving, credible proposal of $\mathcal{N}^o$ that is accepted is played and, hence, that $\mathcal{N}^o$ is stable.

C Proofs of Section 6

Notation: throughout the analysis, for any payoff vector X achieved by some SPE of $\mathcal{N}$, we will denote by $X^{\mathcal{N}}$ the corresponding SPE.

C.1 Proof of Theorem 5 (Sufficient Conditions)

Consider two feasible Pareto points, Q_1 and Q_2, lying at an arbitrarily small but strictly positive distance from P_2 and P_1, respectively, and illustrated by Figure 7. It suffices enough to show that for any ε small enough, there exists a forgiving stable norm $\mathcal{N}$ which includes Q_1 and Q_1 as equilibrium payoffs, that is, norm has elements $Q_1^{\mathcal{N}}$, $Q_2^{\mathcal{N}}$. By public randomization, this will imply that this norm can also be made to contain all payoffs above the segment $[Q_1, Q_2]$. The argument below focuses on the case in which P_2 and P_1 are determined by the minmax payoffs, which is the harder one.[40]

We construct a norm which continuation payoffs just after the public randomization stage (before the action stage) consist of the Pareto frontier contained between Q_1 and Q_2 and of two additional points, A_1 and A_2, respectively lying within ε-proportional distance from Q_1 and Q_2, as indicated on Figure 7. We describe the implementation of $A_1^{\mathcal{N}}$ and $Q_1^{\mathcal{N}}$; $A_2^{\mathcal{N}}$ and $Q_2^{\mathcal{N}}$ have a symmetric implementation.

While Q_1 is taken as given, the location of A_1 depends on ε, and is determined by the following conditions

$$\pi_1(A_1) = \pi_1(Q_1) - K\varepsilon$$

$$\pi_2(A_1) = \pi_2(Q_1) - L\varepsilon, \tag{14}$$

for constants K and L which will be determined ulteriorly.

[40] If, say, $\pi_1(P_2) > \underline{v_1}$, it suffices to set $Q_1 = P_2$ in our construction and use it as as the best proposal for player 2.

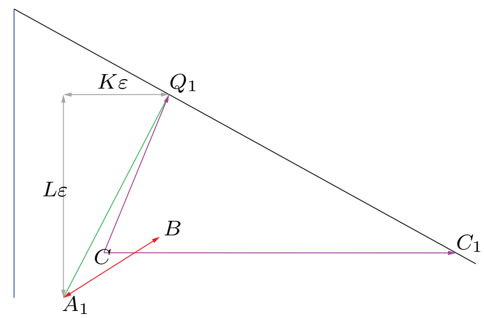

Figure 8: Norm construction with NFD (implementation)

To implement $A_1^{\mathcal{N}}$, players are prescribed to minmax each other. The continuation payoff B after the action stage is a function of the players' realized actions, a_1 and a_2: $B = B(a_1, a_2)$. The implementation is illustrated by Figure 8. For any action a_i of player i the continuation payoff $\pi_i(B(a_i, a_j))$ does not depend on a_j.

Given that player 2 has minmaxed player 1, let $Eu_1(a_1)$ denote 1's expected payoff for the period, as a function of his chosen action, a_1. 1's continuation payoff, $\pi_1(B(a_1, a_2))$, satisfies the promise-keeping condition

$$\pi_1(A_1) = \varepsilon Eu_1(a_1) + (1 - \varepsilon)\pi_1(B(a_1, a_2)).$$

A similar relation holds for 2's continuation payoff. By appropriately choosing players' continuation payoffs $B(a_1, a_2)_{(a_1,a_2)\in\mathcal{A}}$, the construction can make players indifferent between taking *any* action in the game.

Moreover, if the constant K appearing in (14) is large enough, then for any action profile (a_1, a_2), one necessarily has $\pi_1(B(a_1, a_2)) < \pi_1(Q_1))$.[41]

Consider any of the continuation payoffs $B(a_1, a_2)_{(a_1,a_2)\in\mathcal{A}}$ after the action stage—henceforth referred to as 'B' for simplicity. B is a weighted average of three continuation payoffs corresponding to the following events: player 1 makes a proposal, player 2 makes a proposal, no one makes a proposal. Let C denote the continuation payoff in case no one makes a proposal (this payoff is computed before the public randomization taking place in the following period).

For the norm to be forgiving, any rejected proposal results in payoff C. This implies that if player 1 gets to make a proposal, in equilibrium he proposes the element with a Pareto-efficient payoff C_1 which gives 2 her default value $\pi_2(C)$, making player 2 to accept the proposal in equilibrium.

The situation is different if player 2 gets to make a proposal. $B^{\mathcal{N}}$ gives player 1 a lower payoff than $Q_1^{\mathcal{N}}$, and player 2 is prescribed to propose an element $Q_1^{\mathcal{N}}$, which achieves her highest payoff in the norm and also gives player 1 a higher payoff than $C^{\mathcal{N}}$ does.

As shown on Figure 8, at element $B^{\mathcal{N}}$ if player 1 gets a chance to make a proposal, he proposes $C_1^{\mathcal{N}}$, if 2 gets a chance to make a proposal, she proposes $Q_1^{\mathcal{N}}$. B is thus a weighted average of C, C_1 and Q_1. Given any point B, one can find a default option C such that B is indeed the right weighted average, given the probabilities of proposal for each player.

We will verify at the end of this proof that the constants K and L from (14) may be chosen so that C lies to the right of the line (A_1, Q_1). If this is true, $C^{\mathcal{N}}$ may be implemented, before public randomization, as a weighted average of $A_1^{\mathcal{N}}$, $Q_1^{\mathcal{N}}$, and $Q_2^{\mathcal{N}}$.

[41] Indeed, the distance between A and $B(a_1, a_2)$ is proportional to ε, with a coefficient bounded above by the highest absolute value of the payoff of the stage game.

The remaining element of interest, $Q_1^{\mathcal{N}}$, is implemented as follows: players are prescribed to choose the pure-strategy Pareto-efficient payoff northwest of Q_1. If 1 deviates in action, the continuation payoff jumps to B; if 2 deviates, it jumps to the analog of B near Q_2. Players are incentivized to play as prescribed as long as $\frac{\pi_1(Q_1)-\pi_1(B)}{\varepsilon}$ is large enough. This is achieved by judiciously choosing the constants K and L arising in (14), as explained next.

Determination of the constants K and L

First, we observe that for K large enough, the threat of jumping to continuation $B^{\mathcal{N}}$ is enough to incentivize player 1 to play as prescribed in the implementation of $Q_1^{\mathcal{N}}$. We fix such a K—this choice is independent of ε. We now show that for L big enough, for any realization of B (which depends on which actions players choose while implementing $A^{\mathcal{N}}$), the point C will lie to the right of line A_1Q_1, as mentioned earlier.

Since a player's probability of proposal and the distance from B to the Pareto line are both proportional ε, the distance between B and C must be proportional to ε^2. Therefore, if we can show that each continuation point $B(a_1, a_2)$ lies to the right of the line A_1Q_1, at a strictly positive ε-proportional distance, so does the point C, for sufficiently small ε.

The points $B(a_1, a_2)$ are constructed by promise-keeping conditions. Let B^* denote the continuation payoff, out of all continuations $B(a_1, a_2)$, which gives the lowest payoff to player 1 and the highest payoff to player 2. B^* corresponds to the highest value $Eu_1(a_1)$ out of all actions a_1 and to the lowest value $Eu_1(a_2)$ out of all actions a_2. It suffices to show that B^* lies to the right of A_1Q_1. We recall the promise-keeping conditions

$$\pi_1(A_1) = \varepsilon Eu_1(a_1) + (1-\varepsilon)\pi_1(B^*)$$

$$\pi_2(A_1) = \varepsilon Eu_2(a_2) + (1-\varepsilon)\pi_2(B^*)$$

or, equivalently,

$$[\pi_1(A_1) - \pi_1(B^*)] = \varepsilon[Eu_1(a_1) - \pi_1(B^*)]$$

$$[\pi_2(A_1) - \pi_2(B^*)] = \varepsilon[Eu_2(a_2) - \pi_2(B^*)].$$

The ratio of the absolute values of the right-hand sides in the two equations above, $|\frac{Eu_2(a_2)-\pi_2(B^*)}{Eu_1(a_1)-\pi_1(B^*)}|$, determines the tangent of the angle of the vector A_1B^* above the horizontal. Since B^* is at an ε-distance from Q_1, this ratio simplifies to $|\frac{Eu_2(a_2)-\pi_2(Q_1)}{Eu_1(a_1)-\pi_1(Q_2)}|$, plus ε-terms which can be ignored.

Player 1 cannot obtain a higher payoff than his minmax $\underline{v}_1$ (as player 2 is minmaxing him), and player 2 cannot obtain a lower payoff than her lowest possible payoff in the game, which we denote as $\underline{v}$. Therefore, the angle of the vector A_1B^* above the horizontal is no higher than $|\frac{\underline{v}-\pi_2(Q_1)}{\underline{v}_1-\pi_1(Q_2)}|$, a finite value independent of L and ε.

The tangent of the angle of the line (A_1Q_1) above the horizonal is equal to $\frac{L}{K}$. By choosing L high enough, this ratio exceeds twice the ratio $|\frac{\underline{v}-\pi_2(Q_1)}{\underline{v}_1-\pi_1(Q_2)}|$. This guarantees that the vector A_1B^* lies strictly to the right of the line (A_1Q_1), as desired.

There remains to check that the norm satisfies all the conditions of Theorem 5. First, both players are incentivized to propose as prescribed: player 1 proposes the best available option for him, given the default option C. If player 2 wants to improve upon Q_1, she has to propose a continuation which gives her at least η more than her on-path continuation payoff. For ε small enough, however, the only proposals that would achieve this would have to give player 1 less than $\pi_1(C)$, and would therefore be rejected. Second, the continuation payoff, C, is the same when a proposal is rejected, regardless of the identity of the proposer and the nature of the proposal. The norm is thus forgiving. Finally, the point Q_1 is a continuation of the norm both after and before the public randomization, as desired.

C.2 Proof of Theorem 4 (Necessary Conditions)

Consider a forgiving stable norm $\mathcal{N}$. For simplicity, we assume that at each stage of the game—before the action stage, before the proposal stage, and before the public randomization stage—there exist equilibria in the norm with respective payoff vectors A, B, and C, that yield the maximal value of ρ at the corresponding stage.[42] Let α denote the (possibly mixed) action profile corresponding to the first-period play implementing element $A^{\mathcal{N}}$—the continuation before the action stage, and let $v(\alpha)$ denote the expected current payoff resulting from α. Since $\rho(v(\alpha)) \leq \bar{\rho}$, we necessarily have

$$\rho(A) \leq \varepsilon\bar{\rho} + (1-\varepsilon)\rho(B)$$

Point B, which is a continuation payoff before the proposal stage, is the weighted average of the continuation payoffs following accepted proposals, and of the default option. When a player—player 1, say—gets a chance to make a proposal, the expected continuation payoff must lie within at most an $\sqrt{\varepsilon}$-distance from the Pareto line. Otherwise, player 1 could propose a Pareto point which increases both players' payoffs by a value proportional to $\sqrt{\varepsilon}$, and is an equilibrium lying above the minmax.[43] This proposal would then be accepted by player 2 and would be a profitable deviation for player 1. Therefore, if a player gets a chance to make a proposal, which happens with probability $q\varepsilon$, the resulting continuation cannot have a positive value of ρ that exceeds $\sqrt{\varepsilon}$. When no one makes a proposal, the continuation payoff is dictated by the default continuation, whose value of ρ is at most $\rho(C)$. This implies that

$$\rho(B) \leq q\varepsilon \times \sqrt{\varepsilon} + (1-q\varepsilon)\rho(C).$$

Finally, since C is a convex combination of payoffs, obtained by public randomization, of equilibrium payoffs before the action stage whose maximal ρ-value is achieved by A,

$$\rho(C) \leq \rho(A).$$

Combining the above inequalities and getting rid of second-order ε terms shows Theorem 4.

D Comparative statics

Consider any $q > q'$ and any norm $\mathcal{N}$ that is stable for q. We will show the existence of a norm $\mathcal{N}'$, stable for q' and payoff-equivalent to $\mathcal{N}$, which implies that all payoffs implemented by $\mathcal{N}$ are q'-renegotiation-proof for the lower value of q'.

In the new norm $\mathcal{N}'$, any payoff A achieved by $\mathcal{N}$ before the action stage is implemented using the same mixed actions and the same subsequent continuations as prescribed by$\mathcal{N}$. Consider now any vector payoff B, calculated before the proposal stage, implemented by some equilibrium $B^{\mathcal{N}}$ of $\mathcal{N}$. $B^{\mathcal{N}}$ is a mixture of three continuation equilibria: $C_1^{\mathcal{N}}$, which arises when 1 gets a chance to make a proposal and is calculated after the proposal stage; $C_2^{\mathcal{N}}$ which arises if 2 gets to make a proposal; and $C^{\mathcal{N}}$, which arises if no one gets to make a proposal.

With the new negotiation factor q', B is implemented as follows: players are prescribed to make exactly the same proposals (with the same prescribed punishments if someone made an off-path proposal). For B to still to be the weighted average of the continuations occurring after the three proposal events, we change the

[42] If the supremum values are not achieved, the proof can be easily adjusted by taking appropriate limits.

[43] With the more permissive concept of an η-stable norm, the continuation payoff has to lie within a distance of $\sqrt{\varepsilon} + \eta$ from the Pareto line. Otherwise player 1 could make a proposal which gives him η more, and gives player 2 $\sqrt{\varepsilon}$ more than the continuation payoff.

continuation payoff in case no proposal is made: the new continuation payoff in this case, C', has to lie on the line between B and C. The new continuation $C'^{\mathcal{N}}$ can be implemented by using a public randomization, as it lies in the triangle (C, C_1, C_2). This, essentially, yields the new implementation.

There might be a problem, however, with this candidate implementation. One needs to make sure that players are correctly incentivized to make a proposal, when they get an opportunity to do so, rather than to conceal this opportunity. This is the case if $\pi_1(C_1) \geq \pi_1(C)$ and $\pi_2(C_2) \geq \pi_2(C)$, i.e., if each player gets at least as high a payoff when he makes a proposal as when he remains silent. When one moves point C to C', these incentives might get violated, and the construction above must be adjusted as follows.

The new continuation payoff when no proposal is made, C', lies in between C and B. Suppose that it violates 1's incentives to make his prescribed proposal: $\pi_1(C') > \pi_1(C_1)$. Since, in the old norm, we had $\pi_1(C_1) \geq \pi_1(C)$, such a violation is possible only if $\pi_1(C_2) > \pi_1(C_1)$. In this case, we modify the prescribed proposal for player 1 by moving point C_1 towards C_2. As this happens, the value of $\pi_1(C_1)$ increases and the value $\pi_1(C')$ decreases (to keep B the weighted average). When these values become equal, the incentives for player 1 to make a proposal start holding again. With the new continuation payoff C_1' for player 1's proposal and renewed continuation payoff in case of no proposal C'', player 1 is incentivized to make the prescribed proposal. One then can check that both new points can be implemented: the payoff C_1' lies between C_1 and C_2 and therefore can be implemented using public randomization, while point C'' lies within the triangle (C, C_1, C_2) and can therefore also be implemented.

The same procedure can be done for player 2. The modified continuation payoffs can be implemented using public randomization device. The new norm $\mathcal{N}'$ therefore has the same set of payoffs as the old norm $\mathcal{N}$ at any stage; and it is stable given the new value of q'.

E Proofs of Section 7 (Three or More Players)

E.1 Proof of Theorem 6

Since $v \in \mathcal{F}$, the usual Folk Theorem implies that for ε small enough v can be achieved by an SPE of the underlying repeated game. This SPE can be embedded into an equilibrium of the repeated game with renegotiation. In this equilibrium, no proposals are ever prescribed at any stage of the game. If any player i makes a proposal, other players are all prescribed to reject it, and the continuation payoff is player i's punishment, as in the underlying SPE. If only one player $j \neq i$ accepts i's proposal, the continuation is the punishment equilibrium for j. If at least two players accept the proposal, it is implemented. These prescriptions guarantee that any unilateral deviation in action, proposal, or acceptance decision is suboptimal.

E.2 Proof of Proposition 3

We fix one of the two procures and let $\mathcal{F}_k$ denote the set corresponding to the k-th step in the sequential reduction of the set $\mathcal{F}$ under this procedures. We first show that points on the relative Pareto frontier $\mathcal{P}(\mathcal{F}_k)$ of $\mathcal{F}_k$ are never removed by the procedure. Suppose, contrary to the claim, that some point $A \in \mathcal{P}(\mathcal{F}_k)$ was removed by the procedure. Then there would be a player i such that $\pi_i(A) < \underline{\pi_i}(\mathcal{F}_k)$. If A was prescribed as a punishment payoff for any proposal of player i, then for i's optimal proposal with payoff $C \in \mathcal{F}_k$, the punishment payoff A would not be credible as it is removed at the k-th step. That is, any $j \neq i$ has $\pi_j(A) < \pi_j(C)$. Since A lies on the Pareto frontier of $\mathcal{F}_k$, this means that $\pi_i(C) < \pi_i(A)$: C gives i a lower payoff than $\underline{\pi_i}(\mathcal{F}_k)$, which contradicts C's assumed optimality. One could simply prescribe both continuations to have C as their payoff vector, and this would give i a lower payoff than $\underline{\pi_i}(\mathcal{F}_k)$.

The optimal non-prescribed proposal for player i always lies on a Pareto frontier. Indeed, if i makes a non-prescribed proposal with payoff C, which is not Pareto-optimal, there must exist another point C', which is the payoff of another non-prescribed proposal for player i which Pareto dominates C, given that the number of prescribed proposals is finite[44] and given our maintained full-dimensionality assumption. The set $\mathcal{D}(C, \mathcal{F}_k)$ of possible punishment payoffs is strictly larger than the set $\mathcal{D}(C', \mathcal{F}_k)$, since the latter set gives every player $j \neq i$ a higher lower-bound on his payoff. This implies that the proposal C' gives player i a worst punishment payoff $\underline{\pi_i}(C', \mathcal{F}_k)$ at least as high as the proposal associated with payoff C. Therefore, player i can without loss always choose a point on Pareto frontier.

Since no point on the relative Pareto frontier of $\mathcal{F}$ is removed in the sequential reduction , the set of optimal proposals for any player i remains the same along the sequence. However, the set of possible punishments keeps decreasing at each step, which weakly increases, as a result, the minimal value $\underline{\pi_i}(\mathcal{F}_k)$ with k. (Recall that $\underline{\pi_i}(\mathcal{F}_k)$ is i's minimal payoff if he gets a chance to make a proposal). At each step, the set $\mathcal{F}_k$ is characterized by the n lower bounds of the players' payoffs $\{\pi_{min,i}(\mathcal{F}_k)\}_{i \in \{1,\ldots,n\}}$. These lower bounds are weakly increasing at each step, which implies that the procedure converges to a stable point.

E.3 Proof of Theorem 7

Necessity

Suppose that A lies outside of $\mathcal{V}_n$ and, for any small enough $\varepsilon > 0$, there exists a stable norm $\mathcal{N}(\varepsilon)$ such that $A \in \mathcal{N}(\varepsilon)$. A norm $\mathcal{N}(\varepsilon)$ has to satisfy the inequality (8) (if being used as an argument instead of $\mathcal{F}_n^k$), up to an ε-term. Let's limit ε to 0 and consider a sequence of norms $\mathcal{N}(\varepsilon)$ which payoff sets converge. This limit payoff set contains A and satisfies the inequality (8), which in turn means that A should have not been removed from any of the sets $\mathcal{F}_n^k$. However, that would make A an element in $\mathcal{V}_n$, a contradiction.

Sufficiency: $\mathcal{R}$

We first prove that any point in $\mathcal{R}$ is renegotiation-proof. Consider any point A with $\pi_i > \pi_i(P_{-i})$ for any i. As in the two-player case, one can find n points A_i such that for $j \neq i$ $\pi_j(A_i) = \pi_j(A)$ and $\pi_i(A_i) = \pi_i(P_{-i}) + \sqrt{\varepsilon}$. Each point A_i will give the lowest payoff for player i in the constructed norm $\mathcal{N}$. In the equilibrium of the norm $A_i^{\mathcal{N}}$ associated with payoff vector A_i, player i is being minmaxed. Since players other than i may have to use mixed strategies, this generates a set $\mathcal{B}$ of continuation payoffs, following the action stage, which depend on the realization of actions of players other than i. Any continuation $B \in \mathcal{B}$ is implemented as follows: if player i can make a proposal, he is prescribed to propose some continuation with payoff C; other players are prescribed to remain silent; in the absence of any proposal, the continuation returns to $A_i^{\mathcal{N}}$. As in the two-player case, one can guarantee (possibly using the public randomization), that the distance A_iC is of order $\sqrt{\varepsilon}$.

Since the Pareto frontier is connected, so is its truncation to points for which i's payoff lies above $\pi_i(P_{-i})$. One can therefore find a connected subset S_ε of the frontier consisting of all points giving, for each i, a payoff greater than or equal to $\pi_i(A_i) + K\varepsilon$, where K is a constant chosen large enough that players are incentivized not to deviate in actions.

Continuation equilibria with payoffs in S_ε are implemented in such a way that each player i gets at least $\pi_i(A_i) + K\varepsilon$ in all continuations. Players are prescribed to stay silent. Since each point of S_ε is Pareto-efficient, there are no unanimously improving proposals anyway. Moreover, using $A_i^{\mathcal{N}}$ as a punishment if i deviates in actions guarantees that such deviation would be suboptimal.

When implementing $A_i^{\mathcal{N}}$, players are already incentivized to follow the prescribed actions. If i wants to make a non-prescribed proposal, then by construction of S_ε there exists a continuation with a payoff Q_i in set S_ε which gives player i a lower payoff than C. Indeed, the lower bound for π_i at the set S_ε is $\pi_i(A_i) + K\varepsilon$, while $\pi_i(C) - \pi_i(A_i)$ is of order $\sqrt{\varepsilon}$.

Sufficiency: General Conditions

The proof is similar to the two-player case. For any point $A \in \mathcal{V}_s$ with $\pi_i > \pi_{min,i}(\mathcal{V}_s)$, consider the set of points $A_i \in \mathcal{V}_s$ such that for any i $\pi_i(A_i) = \pi_{min,i}(\mathcal{V}_s) + \sqrt{\varepsilon}$ and $\pi_{-i}(A_i) = \pi_{-i}(A)$. The points A_i have a

[44] See Appendix F

smaller i-th coordinate than A provided that ε is small enough. In addition, we have $\pi_i(A_j) - \pi_i(A_i) >> \sqrt{\varepsilon}$ for any $j \neq i$ without loss of generality.

We build a stable norm $\mathcal{N}$ such that A_i gives the lowest payoff to player i in the norm (calculated at the start of the period). At $A_i^{\mathcal{N}}$, player i is minmaxed. Since players other than i may have to mix their actions, we construct a set of continuations with payoffs $B \in \mathcal{B}$, corresponding to the observed actions of players $-i$. For any continuation equilibrium $B^{\mathcal{N}}$ associated with some payoff $B \in \mathcal{B}$, i is prescribed to make a proposal with some payoff vector C, and all other players are prescribed to remain silent. As with the two-player case, C can be assumed to lie at a distance of order $\sqrt{\varepsilon}$ from A_i. When implementing the equilibrium $C^{\mathcal{N}}$ associated with C, players are prescribed to follow a deterministic sequence of actions such that the continuation payoff remains within an ε-distance from C. Players are prescribed not to make any proposals.

The initial point A is also implemented by deterministic actions and no proposals. Moreover, each point in the orthant with lower bounds $\pi_i(A_i) + K\varepsilon$ is included in the norm $\mathcal{N}$ and implemented in such a way that $\pi_i > \pi_i(A_i) + K\varepsilon$: $A_i^{\mathcal{N}}$ is severe enough a punishment for i that it makes it suboptimal for him to deviate in action.

This norm can be shown to be generically stable. The only new issue concerns i's incentives to deviate in proposal. We have reduced (increased the lower bounds on payoffs) the initial set $\mathcal{V}_s$ by an order of $\sqrt{\varepsilon}$. The whole orthant defined by $\pi_i > \pi_i(A_i) + K\varepsilon$ for all i is part of the norm, but some points below are removed from the original set $\mathcal{V}_s$. As a result, the value $\underline{\pi_i}(.)$, which i can guarantee if having a chance to propose, can now be larger. Our goal is to show that, nevertheless, generically the value of $\underline{\pi_i}(.)$ is smaller than $\pi_i(C)$, and therefore player i is incentivized to propose $C^{\mathcal{N}}$.

When building a set $\mathcal{V}_s$ by sequentially removing payoffs with $\pi_{min,i}(.) < \underline{\pi_i}(.)$, the initial set of individually-rational payoffs gets reduced. If for player i the final value of $\pi_{min,i}(\mathcal{V}_s)$ is strictly larger than his minmax payoff $\underline{v_i}$, then the value of $\pi_i(A_i) - \underline{v_i}$ is of order ε^0. This means that the distance A_iC can be made of $\varepsilon^{\frac{1}{4}}$-order. At the same time, the set $\mathcal{V}_s$ (and, respectively, the value $\underline{\pi_i}(.)$) were changed by an order of $\sqrt{\varepsilon}$, guaranteeing that $\underline{\pi_i}(.) < \pi_i(C)$.

If player i's payoff $\pi_{min,i}(\mathcal{V}_s)$ equals to minmax $\underline{v_i}$, this means that i's payoff was not increased when building set $\mathcal{V}_s$. Put it differently, one can consider a hyperplane of the set $\mathcal{V}_s$ with $\pi_i = \underline{v_i}$, and find the maximum payoffs of other players $\overline{\pi_j}$, $j \neq i$ on that hyperplane. The $n-1$-dimensional payoff vector $\{\overline{\pi_j}\}_{j \neq i}$ cannot lie within an interior of $\mathcal{V}_s$ (otherwise, player i could make a proposal dominating $\{\overline{\pi_j}\}_{j \neq i}$ and thus guaranteeing himself a payoff higher than $\underline{v_i}$). When the set $\mathcal{V}_s$ is reduced by (an arbitrarily small) $\sqrt{\varepsilon}$-order, player i can gain incentives to make an off-path proposal, only if the vector $\{\overline{\pi_j}\}_{j \neq i}$ lies exactly on the Pareto frontier of $\mathcal{V}_s$. However, this possibility is not generic.

E.4 Proof of Proposition 4

Intuition. The sets $\mathcal{V}_s$ and $\mathcal{V}_n^q$—necessary conditions depend on q, hence the superscript—are both obtained from $\mathcal{F}$ by sequentially increasing the lower bounds on each player's payoff when he gets a chance to make a proposal. $\mathcal{V}_s$ is obtained by removing payoffs below $\underline{\pi_i}(\cdot)$ at each step, while $\mathcal{V}_n^q$ is obtained by removing payoffs below $\frac{n\underline{v_i}+q\underline{\pi_i}(.)}{n+q}$. When q goes to infinity, the sets of payoffs removed at each step of these procedures converge to each other. As we show below, this implies that $\mathcal{V}_n^q$ converges to the set $\mathcal{V}_s$ as q goes to infinity.

The set of sufficient conditions, $\mathcal{V}_s$, can be characterized by two sets of lower bounds for each player i: $\underline{\pi_i}(\mathcal{V}_s)$ is the lower bound on i's payoff when he gets a chance to make a proposal and $\pi_{min,i}(\mathcal{V}_s)$ is the lower bound for his payoff at the beginning of a period. $\mathcal{V}_s$ was constructed in such a way that $\underline{\pi_i}(\mathcal{V}_s) \leq \pi_{min,i}(\mathcal{V}_s)$.

To capture the above intuition, we first show by induction that $\mathcal{V}_s$ is the largest set $\mathcal{S}$ of individually rational payoffs whose Pareto frontier is equal to $\mathcal{P}(\mathcal{V})$ and such that $\underline{\pi_i}(\mathcal{S}) \leq \pi_{min,i}(\mathcal{S})$ for any i. Consider such a set $\mathcal{S}$. The sequence of sets $\mathcal{F}_s^k$ converging to $\mathcal{V}_s$ starts with $\mathcal{F}_s^0 = \mathcal{F}$, the set of all individually rational points. This implies that $\underline{\pi_i}(\mathcal{S}) \geq \underline{\pi_i}(\mathcal{F}_s^0)$, since $\mathcal{F}_s^0$ contains $\mathcal{S}$ and, hence, the set of punishments if i makes an unprescribed proposal is higher with $\mathcal{F}_s^0$ than with $\mathcal{S}$, resulting in a lower bound $\underline{\pi_i}$. We now show the induction hypothesis: if $\underline{\pi_i}(\mathcal{S}) \geq \underline{\pi_i}(\mathcal{F}_s^k)$, then the same condition holds for $k+1$. Due to the way the payoffs are cut at step k, one has for each i, $\pi_{min,i}(\mathcal{F}_s^{k+1}) = \max\{\pi_{min,i}(\mathcal{F}_s^k), \underline{\pi_i}(\mathcal{F}_s^k)\} \leq \underline{\pi_i}(\mathcal{F}_s^k)$, which does not exceed $\underline{\pi_i}(\mathcal{S}) \leq \pi_{min,i}(\mathcal{S})$. Since the lower bound $\pi_{min,i}(\mathcal{F}_s^{k+1})$ is lower than $\pi_{min,i}(\mathcal{S})$, the set $\mathcal{F}_s^{k+1}$ contains $\mathcal{S}$, and one has that $\underline{\pi_i}(\mathcal{S}) \geq \underline{\pi_i}(\mathcal{F}_s^{k+1})$. By induction, the limit set $\mathcal{V}_s$ contains $\mathcal{S}$.

Let $\mathcal{V}_n$ denote the limit of $\mathcal{V}_n^q$ as q goes to infinity. We wish to show that $\mathcal{V}_n = \mathcal{V}_s$. Consider the sequences $\{\mathcal{F}_n^{k,q}\}_{k=0}^{+\infty}$ resulting from the procedure applied, for any fixed q, to derive necessary conditions for this value of q. Due to the way points are removed at each step, it is easy to check that $\mathcal{F}_n^{k,q'} \subset \mathcal{F}_n^{k,q}$ whenever $q' > q$; by the same logic, it is straightforward to check that $\mathcal{V}_s$ is contained in $\mathcal{V}_n$. To prove the reverse inclusion, note for each q and i, we have $\pi_{min,i}(\mathcal{V}_n^q) \geq \frac{n\underline{v_i} + q\underline{\pi_i}(\mathcal{V}_n^q)}{n+q}$, as this inequality holds at each step k of the procedure. Taking the limit as q goes to infinity, the limiting set $\mathcal{V}_n$ must satisfy for each i $\underline{\pi_i}(\mathcal{V}_n) \leq \pi_{min,i}(\mathcal{V}_n)$. From the previous paragraph, this implies that $\mathcal{V}_s$ contains $\mathcal{V}_n$, which concludes the proof.

E.5 Proof of Proposition 5 (Sketch)

Consider for simplicity the case of three players: player 1 makes a proposal and player 2 responds first, followed by player 3. Depending on responding players' votes, there are four possible continuations, one of which is equal to 1's proposal and arises when 2 and 3 accept the proposal.

The ability to punish 2 for accepting player 1's proposal is constrained by the following issue: if 2 accepts the proposal, 3 will reject it only if the punishment for player 2 gives him at least the same payoff as 1's proposal, which will be implemented if he accepts it. This puts a lower bound on 2's punishment payoff, which is higher than the minmax $\underline{v_2}$.

As a result, 1's punishment for making an off-path proposal is also limited. Since fewer punishments are available, fewer equilibria are renegotiation-proof: sequential voting has more predictive power than simultaneous voting.

By nature of the arguments used to derive necessary and sufficient conditions, these conditions are characterized by upper orthants, even if players randomize their acceptance decision.

Since allowing only two continuations—as simple norms do with simultaneous voting—is a special case of the more numerous continuations allowed by sequential voting, it follows that simple stable norms have more predictive power than the stable norms obtained with sequential voting.

E.6 Proof of Theorem 8 (Sketch)

We construct a forgiving η-stable norm $\mathcal{N}$ as follows. The norm $\mathcal{N}$ includes all Pareto-efficient payoffs which lie at some arbitrary small, but ε-independent distance from the minmax values. The norm $\mathcal{N}$ also includes, for each player, a set of Pareto-inefficient elements used to build a punishment equilibrium for that player, all elements in each set lie within a distance of order ε from the Pareto-efficient elements of the norm. For each player i, there is a Pareto-inefficient payoff vector A_i which gives i his worst payoff in $\mathcal{N}$. The equilibrium $A_i^{\mathcal{N}}$ which achieves payoff A_i, together with its continuations, form the punishment set for player i, as described below.

If players were unable to make any proposal, one could implement payoff A_i as follows. Player i is being minmaxed, which may require other players to use mixed strategies. As described in earlier proofs, this results in a set $\mathcal{B}^1$ of continuation payoffs, (potentially) one for each observed action profile (these various continuations are needed to incentivize the minmaxing strategy). Each continuation payoff $B^1 \in \mathcal{B}^1$ is implemented by minmaxing player i, which again generates several continuation payoffs in the next period, with generic element denoted as B^2. Player i is minmaxed in this way for several periods. In each period i's continuation payoff, π_i, increases by an amount of order ε. One can compute the number T of periods needed to minmax player i, so that π_i exceeds $\pi_i(A_i)$ by a sufficiently high amount that i can be incentivized to play any action by the threat of returning to A_i. The value of T is independent of ε. After these T periods, each continuation payoff B^T can be implemented by playing a deterministic sequence of actions so that the continuation payoff always lies within some ε-proportional distance from B^T. This implementation is an equilibrium, since the payoff A_i prevents any deviation from player i, and any deviation by another player leads to an even larger drop in the continuation payoff of the deviator.

When proposals are re-introduced in the game, there will be changes in the implementation of $A_i^{\mathcal{N}}$, but these changes will be insignificant. After the first round of minmaxing player i, the resulting continuation payoff B^1 is calculated taken into account the possibility of proposals. That is, $B^{1\mathcal{N}}$ is the convex combination of some default option, $C^{1\mathcal{N}}$, if no one makes a proposal, and of proposals payoffs $C_i^{\mathcal{N}}$ for each player, which are chosen to be Pareto efficient elements of the norm $\mathcal{N}$. The distance between the payoffs B^1 and C^1 is of order ε^2—as explained the similar proofs seen earlier. In the next period, the continuation payoff before the actions will be C^1 (instead of B^1, in the previous paragraph). Therefore, if one repeats minmaxing player i for T periods, the resulting continuation payoff compared to the case with no proposals, will differ by an amount of order ε^2, which is negligible as ε becomes arbitrarily small. As players become arbitrarily patient, the modified implementation of A_i, based on minmaxing player i for T periods and then choosing a deterministic sequence of actions, will thus be an equilibrium even with the possibility of proposals.

Finally, the payoff A_i (and, therefore, all the default continuation payoffs C's) can be chosen so as to lie within some distance $K\varepsilon$-distance from the Pareto line. With ε small enough, no player can make an off-equilibrium proposal that would give him a payoff of at least η more than the equilibrium proposal, while keeping all other players at least at well as off as with the default payoff C. Therefore, the constructed norm is η-stable. Using initial public randomization, one can then include in the norm any point in the convex hull $Co(\mathcal{P}')$, which concludes the proof.

F Restricting plans

We refine steps b1), b2), and c) as follows. There exist constants k and k' with $k' > 3k > 0$ [45]such that

b1) for any z, $\mathbf{a}$, and i, the support of $\mu_i[m](z, \mathbf{a})$ contains at most k proposals;

b2) $-i$ is prescribed to surely accept (reject) any on-path (off-path) proposal;

c) for any z and $\mathbf{a}$, the set $\mathcal{M}^{+1}[m](z, \mathbf{a}, \mu)$ of possible continuation plans has at most k' elements and contains, in this order: the current proposal, μ; a default option (used if no proposal was made); the set of possible on-path proposals for each player (at most $2k$ elements); a finite number of alternative continuation plans, pre-specified by m as a function of z and $\mathbf{a}$, which may be used as punishments against deviating proposers.

The condition $k' > 3k$ guarantees that the list of continuation plans can indeed be included in $\mathcal{M}^{+1}[m](z, \mathbf{a}, \mu)$ as long as there are no more than $k' - 2k - 2 > 0$ alternative continuations plans.

To specify the continuation plan selected for the next period, we partition the proposal stage according to i) who (if any) got to make a proposal, ii) in the event that a proposal was made, whether the proposal was prescribed (out of at most k possibilities) or off-path, iii) whether the proposal was accepted or rejected. As is readily checked, this partition consists of at most $\hat{k} = 1 + 2 \times (k + 1) \times 2$ elements. In addition, we also compare the current proposal, μ (or the default option in case of no proposal), to each of the (at most) $k' - 1$ other feasible continuation plans. Specifically, we consider, for each player and alternative continuation, which of the current proposal or the alternative continuation plan gives the higher payoff to that player. Allowing three comparison outcomes ($\{=, <, >\}$) for each player yields $3^{2(k'-1)}$ combinations.

Let $\mathcal{E}$ denote the set of possible *events* at the proposal stage, and $\bar{k} = \hat{k} \times 3^{2(k'-1)}$ denote its cardinality. A *choice rule* g determines, for each event, the continuation plan—an element of $\mathcal{M}^{+1}[m](z, \mathbf{a}, \mu)$—for the

[45]The plan for $n > 2$ players is described similarly, with increased constants k and $k' > (n + 1)k > 0$.

next period. The choice rule $g[m](z, \mathbf{a})$ is pre-specified by m as a function of z and $\mathbf{a}$; and it is an element of a finite set $\mathcal{G}$ with a cardinality $(k')^{\bar{k}}$.

In summary, each plan prescribes, for each realization of z and $\mathbf{a}$, a successor (a plan for the following period) as a function of the choice rule (an element from $\mathcal{G}$) and of the event which occurred. Letting $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$ denote the (finite) set of action profiles, this stage adds $\mathcal{M}_{prop} = (k'^{\bar{k}|\mathcal{G}|})^{[0,1]\times\mathcal{A}}$ elements to the prescription, which has the cardinality of the power set of the continuum.[46]

For any set M, let $T(M)$ denote the set of plans obtained from the recursive construction above, choosing proposals from M and continuation plans from M. $T(M)$ is the Cartesian product of the prescriptions obtained at each step of the construction. We now show that there exist sets M's for which $T(M)$ has the same cardinality as M. These sets have cardinality $\beth_2 = |2^{\mathfrak{c}}|$, i.e., the cardinality of the power set of the continuum, which is also the cardinality of $\mathbb{R}^{\mathbb{R}}$—the set of all real-valued functions over $\mathbb{R}$—see, e.g., Forster (1995).[47]

Proposition 7 *If M has cardinality $|2^{\mathfrak{c}}|$, then so does $T(M)$.*

This implies that any set M with cardinality $|2^{\mathfrak{c}}|$ is in bijection with $T(M)$.[48] The set $\mathcal{M}$ of plans is then structured as follows: let ϕ denote the bijection between $2^{\mathfrak{c}}$ and $T(2^{\mathfrak{c}})$. To any $m \in 2^{\mathfrak{c}}$, we can associate the plan, defined recursively through $\phi(m)$, which specifies mixed strategies for each realization z of the randomization device, proposals as a function of z and of the observed action profile $\mathbf{a}$, and continuation plans which are elements of $2^{\mathfrak{c}}$. Thus, each element of $\mathcal{M} = 2^{\mathfrak{c}}$ specifies a plan, which is defined recursively.
Proof. Given the cardinality of $\mathcal{M}$, for each part from a) to c) in the prescription of a plan, one must find the cardinality added to the plan choice from that part. Since these parts are related by a Cartesian product and we are dealing with infinite sets, the cardinality of the Cartesian product coincides with the cardinality of the largest component of the product.[49] Part a) maps real-line outcomes (randomization device) into mixed strategies over $\mathcal{A}$ for both players, which is a subset of $|\mathcal{A}|^{\mathbb{R}}$. Since $\mathbb{R}^{\mathbb{R}}$ has the cardinality of $2^{\mathfrak{c}}$, this has the same cardinality as $\mathcal{M}$. Part b1) maps any outcome $(z, \mathbf{a})$ and proposer into k possible plans in $\mathcal{M}$ and a distribution over these plans, and thus has $\mathcal{M}^{k\times[0,1]\times\mathcal{A}} \times \Delta_k(\mathbb{R})^{[0,1]\times\mathcal{A}}$ elements where $\Delta_k(\mathbb{R})$ is the probability simplex in $\mathbb{R}^k$. Again, this set has the cardinality of $\beth_2$ since $\mathcal{M}^{\mathbb{R}}$ is equivalent to $\mathbb{R}^{\mathbb{R}}$. Part b2) does not add any cardinality since the prescription is to accept only prescribed proposals. Similarly, part b3) does not add any cardinality. Part c) adds the choice rule specification with a cardinality $\mathcal{G}^{[0,1]\times\mathcal{A}} = 2^{\mathfrak{c}}$ and the set $\mathcal{M}^{+1}$ of continuations with the cardinality of $\mathcal{M}^{k'\times[0,1]\times\mathcal{A}}$, which is the same as $\mathcal{M}$'s, concluding the proof. ■

Special cases used in the analysis

1) Babbling equilibria. These are the SPEs of the underlying repeated game, ignoring any renegotiation of continuation play. Babbling equilibria are captured by the choice rule which imposes the default continuation (the second element of $\mathcal{M}^{+1}[m](z,\ a, \mu)$), no matter what happens during the proposal stage.
2) Stable norms. We can use choice rules such that if a proposal is accepted it is used as the continuation plan

[46]Since $[0, 1]$ has the cardinality of the continuum, $\mathcal{M}_{prop}$ has the same cardinality as the set of functions which maps real numbers into a finite set, which is the same as the cardinality of $2^{\mathfrak{c}}$, the power set of the continuum.

[47]Perhaps a simple way to see this is the chain $|\mathbb{R}^{\mathbb{R}}| = |(2^{\aleph_0})^{2^{\aleph_0}}| = |2^{\aleph_0 2^{\aleph_0}}| = |2^{2^{\aleph_0}}| = |2^{\mathbb{R}}|$, where the third inequality holds because $\mathbb{N} \times \mathbb{R}$ is no larger than $\mathbb{R}$.

[48]By definition, two sets have the same cardinality if there exists a bijection between them. See, e.g., Kuratowski and Mostowski (1968).

[49]In particular, for $\beth_0 = \mathbb{N}$, we have $|\mathbb{N} \times \mathbb{N}| = |\mathbb{N}|$. The same is true for higher beth numbers, such as $\beth_1 = \mathbb{R}$ and $\beth_2 = 2^{\mathfrak{c}}$.

3) Forgiving norms. The choice rule specifies that if a proposal is rejected, the default continuation is played regardless of the proposal and the identity of the proposer.

G Extensions: truthful equilibria and asymmetric proposal probabilities

Truthful equilibria

Appendix F restricts all plans to prescribe a bounded number of proposals and a bounded of possible continuations, and to prescribe that only on-path proposals be accepted. In particular, any equilibrium prescribed by a plan must be *truthful*: any on-path proposal is always accepted and implemented.

This result is general: we now show that for any well-defined, possibly more permissive plan prescriptions and associated stable norm $\mathcal{N}$, there is a payoff-equivalent stable norm, using the same concept of stability, that consists entirely of the more restrictive plans used in this paper.

The argument applies to any number of players and any concept of norm stability: consider any stable norm $\mathcal{N}$ which uses plans with milder restrictions. When some player i gets a chance to make a proposal, he can make any number of proposals in equilibrium, the expectation of which is some continuation payoff C. We alter the equilibrium by prescribing player i to make only one proposal with payoff C. The altered equilibrium prescribes all other players to accept the proposal and C to be implemented regardless of the acceptance decision. The payoff C can be implemented using public randomization.

If i deviates and proposes a Pareto improvement relative to C, everyone is prescribed to reject it. With two players, the new norm prescribes to have the same rejection continuation as in the original norm, $\mathcal{N}$. The incentives to accept the proposal have not been changed by the transformation, so the other player everyone is incentivized to reject an off-path proposal. With more than two players rejecting the off-path proposal is an equilibrium.

When the norm is simple and players vote for the payoff-improving proposal, as in Assumption 1, player i is still prescribed to propose C. It remains to be shown that if one has enough continuations to incentivize the players to reject any off-path proposal in a simple norm with more permissive plans, it is possible to use only $n-1$ of them. If one considers all possible off-path proposals of player i, for each such a proposal there is at least one player $j \neq i$ who rejects it. This means that for any off-path proposal rejected by player j there is an element $C_j^{\mathcal{N}}$ in the norm which gives j at least as high payoff as the off-path proposal. Any norm can be expanded to have all the payoffs in its closure by properly designing the prescriptions. Thus, one can consider the element $C_j^{\mathcal{N}}$ with the highest payoff for j and use it as a punishment; and have at most $n-1$ continuations in total for the simple norm to be stable.

The equivalence with more permissive plans also holds if one does not impose the stability requirement. Player i is prescribed to propose $C^{\mathcal{N}}$ and the continuation is prescribed to be $C^{\mathcal{N}}$ regardless of acceptance decision. If player i makes another proposal, a default option is always played. This makes accepting $C^{\mathcal{N}}$ and rejecting any other proposal an equilibrium.

Asymmetric proposing probabilities

It is easy to extend the analysis to a protocol in which one of the players has a higher probability factor q_i of proposal than the other player. The sufficient conditions are unchanged in this setting, but the necessary conditions become tighter for the player whose proposal probability is higher, which translates into a higher minimal guaranteed payoff for that player, across all renegotiation-proof equilibria. To see this starkly suppose that $\underline{v}_1 < \pi_1(P_2)$ and $\underline{v}_2 < \pi_2(P_1)$ (configuration (a) in Figure 2), so that renegotiation potentially benefits both players, compared to the minmax payoffs, and consider the case in which 1 can make frequent proposals while 2 never gets a chance to make a proposal (i.e., q_1 is arbitrarily large while $q_2 = 0$). Then, 2's minimal guaranteed renegotiation-proof payoff collapses to her minmax payoff, while 1 is guaranteed to get a payoff of at least $\pi_1(P_2)$. More generally, player i's minimal payoff, given by (2), is calculated using the probability q_i that he gets an opportunity to make a proposal, and is independent of the other player's probability of getting that opportunity. As q_i increases, player i's guaranteed continuation payoff increases as well, and vice versa.

H Observability of mixing strategies

We have assumed throughout the paper that when a player randomizes across several actions or proposals, only the outcome of this randomization is observed by the other player. In particular, players' continuation values cannot directly depend on their choice of mixed strategy. Our results do not change if instead we assume that mixed strategies are observable. For sufficient conditions, this fact is straightforward because our construction is clearly compatible with players observing more information. For necessary conditions, payoff lower bounds were computed using only that any player can guarantee himself at least his minmax payoff during the action stage and at least some particular payoff during the proposal stage which satisfies the responder. These lower bounds do not change when mixing is observable.

The observability of mixed strategies does affect, however, the set of weakly renegotiation-proof (WRP) equilibria defined by Farrell and Maskin (1989), as follows. An SPE σ is *weakly-renegotiation proof* if there do not exist continuation equilibria σ^1, σ^2 of σ such that σ^1 strictly Pareto dominates σ^2. If a payoff vector arises as players' continuation payoff following some history of a WRP equilibrium, we will also say that these payoffs are WRP.

Assuming that mixing probabilities are observable, Farrell and Maskin obtained a sufficient condition for any feasible payoff to be WRP in the context of two-player repeated games. To formulate this condition, they define $c_i(\alpha) = max_{a'_i} \pi_i(a'_i, \alpha_{-i})$ as the *cheating* payoff of player i when he chooses a best response to the (mixed) action α_{-i}, and establish the following result.

PROPOSITION 8 *Let $\pi = (\pi_1, \pi_2)$ denote a feasible payoff. If there exist (mixed) action pairs $\alpha^i = (\alpha_1^i, \alpha_2^i)$ (for $i = 1, 2$) such that $c_i(\alpha^i) < \pi_i$, and $\pi_{-i}(\alpha^i) \geq \pi_{-i}$, then the payoff π is WRP if players are sufficiently patient.*

We now present an example with a Pareto-efficient payoff that satisfies the requirement of the above proposition, but cannot be WRP if mixing probabilities are unobserved.[50] The stage game is as follows:

9,-4	-2,-4
-2,-4	9,-4
0,8	8,0

The payoffs $(0, 8)$ and $(8, 0)$ are Pareto efficient and the minmax values of players are $\underline{v}_1 = \frac{72}{19} = 4 - \frac{4}{19}$ and $\underline{v}_2 = -4$, as is easily checked (for $\underline{v}_1$, 2 mixes so as to make 1 indifferent between the first and last rows). The Pareto-efficient, individually-rational point A with payoffs $\pi_1 = 4 + \frac{1}{100}, \pi_2 = 4 - \frac{1}{100}$ satisfies the premise of Proposition 8 with α^1 defined by player 1 choosing the last row and player 2 mixing equally between two columns and α^2 defined by player 1 choosing the first row and player 2 choosing the first column. We will

[50]The definition of WRP is the same as before. The only difference is that equilibrium strategies now depend only on the history of realized actions rather than on the history that included mixed strategies.

nevertheless show that A cannot be WRP for low enough ε (arbitrarily patient players), even if players have access to a public randomization device.

Counter-Example 1 *With unobservable mixed strategies, A is not WRP.*

Suppose, by way of contradiction, that A is the continuation payoff of some WRP equilibrium σ, and consider the payoff vector A' corresponding to player 1's lowest payoff and, hence, player 2's highest payoff among all continuations payoff of σ before public randomization.[51] When implementing A', depending on the outcome of public randomization, player 2 plays a pure strategy with some probability β and mixes with the complement probability $(1-\beta)$.

Since A' gives 1 his lowest possible payoff, when implementing A' player 1 cannot get a period payoff higher than $\pi_1(A')$, even if he always plays a stage-game best response. Otherwise, the promise-keeping constraint would have to prescribe a continuation giving 1 a payoff lower than $\pi_1(A')$. If player 2 chooses a pure strategy, player 1 can guarantee himself a payoff of at least 9. If player 2 chooses a mixed strategy, player 1 can guarantee himself his minmax payoff of $\frac{72}{19}$. This puts an upper bound on the probability β of player 2 choosing pure strategy:

$$\pi_1(A') \geq 9\beta + \frac{72}{19}(1-\beta) \tag{15}$$

The continuation payoff $\pi_2(A')$ of player 2 is a mixture between continuation payoffs $\pi_{2,pure}$ and $\pi_{2,mixed}$ conditional on her playing a pure and a mixed strategy:

$$\pi_2(A') = \pi_{2,pure}\beta + \pi_{2,mixed}(1-\beta) \tag{16}$$

If 2 mixes between the two columns, by indifference any choice has to give her the same payoff $\pi_{2,mixed}$. Player 2 cannot get more than 0 when choosing the right column, and the continuation payoff from the next period onward cannot exceed $\pi_2(A')$. This puts an upper bound on 2's continuation payoff:

$$\pi_{2,mixed} \leq 0 \times \varepsilon + (1-\varepsilon)\pi_2(A')$$

Similarly, since 2 cannot get a payoff higher then 8, we have

$$\pi_{2,pure} \leq 8 \times \varepsilon + (1-\varepsilon)\pi_2(A')$$

Combining these inequalities with (16) yields

$$\pi_2(A') \leq 8 \times \varepsilon\beta + (1-\varepsilon)\pi_2(A').$$

Rearranging, we get the following lower bound for β:

$$\pi_2(A') \leq 8\beta \tag{17}$$

Combining (15) and (17) yields $\pi_2(A') \leq \frac{8}{99}(19\pi_1(A') - 72)$. Since $\pi_2(A) \leq \pi_2(A')$ and $\pi_1(A) \geq \pi_1(A')$ this implies that

$$4 - \frac{1}{100} \leq \frac{8}{99}\left(19\left(4 + \frac{1}{100}\right) - 72\right),$$

which is false (the right-hand side is approximately equal to 0.34) and yields the desired contradiction.

[51]Since σ is WRP, 1's lowest continuation payoff is achieved for 2's highest continuation payoff. The proof can be easily adjusted if σ's payoff extrema are not achieved.