

Voting in Collective Stopping Games

P. Jean–Jacques Herings and Arkadi Predtetchinski*

January 28, 2013

Abstract

At each moment of time an alternative from a finite set is chosen by some dynamic process. Players observe the alternative selected and sequentially cast a yes or a no vote. If the set of players casting a yes–vote is decisive for the alternative in question, the alternative is accepted and the game ends. Otherwise the next period begins. We refer to this class of problems as collective stopping problems. Collective choice games, quitting games, and coalition formation games are particular examples that fit nicely into this more general framework.

When an appropriately defined notion of the core of this game is non–empty, a stationary equilibrium in pure strategies is shown to exist. But in general, stationary equilibria may not exist in collective stopping games and if they do exist, they may all involve mixed strategies.

We consider strategies that are pure and action–independent, and allow for a limited degree of history dependence. Under such individual behavior, aggregate behavior can be conveniently summarized by a collective strategy. We consider collective strategies that are simple and induced by two–step game–plans with a finite threshold. The collection of such strategies is finite. We provide a constructive proof that this collection always contains a subgame perfect equilibrium. The existence of such an equilibrium is shown to imply the existence of a sequential equilibrium in an extended model with incomplete information.

Collective equilibria are shown to be robust to perturbations in utilities. When we apply our construction to the case with three alternatives exhibiting a Condorcet cycle, we obtain immediate acceptance of any alternative.

*We would like to thank Paulo Barelli, Francis Bloch, John Duggan, János Flesch, Jean Guillaume Forand, Julio González Díaz, Harold Houba, Tasos Kalandrakis, and William Thomson for very their helpful comments on the paper.

JEL CLASSIFICATION CODE: C62, C72, C73, C78.

KEYWORDS: Voting, Collective choice, Coalition formation, Subgame perfect equilibrium, Stationary equilibrium.

1 Introduction

In this paper, we study a class of games with the following features. At each moment in time, an alternative is selected by means of some dynamic process, and the chosen alternative is put to a vote. If the set of players who voted in favor of the alternative is decisive for the alternative in question, the alternative is implemented, each player receives a payoff as determined by the alternative, and the game ends. Otherwise, the game moves to the next period with the dynamic process selecting a new alternative, and so on, and so forth. This setup is sufficiently general to encompass classical problems in the literatures on collective choice, quitting games, and coalition formation as special cases.

Many problems in collective choice involve the sequential evaluation of alternatives by a committee. When a given alternative is rejected, the committee members will consider a new alternative in a subsequent time period. Moreover, in many cases, none of the committee members has full control over the contents of the new proposal. Compte and Jehiel (2010a) study such problems and mention recruitment decisions in which candidates are examined one by one, business decisions where financing capacity for projects is scarce, and family decisions concerning housing as particular examples.

The decision the committee has to make is whether to accept the current proposal and stop searching, or reject the current proposal and wait for a better alternative to arrive. Penn (2009) argues that, even though in real-life political processes the proposals might be chosen strategically, it is a useful simplifying assumption to consider them as exogenously generated as the outcome of some random process. Indeed the rules of the agenda-setting might be extremely complex, and modeling such an agenda-setting process faithfully might be unnecessary and even undesirable. Roberts (2007), Penn (2009), and Compte and Jehiel (2010a) all assume that the new alternative is drawn from a fixed probability distribution. Nevertheless, in many cases it is more realistic to assume that the probability by which a particular new alternative is selected may depend on the characteristics of the current alternative. We allow for this much richer class of selection dynamics in this paper.

In many collective choice problems, it is natural to assume that the decision-making body and its approval rules are fixed. In other stopping problems, there is not a single decision-making body that can stop the process. In quitting games as introduced in Solan

and Vieille (2001), see also Solan (2005) and Mashiah–Yaakovi (2009), there is at each stage a single player who has the choice between continuing and quitting. The game ends as soon as at least one player chooses to quit. Such games are a variation on wars of attrition models, first analyzed by Maynard–Smith (1974), with economic applications like patent races, oligopoly exit games, and all-pay auctions.

In quitting games, there is a single decision maker who decides to stop or to continue. But in economic applications like oligopoly exit games, such decisions are usually taken by management teams that consist of several decision makers and where decision making takes place by majority or unanimity rule. This motivates the study of stopping problems where at each stage there is a collection of coalitions that can decide to stop the process. The stopping problems in collective choice literature mentioned before correspond to the special case where the collection of decisive coalitions does not change over time.

We allow the collection of decisive coalitions to depend on the alternative that is up for voting. The class of games that we study is therefore also sufficiently general to admit an interpretation of coalition formation. Under this interpretation, in each time period nature selects a coalition and an allocation of payoffs, which is implemented if the coalition members all approve. Contrary to standard non-cooperative models of coalition formation as in Bloch (1996), Okada (1996), Ray and Vohra (1999), or Compte and Jehiel (2010b), nature does not first select a player, who next proposes a coalition, but directly selects the coalition itself.

We do not make any convexity assumptions on the set of feasible alternatives and we do not make any concavity assumptions on the utility functions. Moreover, to have a model that is sufficiently rich to incorporate quitting games and coalition formation games, we do not only study approval rules following from decision making by qualified majorities, but allow for a general collection of decisive coalitions that is associated with each alternative. We first analyze the basic version of the model where the order of voting is history-independent and all actions taken previously are observed by every player. We then extend this basic framework by allowing for incomplete information.

The core consists of those alternatives for which no alternative and associated decisive coalition exists that gives each decisive coalition member a strictly higher utility. We find that each core element, if it exists, naturally induces a subgame perfect equilibrium in pure stationary strategies. When the core is empty, however, subgame perfect equilibrium need not even exist in mixed stationary strategies, where the intuition for non-existence is closely related to the logic of the Condorcet paradox. Any general existence result therefore requires the strategies to exhibit some degree of history-dependence. In the presence of breakdown, stationary equilibria do exist, but might require the use of mixing. Once again,

we will allow strategies to feature some degree of history–dependence to obtain equilibrium existence in pure strategies. In the case with breakdown, an interesting alternative that guarantees the existence of pure strategy stationary equilibria has been proposed in Duggan and Kalandrakis (2010) and consists of introducing small utility shocks to the preferences of players.

For the same reasons Maskin and Tirole (2001) provide for the study of stationary strategies, i.e. reducing the multiplicity of equilibria in dynamic games, reducing the number of parameters to be estimated in econometric models, and amenability to simulation techniques, we are interested in the question whether there are relatively “simple” subgame perfect equilibria. Ideally we would like to know that some relatively small set of well–behaved strategies always contains at least one subgame perfect equilibrium. Notwithstanding the fact that the set of behavior strategies in the games we consider is vastly infinite, we will provide a procedure that determines a subgame perfect equilibrium in a finite number of steps.

We start our analysis by restricting attention to pure action–independent strategies. Action–independence says that the vote of a player cannot be conditioned on the votes previously cast, whether in the same round of voting or in the past. Thus in order to play an action–independent strategy, a player need not “remember” how each individual player has voted so far, but only what alternatives have been voted on, and in fact turned down, so far. The condition of action–independence guarantees a degree of robustness of our result with respect to the precise specification of the voting stage of the game. In the basic model described above, the voting order is exogenously given and is fixed throughout the game, and players observe all previously taken actions. Under action–independence, however, each of these assumptions can be relaxed. Our existence result extends without any difficulty to a more general model where the voting order is history–dependent and/or probabilistic, and information on previously cast votes might be incomplete.

Under action–independence, the play of the game can be conveniently summarized by a so–called collective strategy. A collective strategy describes whether, after each history of alternatives generated by nature, the current proposal must be accepted or rejected. In particular, a collective strategy tells us how to continue the play of the game following a deviation, i.e. a rejection of an alternative that, according to the strategy, should have been accepted. We refer to such a deviation as an “unlawful” rejection. We show how to construct strategies for the players from a given collective strategy and how the concept of a subgame perfect equilibrium can be reformulated in terms of collective strategies, leading to the concept of a collective equilibrium.

A collective strategy is said to be simple if an unlawful rejection of a given alternative

at any point in the game triggers the same continuation play. A simple strategy can thus be described by a relatively small amount of data: namely, the main game-plan according to which the game will be played until the first unlawful rejection occurs, and for every alternative a continuation game-plan that will be played following an unlawful rejection of the given alternative.

A stationary collective strategy consists of some target set of alternatives that are deemed acceptable. As soon as nature generates an alternative from the target set, it is accepted. A stationary collective strategy is clearly simple. If an alternative in the target set is unlawfully rejected, the players simply wait for nature to generate the next alternative in the target set. More general than a stationary collective strategy is a simple strategy that is induced by stationary game-plans, since each unlawfully rejected alternative may lead to a particular set of alternatives that is targeted next. However, we show that a collective equilibrium in simple strategies induced by stationary game-plans may not exist.

A two-step game-plan is only slightly more complicated than a stationary game-plan. Instead of having one target set of alternatives, the players have two sets, a set of alternatives X^1 and a larger set of alternatives X^2 . In a two-step game-plan, players wait m periods for nature to select an alternative from X^1 . If no such alternative is chosen from X^1 in the first m rounds then, the players wait for an alternative from the set X^2 to be chosen. When the threshold m is equal to zero, a two-step game-plan collapses to a stationary game-plan.

We put an a priori upper bound on the threshold. The set of two-step game plans satisfying this bound is finite. Our main result claims that there is a collective equilibrium that is induced by two-step game-plans satisfying the upper bound on the threshold. Moreover, we specify an iterative procedure that terminates in a finite number of steps with the two-step game-plans that induce the collective equilibrium. We also show that the main game-plan according to which the game is played until the first unlawful rejection occurs can be chosen to be stationary. We demonstrate that, unlike stationary equilibrium, our equilibrium concept is generically robust to small perturbations in utilities.

We examine an example with three alternatives and an arbitrary number of players that exhibits the Condorcet paradox: a decisive coalition of players prefers the first alternative to the second, another decisive coalition of player the second alternative to the third, and another decisive coalition prefers the third alternative to the first. We show that our collective equilibrium has the immediate acceptance property in this example. The first alternative that is generated by nature is accepted.

The rest of the paper is organized as follows. Section 2 introduces the basic model and Section 3 establishes the one-shot deviation principle. Section 4 discusses stationary

strategies and relation between the equilibria and the core. We also discuss examples that admit no subgame perfect equilibria in stationary strategies. Section 5 is devoted to pure action-independent strategies, Sections 6 to collective strategies, Section 7 to simple collective strategies, and Section 8 to two-step game-plans. Section 9 provides the main result, existence of a collective equilibrium, and Section 10 presents the case with three alternatives as an example. Section 11 extends the result to an incomplete information setting and Section 12 concludes.

2 The model

We consider a dynamic game $\Gamma = (N, X, \mu^0, \mu, \mathcal{C}, u)$. The set of players is N , a set with cardinality n . In each period $t = 0, 1, \dots$ nature draws an alternative x from a non-empty, finite set X . In period 0 alternatives are selected according to the probability distribution μ^0 on X . In later periods, the selection of an alternative is determined by the Markov process μ , where $\mu(x | \bar{x})$ denotes the probability that the current alternative is x conditional on previous period's alternative being \bar{x} . To keep notation and proofs as simple as possible, we assume that μ is irreducible, i.e. given any current alternative \bar{x} there is positive probability to reach any other alternative x at some point in the future.

After the selection of an alternative, all players vote sequentially, each player casting a “y” or an “n” vote. For the sake of expositional simplicity, we assume that the order of voting \prec is independent of the history of play and that each player observes the entire history of play preceding his own move, assumptions that can easily be avoided as we demonstrate in an extended model in Section 13.

The correspondence $\mathcal{C} : X \rightarrow 2^N$ associates to each alternative x a collection of decisive coalitions $\mathcal{C}(x)$, a collection of subsets of N . Alternative x is accepted if and only if the set of players who vote “y” on x is a member of $\mathcal{C}(x)$. After the acceptance of an alternative, the game ends. Otherwise, the game proceeds to the next time period. The collection $\mathcal{C}(x)$ is assumed to be non-empty and monotonic. If $C \in \mathcal{C}(x)$ and $D \subset N$ is a superset of C , then $D \in \mathcal{C}(x)$. In case $\emptyset \in \mathcal{C}(x)$, alternative x corresponds to a breakdown alternative. The monotonicity assumption on $\mathcal{C}(x)$ implies that, irrespective of the voting behavior of the players, the game ends after a breakdown alternative is selected. We allow for the existence of multiple breakdown alternatives.

Player $i \in N$ has a utility function $u_i : X \rightarrow \mathbb{R}_+$, where $u_i(x)$ is the utility player i derives from the implementation of alternative x . In case of perpetual disagreement, every player's utility is zero. The profile of utility functions $(u_i)_{i \in N}$ is denoted by u . To make

our problem non-trivial, we assume that there is at least one alternative $x \in X$ such that $u(x) \neq 0$. Although players are assumed not to discount utilities, the closely related model with a positive probability $1 - \delta$ of breakdown in every period is a special case of our model. It suffices to specify that one of the alternatives x in X is a breakdown alternative which is selected with probability $1 - \delta$ in every period and leads to utility $u(x) = 0$.

Let $A = \{y, n\}^N$ denote the set of players' joint actions in a voting stage of the game. The subset $A^*(x)$ of A defined by

$$A^*(x) = \{a \in A \mid \{i \in N \mid a^i = y\} \in \mathcal{C}(x)\}$$

is the set of joint actions which leads to the acceptance of alternative x .

Let H_i be the set of all histories where player i makes a decision. We define H_i as the set of all sequences $(s_0, a_0, \dots, s_{t-1}, a_{t-1}, s_t, a_t^{\prec i})$ where s_0, \dots, s_t are alternatives in X , a_0, \dots, a_{t-1} are the actions by the players in the voting stages $0, \dots, t-1$, and $a_t^{\prec i}$ is an element of the set $\{y, n\}^{\prec i}$, where $\prec i = \{j \in N \mid j \prec i\}$ is the set of players who vote before player i . Notice that only sequences $(s_0, a_0, \dots, s_{t-1}, a_{t-1}, s_t, a_t^{\prec i})$ such that $\mu^0(s_0) > 0$, $a_k \notin A^*(s_k)$ for every $k \in \{0, \dots, t-1\}$, and $\mu(s_{k+1} \mid s_k) > 0$ for every $k \in \{0, \dots, t-1\}$ can occur with positive probability. A behavioral strategy for player i is a function $\sigma_i : H_i \rightarrow [0, 1]$ where $\sigma_i(h)$ is the probability for player i to play "y" at history $h \in H_i$.

Three important special cases of the model are collective choice games, quitting games, and coalition formation games. Collective choice games are obtained as follows. Suppose that for every alternative x , $\mathcal{C}(x)$ consists of all coalitions with at least q players. This represents a quota voting rule with q being the size of the majority required for the approval of an alternative. Under this specification of decisive coalitions and μ being such that proposals are drawn from a fixed probability distribution in each period, our model is the discrete analogue of the model of Compte and Jehiel (2010a) and the alternative to Banks and Duggan (2000) where proposals are generated by an exogenous dynamic process rather than chosen endogenously.

In a simple example of a quitting game with perfect information as defined in Solan and Vieille (2001), we have $X = \{x_1, \dots, x_n\}$, $\mu(x_i) = 1/n$ for every $i \in N$, and $\mathcal{C}(x_i) = \{C \subset N \mid i \in C\}$. In this case an alternative x_i is accepted if and only if player i votes in favor of acceptance.

Bloch and Diamantoudi (2011) consider coalition formation in hedonic games. The formation of a coalition $C \in 2^N \setminus \{\emptyset\}$ leads to utilities $u_i(C)$ for the members of coalition C and to utility zero for players in $N \setminus C$. The variation on the game of Bloch and Diamantoudi (2011) where the game ends as soon as the first coalition forms and where nature chooses the

coalition rather than the proposer is a special case of our setup. It is obtained by choosing the non-empty coalitions as the alternatives, and the collection of decisive coalitions for alternative C is given by $\mathcal{C}(C) = \{D \subset N \mid D \supset C\}$.

More generally, since our set of decisive coalitions is allowed to depend on the alternative, we can also interpret our setup as a coalition formation game. Let $X(C)$ denote the subset of alternatives in X for which $C \in \mathcal{C}(x)$. Then $u(X(C))$ corresponds to the payoff set of coalition C . Compared to standard models of non-transferable utility games, our approach also specifies the payoffs to non-coalition members, so externalities are allowed for. Our monotonicity assumption on \mathcal{C} leads to a monotonicity assumption on the sets of payoffs: if $\bar{u} \in u(X(C))$ and $C \subset D$, then $\bar{u} \in u(X(D))$.

The game Γ belongs to the class of stochastic games with perfect information and recursive payoffs. These are stochastic games where each state is controlled by one player, the payoffs in the transient states are all zero, and the payoffs in the absorbing states are non-negative. The main result in Flesch, Kuipers, Schoenmakers, and Vrieze (2010) implies that Γ admits a subgame perfect ϵ -equilibrium for every $\epsilon > 0$. As is explained in detail in the following sections, we exploit the special features of the game Γ to obtain significantly stronger results.

3 The one-shot deviation principle

The one-shot deviation principle claims that a joint strategy $\sigma = (\sigma_i)_{i \in N}$ is a subgame perfect equilibrium if and only if no player i has a strategy σ'_i such that σ'_i agrees with σ_i after all histories in H_i except some history h , and conditional on history h , σ'_i yields player i a higher payoff against σ_{-i} than σ_i . A strategy σ'_i as above is said to be a *profitable one-shot deviation* from σ . It is well-known that the one-shot deviation principle holds in any dynamic game where the payoff function is continuous at infinity, see for instance Fudenberg and Tirole (1991), Theorem 4.2. Unfortunately, the payoff function in our game Γ is not continuous at infinity. Nevertheless, as demonstrated below, our game does satisfy the one-shot deviation principle.

Let a joint strategy σ and a history $h \in H_i$ be given. We let $\pi(t, x \mid \sigma, h)$ denote the probability that the play of the game will terminate in period t with the acceptance of alternative x , conditional on the fact that history h has taken place. The expected payoff of player i conditional on history h is given by

$$v_i(\sigma \mid h) = \sum_{x \in X} u_i(x) \sum_{t=0}^{\infty} \pi(t, x \mid \sigma, h).$$

A joint strategy σ is a *subgame perfect equilibrium* of the game Γ if for every player i , each history $h \in H_i$, and each strategy σ'_i it holds that

$$v_i(\sigma'_i, \sigma_{-i} \mid h) \leq v_i(\sigma_i, \sigma_{-i} \mid h).$$

THEOREM 3.1: (The one-shot deviation principle) *A joint strategy σ is a subgame perfect equilibrium of Γ if and only if for every player i , each $h' \in H_i$, and each strategy σ'_i such that $\sigma'_i(h) = \sigma_i(h)$ for every $h \in H_i \setminus \{h'\}$, it holds that*

$$v_i(\sigma'_i, \sigma_{-i} \mid h') \leq v_i(\sigma_i, \sigma_{-i} \mid h').$$

The “only if” part of the theorem is trivial. The lemma below is a crucial step towards the proof of the “if” part. The result claims that if a player i can improve upon his payoff obtained under σ , he can do so by deviating from σ at finitely many histories in H_i only.

Let σ_i and σ'_i be strategies for player i . For each $k \in \mathbb{N}$, we define the strategy σ_i^k for player i by the following rule: The strategy σ_i^k agrees with σ'_i for histories $h \in H_i$ in periods $0, \dots, k$, and agrees with σ_i for histories $h \in H_i$ in periods $k+1, k+2, \dots$. Formally, for a history $h = (s_0, a_0, \dots, s_{t-1}, a_{t-1}, s_t, a_t^{\leftarrow i})$ in H_i , set $\sigma_i^k(h) = \sigma'_i(h)$ if $t \leq k$ and set $\sigma_i^k(h) = \sigma_i(h)$ otherwise.

LEMMA 3.2: *Let σ_i and σ'_i be strategies for player i and let σ_{-i} be a tuple of strategies for the other players. Suppose that for some $h \in H_i$ we have that*

$$v_i(\sigma'_i, \sigma_{-i} \mid h) > v_i(\sigma_i, \sigma_{-i} \mid h).$$

Then there exists a $k \in \mathbb{N}$ such that

$$v_i(\sigma_i^k, \sigma_{-i} \mid h) > v_i(\sigma_i, \sigma_{-i} \mid h).$$

PROOF: We write $\sigma = (\sigma_i, \sigma_{-i})$, $\sigma' = (\sigma'_i, \sigma_{-i})$, and $\sigma^k = (\sigma_i^k, \sigma_{-i})$. Choose $\epsilon > 0$ such that $v_i(\sigma' \mid h) > v_i(\sigma \mid h) + \epsilon$. Choose $c > 0$ such that for every $i \in N$, $x \in X$, it holds that $u_i(x) \leq c$. Let t_h denote the time period of history h . The sum

$$\sum_{t=t_h}^{\infty} \sum_{x \in X} \pi(t, x \mid \sigma', h)$$

is the probability that the game eventually ends with the acceptance of some alternative, and hence is bounded from above by 1. Therefore, there is a time period $k \geq t_h$ such that

$$\sum_{t=k+1}^{\infty} \sum_{x \in X} \pi(t, x \mid \sigma', h) \leq \frac{\epsilon}{c}$$

and it holds that

$$v_i(\sigma' | h) - \sum_{t=t_h}^k \sum_{x \in X} u_i(x) \pi(t, x | \sigma', h) = \sum_{t=k+1}^{\infty} \sum_{x \in X} u_i(x) \pi(t, x | \sigma', h) \leq \epsilon.$$

Since the joint strategies σ^k and σ' agree on all histories up to and including period k , we have $\pi(t, x | \sigma^k, h) = \pi(t, x | \sigma', h)$ whenever $t_h \leq t \leq k$. Hence

$$v_i(\sigma^k | h) \geq \sum_{t=t_h}^k \sum_{x \in X} u_i(x) \pi(t, x | \sigma^k, h) = \sum_{t=t_h}^k \sum_{x \in X} u_i(x) \pi(t, x | \sigma', h) \geq v_i(\sigma' | h) - \epsilon > v_i(\sigma | h),$$

where the first inequality follows since $u(x) \geq 0$ for every $x \in X$. \square

The one-shot deviation principle can now be derived using a standard technique.

PROOF OF THEOREM 3.1: Suppose there is a player i , a history $h \in H_i$, and a strategy σ'_i such that

$$v_i(\sigma'_i, \sigma_{-i} | h) > v_i(\sigma_i, \sigma_{-i} | h).$$

By Lemma 3.2 there exists k such that $v_i(\sigma_i^k, \sigma_{-i} | h) > v_i(\sigma_i, \sigma_{-i} | h)$. Either there is a history h' in period k where player i has a profitable one-shot deviation from σ_i , or

$$v_i(\sigma_i^{k-1}, \sigma_{-i} | h) > v_i(\sigma_i, \sigma_{-i} | h).$$

This process terminates in finitely many steps with a history where player i has a profitable one-shot deviation. \square

4 Stationary strategies

It has been shown in Fink (1964), Takahashi (1964), and Sobel (1971) that a stochastic game with discounting admits a subgame perfect equilibrium in stationary strategies. When at least one of the alternatives in our model is a breakdown alternative, the techniques of the stochastic game literature with discounting can be used to show the existence of a subgame perfect equilibrium in stationary strategies. However, it is well-known that even in the presence of discounting, subgame perfect equilibria in pure stationary strategies may not exist.

In the absence of discounting, even when allowing for mixed strategies, non-existence of a Nash equilibrium has been noted by Blackwell and Ferguson (1968). This result obviously implies the non-existence of a subgame perfect equilibrium in stationary strategies. This result has spurred an extensive literature on the existence of weaker notions of Nash equilibrium in special classes of stochastic games with the average reward criterion. An example is the class of recursive games with positive payoffs as introduced in Flesch et al. (2010), a class for which they show the existence of a subgame perfect ϵ -equilibrium. Since our model belongs to the class of recursive games with positive payoffs, this result immediately applies.

We study next whether Γ has subgame perfect equilibria in stationary strategies. We say that a strategy is stationary if the probability to vote in favor of a given alternative x depends only on x and is otherwise independent of the history of play.¹

DEFINITION 4.1: A strategy σ_i for player i is *stationary* if for all histories $h = (s_0, a_0, \dots, s_{t-1}, a_{t-1}, s_t, a_t^{\prec i})$ and $h' = (s'_0, a'_0, \dots, s'_{t'-1}, a'_{t'-1}, s'_{t'}, a'_{t'}^{\prec i})$ in H_i such that $s_t = s'_{t'}$ it holds that $\sigma_i(h) = \sigma_i(h')$.

Let x and y be points of X . The alternative x is said to *strictly dominate* y if

$$\{i \in N \mid u_i(x) > u_i(y)\} \in \mathcal{C}(x).$$

Notice that under the maintained assumptions it is possible that x strictly dominates y while at the same time y strictly dominates x . The set of alternatives that strictly dominate y is denoted by $\text{SD}(y)$. An alternative x is said to have the *core property* if it is not strictly dominated by any other alternative. The *core* consists of all alternatives with the core property.

For non-transferable utility games, the core is defined in terms of utilities rather than alternatives, and more precisely as those utilities $\bar{u} \in u(X(N))$ for which there is no $C \subset N$ with $\hat{u} \in u(X(C))$ such that $\hat{u}_i > \bar{u}_i$ for all $i \in C$. The monotonicity property of \mathcal{C} implies the consistency of our definition of the core with the one in the theory on non-transferable utility games.

It follows directly from the definition that a breakdown alternative strictly dominates any alternative (including itself). Non-emptiness of the core therefore implies the absence of breakdown alternatives.

¹Our notion of stationary is somewhat more stringent than the usual one in the literature. Following the approach in Maskin and Tirole (2001) would lead to a notion of stationarity where the voting decision of a player is allowed to depend on the votes cast previously in the current round of voting. Our negative results regarding the existence of stationary equilibria carry over to this weaker notion of stationarity.

Let Γ have a non-empty core, and let $\bar{x} \in X$ be an alternative with the core property. We define the following pure stationary strategy for player $i \in N$:

$$\sigma_i(s_0, a_0, \dots, s_{t-1}, a_{t-1}, s_t, a_t^{\prec i}) = \begin{cases} y, & \text{if } s_t = \bar{x} \text{ or } u_i(s_t) > u_i(\bar{x}), \\ n, & \text{otherwise.} \end{cases} \quad (4.1)$$

Under the joint strategy σ , all players vote in favor of alternative \bar{x} , and since $\mathcal{C}(\bar{x})$ is non-empty and monotonic, \bar{x} is accepted whenever drawn by nature.

THEOREM 4.2: *Let Γ have a non-empty core. The joint strategy σ as defined in (4.1) is a subgame perfect equilibrium in pure stationary strategies.*

PROOF: Let $\bar{x} \in X$ be an alternative with the core property. We define the pure stationary strategy σ_i for player $i \in N$ by (??). It clearly holds that \bar{x} is accepted whenever drawn by nature under the joint strategy σ .

Suppose there is another alternative, say x , which is accepted when drawn by nature. By definition of σ it holds that $\{i \in N \mid u_i(x) > u_i(\bar{x})\} \in \mathcal{C}(x)$, which means that x strictly dominates \bar{x} , a contradiction. Consequently, \bar{x} is the only alternative that will ever be accepted.

Since μ is irreducible and \bar{x} is the only alternative that will ever be accepted, it holds that $v(\sigma) = u(\bar{x})$. Moreover, the expected utility to the players following the rejection of any alternative is given by $u(\bar{x})$.

We verify next that there are no profitable one-shot deviations from σ . Consider a history h at which, according to σ , player i has to vote in favor of an alternative x . By definition of σ it holds that $u_i(x) \geq u_i(\bar{x})$. When the play resulting from σ leads to the acceptance of x , we have that $v_i(\sigma \mid h) = u_i(x)$, whereas a one-shot deviation by player i to a vote against either still results in the acceptance of x , or to a rejection and utility $u_i(\bar{x}) \leq u_i(x)$, so is not profitable. When playing according to σ leads to the rejection of x , a one-shot deviation by player i to a vote against will still lead to the rejection of x by monotonicity of $\mathcal{C}(x)$, and is therefore not profitable. Consider a history h at which, according to σ , player i has to vote against an alternative x . By definition of σ we have that $u_i(x) \leq u_i(\bar{x})$. By monotonicity of $\mathcal{C}(x)$, a one-shot deviation by player i to a vote in favor will either not make a difference or change a rejection of x into an acceptance and lead to utility $u_i(x) \leq u_i(\bar{x})$, so is not profitable.

By Theorem 3.1 we conclude that σ is a subgame perfect equilibrium. \square

The following example of a collective choice game illustrates that even when the core is non-empty, there may be subgame perfect equilibria in pure stationary strategies leading

to the acceptance of an alternative that does not belong to the core. In fact, even in the presence of an alternative that is unanimously preferred to all other alternatives, some of the other alternatives might be accepted.

EXAMPLE 4.3: There are 3 players and 4 alternatives with payoffs given by Table 1 on the left. The collection of decisive coalitions consists of all coalitions with two or more players, so is obtained by an application of simple majority rule. In every time period all alternatives have an equal chance to be selected by nature. Notice that alternative x_4 is strictly preferred by all players to any other alternative. The alternative x_4 has the core property and Theorem 4.2 implies that there is a subgame perfect equilibrium in pure stationary strategies where alternative x_4 is always implemented. However, there is another subgame perfect equilibrium in pure stationary strategies given by Table 1 on the right. In this equilibrium all four alternatives are immediately accepted, resulting in an expected payoff of $15/4$ for every player. Since $15/4 < 4$ and 4 is the minimum utility of a player who casts a yes vote, it is easy to verify that the one-shot deviation principle is satisfied, and the strategy is an equilibrium indeed. \square

	x_1	x_2	x_3	x_4
1	5	4	0	6
2	0	5	4	6
3	4	0	5	6

	x_1	x_2	x_3	x_4
1	y	y	n	y
2	n	y	y	y
3	y	n	y	y

Table 1: Payoffs and strategies in an example with a non-empty core.

The following two examples illustrate that when Γ has an empty core, there might not be a subgame perfect equilibrium in stationary strategies. The first example corresponds to a quitting game, the second example is a collective choice game that exhibits the Condorcet paradox. The second example can also be reformulated as a coalition formation game. This example is closely related to an example on coalition formation where Bloch (1996) shows non-existence of a subgame perfect equilibrium in stationary strategies.

EXAMPLE 4.4: This is an example of a 3-player quitting game with perfect information as studied in Solan (2005). The set of players is $N = \{1, 2, 3\}$ and the set of alternatives is $X = \{x_1, x_2, x_3\}$. The utilities of these alternatives are specified in Table 2. In every time period, every alternative is chosen by nature with probability $1/3$. Player i is decisive for alternative x_i , that is $\mathcal{C}(x_i) = \{C \subset N \mid i \in C\}$. It holds that x_3 strictly dominates x_1, x_2 .

	x_1	x_2	x_3
1	1	0	3
2	7	4	0
3	0	7	4

Table 2: Payoffs in Example 2.

strictly dominates x_2 , and x_2 strictly dominates x_3 , so the core of the game is empty, and Theorem 4.2 cannot be applied.

In this example, a stationary strategy of player i can be represented by a number $\alpha_i \in [0, 1]$, being the probability for player i to vote in favor of alternative x_i . In a stationary strategy, the vote of player i on other alternatives is inconsequential, and can therefore be ignored. For $i \in N$, it holds that $v_i(0, 0, 0) = 0$ and

$$v_i(\alpha) = \frac{1}{\alpha_1 + \alpha_2 + \alpha_3}(\alpha_1 u_i(x_1) + \alpha_2 u_i(x_2) + \alpha_3 u_i(x_3)), \quad \alpha \neq (0, 0, 0).$$

By stationarity, $v(\alpha) = (v_1(\alpha), v_2(\alpha), v_3(\alpha))$ is also the expected utility conditional on the rejection of any alternative. By Theorem 3.1 it holds that the joint stationary strategy α is subgame perfect if and only if for every player i

$$\begin{aligned} \alpha_i > 0 & \text{ implies } u_i(x_i) \geq v_i(\alpha), \\ \alpha_i < 1 & \text{ implies } u_i(x_i) \leq v_i(\alpha). \end{aligned}$$

We claim that the game has no subgame perfect equilibrium in stationary strategies. For suppose α is such a strategy. We split the argument into four cases depending on the number m of non-zero components of α .

Case $m = 0$. We have that $\alpha = (0, 0, 0)$ and $v(\alpha) = (0, 0, 0)$. Since $u_1(x_1) = 1 > 0 = v_1(\alpha)$, we must have $\alpha_1 = 1$, a contradiction.

Case $m = 1$. Suppose first that $\alpha_1 > 0$. Then $v(\alpha) = u(x_1)$. But then $u_3(x_3) = 4 > 0 = v_3(\alpha)$ and so we must have $\alpha_3 = 1$, implying that $m \geq 2$. Similarly, if $\alpha_2 > 0$ a contradiction arises since $v(\alpha) = u(x_2)$ and $u_1(x_1) = 1 > 0 = v_1(\alpha)$, and if $\alpha_3 > 0$ a contradiction arises since $v(\alpha) = u(x_3)$ and $u_2(x_2) = 4 > 0 = v_2(\alpha)$.

Case $m = 2$. Suppose first that $\alpha_3 = 0$. Then $v(\alpha)$ is a strictly convex combination of $u(x_1)$ and $u(x_2)$ and hence $v_2(\alpha) > u_2(x_2)$. Hence $\alpha_2 = 0$, implying that $m \leq 1$, a contradiction. Similarly, if $\alpha_2 = 0$ a contradiction arises since $v_1(\alpha) > u_1(x_1)$, and if $\alpha_1 = 0$ a contradiction arises since $v_3(\alpha) > u_3(x_3)$.

Case $m = 3$. We have $\alpha_1, \alpha_2, \alpha_3 > 0$ and so

$$u_i(x_i) \geq v_i(\alpha) = \frac{1}{\alpha_1 + \alpha_2 + \alpha_3}(\alpha_1 u_i(x_1) + \alpha_2 u_i(x_2) + \alpha_3 u_i(x_3)), \quad i \in N.$$

Rewriting leads to the inequalities

$$\alpha_3/\alpha_2 \leq 1/2, \quad \alpha_1/\alpha_3 \leq 4/3, \quad \text{and} \quad \alpha_2/\alpha_1 \leq 4/3,$$

and therefore

$$1 = \frac{\alpha_3}{\alpha_2} \frac{\alpha_1}{\alpha_3} \frac{\alpha_2}{\alpha_1} \leq \frac{1}{2} \cdot \frac{4}{3} \cdot \frac{4}{3} = \frac{8}{9},$$

a contradiction. □

EXAMPLE 4.5: All the primitives are the same as in the preceding example, except for the collection of decisive coalitions. Suppose that the votes of two out of three players are sufficient for the acceptance of an alternative, so $\mathcal{C}(x)$ consists of all subsets of N with at least two players. This is an example of a Condorcet paradox, where the majority induced preference relation is intransitive. In a pairwise comparison, alternative x_1 beats x_2 , alternative x_2 beats x_3 , and alternative x_3 beats x_1 . Herings and Houba (2010) study this example under the alternative model where the proposer is selected by nature, rather than the alternative itself.

We claim that the game admits no subgame perfect equilibrium in stationary strategies. A stationary strategy of player i in this game specifies, for every alternative x , the probability for player i to vote in favor of x .

Suppose the joint stationary strategy σ is a subgame perfect equilibrium. It is rather straightforward, though tedious, to show that the cases where σ leads to the rejection of all alternatives, or to the acceptance of at most one alternative with positive probability, are not compatible with equilibrium. Equilibrium utilities are therefore strictly in between the utility of the worst and the best alternative for every player. Next it is rather straightforward, though tedious as well, to show that each player votes in favor of his best alternative and against his worst alternative with probability one.

Hence, in essence, every player i is decisive for alternative x_i since there is one more player who votes in favor of x_i and one more who votes against x_i . Thus the joint stationary strategy σ induces a subgame perfect equilibrium of the game of Example 4.4, contradicting the conclusion that no stationary subgame perfect equilibria exist in that game. Consequently, no joint stationary strategy σ can be a subgame perfect equilibrium in the game of Example 4.5. □

	x_0	x_1	x_2	x_3
1	0	1	0	5
2	0	5	1	0
3	0	0	5	1

Table 3: Payoffs in Example 4.6.

Example 4.5 studies the case of majority voting in a three-player setup with an empty core, or equivalently, the absence of a Condorcet winner. It has been shown in the literature that the occurrence of the Condorcet paradox is not an artifact. Work by Plott (1967), Rubinstein (1979), Schofield (1983), Cox (1984), and Le Breton (1987) shows that the core is generically empty in majority voting situations with three or more players in a setup where alternatives are a compact, convex subset of some Euclidean space. Such voting situations can be approximated arbitrarily closely in our setup with a finite set of alternatives.

The next example of a quitting game shows that in the presence of a breakdown alternative, subgame perfect equilibria in pure stationary strategies may fail to exist, and that subgame perfect equilibria in mixed stationary strategies, which do exist in this case, are Pareto inefficient.

EXAMPLE 4.6: In this example the set of players is $N = \{1, 2, 3\}$ and the set of alternatives is $X = \{x_0, x_1, x_2, x_3\}$. The utilities of these alternatives are specified in Table 3. Every period a breakdown alternative x_0 is selected with a probability $1 - \delta$ strictly in between 0 and 1. For the other alternatives it holds that player i is decisive for alternative x_i , which is selected with probability $\delta/3$ in every period.

In this example, a stationary strategy of player i can be represented by a number $\alpha_i \in [0, 1]$, being the probability for player i to vote in favor of alternative x_i . In a stationary strategy, the vote of player i on other alternatives is inconsequential, and can therefore be ignored. For $i \in N$, it holds that

$$v_i(\alpha) = \frac{1}{3 - 3\delta + \delta\alpha_1 + \delta\alpha_2 + \delta\alpha_3} (\delta\alpha_1 u_i(x_1) + \delta\alpha_2 u_i(x_2) + \delta\alpha_3 u_i(x_3)).$$

By stationarity, $v(\alpha) = (v_1(\alpha), v_2(\alpha), v_3(\alpha))$ is also the expected utility conditional on the rejection of any alternative. By Theorem 3.1 it holds that the joint stationary strategy α is subgame perfect if and only if for every player i

$$\begin{aligned} \alpha_i > 0 & \text{ implies } u_i(x_i) \geq v_i(\alpha), \\ \alpha_i < 1 & \text{ implies } u_i(x_i) \leq v_i(\alpha). \end{aligned}$$

It is a routine exercise to verify that there is a unique subgame perfect equilibrium in stationary strategies. If $\delta \leq 1/2$, then it is given by the pure strategy $\alpha_1 = \alpha_2 = \alpha_3 = 1$ with expected payoffs $v(\alpha) = (2\delta, 2\delta, 2\delta)$. If $\delta > 1/2$, then the subgame perfect equilibrium in stationary strategies is given by the mixed strategy $\alpha_1 = \alpha_2 = \alpha_3 = (1 - \delta)/\delta$ with expected payoff $v(\alpha) = (1, 1, 1)$ irrespective of δ . A strategy profile that would lead to immediate acceptance has payoffs $(2\delta, 2\delta, 2\delta)$, so the delay in the unique stationary subgame perfect equilibrium causes substantial inefficiencies. The expected delay before reaching an agreement is equal to $(2\delta - 2)/(2 - 2\delta)$ periods, which tends to infinity as δ tends to 1. In the limit, every alternative is rejected with probability 1. \square

5 Action–independence

The objective of this paper is to prove the existence of a subgame perfect equilibrium in the game Γ . Moreover, we would like to construct a subgame perfect equilibrium in pure strategies that exhibit a relatively small amount of history–dependence and that are computable in finitely many steps. To guarantee the existence of a subgame perfect equilibrium, strategies need to exhibit some amount of history–dependence as is evidenced by Examples 4.4 and 4.5. Example 4.6 shows that even in the presence of breakdown alternatives, some history–dependence is needed to guarantee the existence of a subgame perfect equilibrium in pure strategies.

We will consider strategies that are pure and action–independent. A strategy is said to be *action–independent* if it prescribes the same action after the same sequence of moves by nature. In other words, under action–independence a player is not allowed to condition his vote on the actions of the players, but only on the moves by nature.

DEFINITION 5.1: A strategy σ_i for player i is *action–independent* if for all histories $h = (s_0, a_0, \dots, s_{t-1}, a_{t-1}, s_t, a_t^{\prec i})$ and $\bar{h} = (\bar{s}_0, \bar{a}_0, \dots, \bar{s}_{t-1}, \bar{a}_{t-1}, \bar{s}_t, \bar{a}_t^{\prec i})$ in H_i such that $s_0 = \bar{s}_0, \dots, s_t = \bar{s}_t$, it holds that $\sigma_i(h) = \sigma_i(\bar{h})$.

To illustrate, consider the case where $N = \{1, 2, 3\}$, $1 \prec 2 \prec 3$, and the approval of two players is sufficient to accept an alternative. Nature has selected some alternative s_0 in period 0 and some alternative s_1 in period 1. Action–independence requires that player 3’s vote on alternative s_0 in period 0 be the same after each of the histories (s_0, n, n) , (s_0, y, n) , (s_0, n, y) , and (s_0, y, y) , and that player 1’s vote on alternative s_1 in period 1 is the same after each of the histories (s_0, a_0, s_1) , irrespective of the choice of $a_0 \in A$.

The requirement of action–independence guarantees the robustness of our solution with respect to the exact setup of the voting stage of the game. We presently assume that the

players vote on each proposal sequentially, that the order of voting is fixed, and that the players observe all moves previously made. Action-independence implies that the order of voting is inessential. Accordingly, our existence result carries over without any difficulty to a game where the order of voting is random or history-dependent. Moreover, under action-independence it is inessential whether the players are indeed able to observe previously cast ballots or not. Consequently, our existence result extends to a game where the players have various degrees of incomplete information about the voting behavior of other players. In particular, we encompass the situation where the players cast their votes simultaneously and the votes are not, partially, or completely disclosed at the end of the voting stage. Section 13 presents all the details about such an extension to incomplete information settings.

6 Collective strategies

Under action-independence, the play of the game can be conveniently summarized by means of so-called collective strategies. A collective strategy is a complete contingent plan of actions which specifies whether a given alternative has to be accepted or rejected, following a given sequence of alternatives selected by nature.

Let S be the set of all finite sequences of elements of X . A *collective strategy* is a function f from S to $\{0, 1\}$, where 0 corresponds to a rejection and 1 to an acceptance. We define $f(s)$ for every element $s \in S$, so f also specifies how the game is played for counterfactual situations where the game proceeds after the acceptance of an alternative. Let F be the set of functions from S to $\{0, 1\}$.

Consider a joint strategy σ where σ_i is pure and action-independent for each player $i \in N$. The *collective strategy* $f_\sigma : S \rightarrow \{0, 1\}$ induced by σ is defined by

$$f_\sigma(s_0, \dots, s_t) = \begin{cases} 1, & \text{if } \{i \in N \mid \sigma_i(s_0, \dots, s_t) = y\} \in \mathcal{C}(s_t), \\ 0, & \text{otherwise.} \end{cases} \quad (6.1)$$

Here we have used action-independence to treat σ_i as a function with domain S rather than H_i . Thus $f_\sigma(s) = 1$ if and only if s is accepted according to the joint strategy σ . Notice that $f_\sigma(s) = 1$ if s_t is a breakdown alternative, irrespective of the actual voting behavior.

Given an $x \in X$ we let Γ_x denote the game where the initial distribution μ^0 over X is given by $\mu^0(y) = \mu(y|x)$ for every $y \in X$. We let $V_x(f)$ denote the vector of expected payoffs in the beginning of the game Γ_x , if the play proceeds according to a collective

strategy f . We shall often consider the case where the Markov process μ is stationary, that is $\mu(y|x) = \mu(y|z)$ for all x, y , and z . In this special case $V_x(f)$ is clearly the same for each $x \in X$, and we shall write simply $V(f)$ to denote this common value.

Consider now a sequence $s = (s_0, \dots, s_t)$ of elements of X . It will be important to compute the expected payoffs in the subgame that begins with a move of nature following the rejection of the alternatives in the sequence s_0, \dots, s_t . Notice that this subgame is identical to the game Γ_{s_t} . We let $f[s]$ denote the continuation collective strategy after the rejection of the alternatives in s . It is given by the equation

$$f[s](r) = f(s \oplus r)$$

for each $r \in S$ where $s \oplus r$ is the concatenation of s and r . Therefore the expected payoffs after the rejection of the sequence s_0, \dots, s_t of alternatives are then given by $V_{s_t}(f[s_0, \dots, s_t])$.

EXAMPLE 6.1: Consider the profile σ of action-independent strategies given by (4.1). The corresponding collective strategy f_σ is given by $f_\sigma(s_0, \dots, s_t) = 1$ if and only if $s_t \in \{\bar{x}\} \cup \text{SD}(\bar{x})$. If \bar{x} has the core property then $\text{SD}(\bar{x})$ is empty. In this case $V_x(f_\sigma) = u(\bar{x})$ for every $x \in X$. This equation holds since by our assumption the Markov process μ is irreducible, so starting from any alternative x the process arrives at \bar{x} with a non-zero probability. \square

Define

$$\begin{aligned} \text{SD}(f) &= \{x \in X : \{i \in N : x_i > V_{x,i}(f)\} \in \mathcal{C}(x)\} \\ \text{WD}(f) &= \{x \in X : \{i \in N : x_i \geq V_{x,i}(f)\} \in \mathcal{C}(x)\}. \end{aligned}$$

The alternatives in $\text{SD}(f)$ are said to *strictly dominate* the collective strategy f and those in $\text{WD}(f)$ are said to *weakly dominate* it. Thus x strictly dominates f if the decisive set of players prefers accepting the alternative x over rejecting it and playing the rest of the game in accordance with a collective strategy f . The definition of strict dominance extends that given in the previous section: The alternative x strictly dominates the alternative y if and only if x strictly dominates the collective strategy f defined by setting $f(s_0, \dots, s_t) = 1$ if and only if $s_t = x$. (This equivalence holds because $V_x(f) = u(y)$).

We now present a notion of equilibrium that is in terms of collective strategies only.

DEFINITION 6.2: The collective strategy $f \in F$ is a *collective equilibrium* if for every sequence $s = (s_0, \dots, s_t) \in S$ it holds that

$$s_t \in \text{SD}(f[s]) \text{ implies } f(s) = 1, \tag{6.2}$$

$$s_t \notin \text{WD}(f[s]) \text{ implies } f(s) = 0. \tag{6.3}$$

Notice that a breakdown alternative is never rejected in a collective equilibrium as a breakdown alternative strictly dominates any collective strategy. The next result shows that a collective equilibrium induces a subgame perfect equilibrium in pure action-independent strategies.

THEOREM 6.3: *Let $f \in F$ be a collective equilibrium. Then the pure action-independent joint strategy defined for $i \in N$ and $s = (s_0, \dots, s_t) \in S$ by*

$$\sigma_i(s) = \begin{cases} y, & \text{if } u_i(s_t) > V_{s_t,i}(f[s]), \\ y, & \text{if } u_i(s_t) = V_{s_t,i}(f[s]) \text{ and } f(s) = 1, \\ n, & \text{otherwise,} \end{cases} \quad (6.4)$$

is a subgame perfect equilibrium of the game Γ with $f_\sigma = f$.

PROOF: We verify that $f_\sigma = f$. Consider some $s = (s_0, \dots, s_t) \in S$ and let $C = \{i \in N \mid \sigma_i(s) = y\}$ be the set of players voting in favor after the sequence s .

In case $f(s) = 1$, we have by definition of σ that $C = \{i \in N \mid u_i(s_t) \geq V_{s_t,i}(f[s])\}$. By (6.3) it holds that $s_t \in \text{WD}(f[s])$, so C belongs to $\mathcal{C}(s_t)$. We conclude that $f_\sigma(s) = 1$.

In case $f(s) = 0$, we have by definition of σ that $C = \{i \in N \mid u_i(s_t) > V_{s_t,i}(f[s])\}$. By (6.2) it holds that $s_t \notin \text{SD}(f[s])$ which means that C does not belong to $\mathcal{C}(s_t)$. We conclude that $f_\sigma(s) = 0$.

It is straightforward to verify that the joint strategy σ satisfies the one-shot deviation property. We invoke Theorem 3.1 to conclude that the joint strategy σ is a subgame perfect equilibrium. \square

The next result shows that any subgame perfect equilibrium in pure action-independent strategies induces a collective equilibrium.

THEOREM 6.4: *Let σ be a subgame perfect equilibrium of Γ in pure action-independent strategies. Then f_σ is a collective equilibrium.*

PROOF: Let $f = f_\sigma$. We show first that Condition (6.2) holds. Let $s = (s_0, \dots, s_t) \in S$ be such that $s_t \in \text{SD}(f[s])$. If s_t is a breakdown alternative, then $f(s) = 1$, so Condition (6.2) holds. Consider the case where s_t is not a breakdown alternative. Suppose that $f(s) = 0$. Since $s_t \in \text{SD}(f[s])$, the set of players $C = \{i \in N \mid u_i(s_t) > V_{s_t,i}(f[s])\}$ belongs to $\mathcal{C}(s_t)$. Now label the players in C as $\{i_1, \dots, i_m\}$, where $i_1 \prec \dots \prec i_m$. For $i \in N$, let $a_i = \sigma_i(s)$ and $a = (a_1, \dots, a_n)$. Consider the action profiles a^0, a^1, \dots, a^m where $a^0 = a$ and, for

every k , $1 \leq k \leq m$, we define

$$a_i^k = \begin{cases} y, & \text{if } i \in \{i_1, \dots, i_k\}, \\ a_i, & \text{otherwise.} \end{cases}$$

Since by our supposition the strategy σ results in the rejection of alternative s_t , we have $a^0 \notin A^*(s_t)$. On the other hand, since $a_i^m = y$ for every $i \in C$ and since $C \in \mathcal{C}(s_t)$, we have that $a^m \in A^*(s_t)$. Let k^* be the least integer for which $a^{k^*} \in A^*(s_t)$, and let $i^* = i_{k^*}$. Notice that $\sigma_{i^*}(s) = a_{i^*} = n$ for otherwise it would be the case that $a^{k^*-1} = a^{k^*}$, which contradicts the choice of k^* .

Consider now any history h in H_{i^*} where the sequence of alternatives is s and the vote of players $i \prec i^*$ in period t is equal to $a_i^{k^*}$. Since σ is action-independent, an n -vote by player i^* at h leads to the action profile a^{k^*-1} in period t and hence to the rejection of s_t . A y -vote by i^* at h leads to the action profile a^{k^*} and hence to the acceptance of s_t . Since $u_{i^*}(s_t) > V_{s_t, i^*}(f[s])$, subgame perfection requires that player i^* cast a y -vote at h , whereas $\sigma_{i^*}(s) = n$, a contradiction.

We now prove that Condition (6.3) holds. Let $s = (s_0, \dots, s_t) \in S$ be such that $s_t \notin \text{WD}(f[s])$. It follows that s_t is not a breakdown alternative. Suppose that $f(s) = 1$. Since $s_t \notin \text{WD}(f[s])$, the set $C = \{i \in N \mid u_i(s_t) < V_{s_t, i}(f[s])\}$ has a non-empty intersection with every member of $\mathcal{C}(s_t)$. Now label the players in C as $\{i_1, \dots, i_m\}$, where $i_1 \prec \dots \prec i_m$. For $i \in N$, let $a_i = \sigma_i(s)$ and $a = (a_1, \dots, a_n)$. Consider the action profiles a^0, a^1, \dots, a^m where $a^0 = a$ and, for every k , $1 \leq k \leq m$, we define

$$a_i^k = \begin{cases} n, & \text{if } i \in \{i_1, \dots, i_k\}, \\ a_i, & \text{otherwise.} \end{cases}$$

Since by our supposition the strategy σ results in the acceptance of alternative s_t , we have $a^0 \in A^*(s_t)$. On the other hand, since $a_i^m = n$ for every $i \in C$ and the intersection of C with every member of $\mathcal{C}(s_t)$ is non-empty, we have that $a^m \notin A^*(s_t)$. Let k^* be the least integer for which $a^{k^*} \notin A^*(s_t)$, and let $i^* = i_{k^*}$. Notice that $\sigma_{i^*}(s) = a_{i^*} = y$ for otherwise it would be the case that $a^{k^*-1} = a^{k^*}$, which contradicts the choice of k^* .

Consider now any history h in H_{i^*} where the sequence of alternatives is s and the vote of players $i \prec i^*$ in period t is equal to $a_i^{k^*}$. Since σ is action-independent, a y -vote by player i^* at h leads to the action profile a^{k^*-1} in period t and hence to the acceptance of s_t . An n -vote by i^* at h leads to the action profile a^{k^*} and hence to the rejection of s_t . Since $u_{i^*}(s_t) < V_{s_t, i^*}(f[s])$, subgame perfection requires that player i^* cast an n -vote at h , whereas $\sigma_{i^*}(s) = y$, a contradiction. \square

As a consequence of Theorems 6.3 and 6.4, we can perform our analysis using collective rather than individual strategies and we refer to members of S as collective histories. Individual strategies can be recovered from a given collective strategy using Equation (6.4).

7 Simple collective strategies

Consider a collection \mathcal{F} consisting of collective strategies f_0 and f_x , for each $x \in X$. Define a new collective strategy f as follows: Fix an infinite sequence s_0, s_1, \dots of the alternatives and proceed as follows:

1. Follow the strategy f_0 until the first time a deviation from f_0 occurs.

Two types of deviations are possible: an acceptance whereas a rejection should (according to f_0) take place, and a rejection whereas an acceptance should take place (we refer to this second type of deviations as an unlawful rejection). Since after an acceptance the game ends we only have to specify what happens after an unlawful rejection.

2. If the alternative at time t^0 must be accepted according to f_0 but it is rejected instead, switch to the strategy $f_{s_{t^0}}$ forgetting the history s_0, \dots, s_{t^0} . Follow $f_{s_{t^0}}$ until the first time a deviation occurs.
3. If the alternative at time t^1 must be accepted according to $f_{s_{t^0}}$ but it is rejected instead, switch to the strategy $f_{s_{t^1}}$ forgetting the history s_0, \dots, s_{t^1} . Follow $f_{s_{t^1}}$ until the first time a deviation occurs.

And so on. The strategy f thus constructed is said to be *induced* by collection \mathcal{F} . A collective strategy is said to be *simple* if it is induced by some collection \mathcal{F} of strategies (this terminology is motivated by the resemblance of our definition to that in Abreu (1988), see the discussion at the end of the section).

The formal definition of the collective strategy f is by induction on the length of the sequence. Define $f(x) = f_0(x)$ for each $x \in X$. Now suppose that f has already been defined on each sequence of alternatives of length of at most t . Consider a sequence $s = (s_0, \dots, s_t)$ of length $t + 1$. Let

$$K(s) = \{k \in \{0, \dots, t-1\} : f(s_0, \dots, s_k) = 1\}.$$

and define

$$f(s_0, \dots, s_t) = \begin{cases} f_0(s_0, \dots, s_t) & \text{if } K(s) = \emptyset \\ f_{s_k}(s_{k+1}, \dots, s_t) & \text{if } K(s) \neq \emptyset \text{ and } k = \max K(s). \end{cases}$$

t	0	1	2	3	4	5
s_t	x_1	x_3	x_1	x_2	x_1	x_1
$f(s_0, \dots, s_t)$	1	0	0	1	1	0

Table 4: The collective strategy f .

The following example illustrates our definitions.

EXAMPLE 7.1: Let $X = \{x_1, x_2, x_3\}$. Consider the collective strategies g and h where $g(s) = 1$ for every $s \in S$ and

$$h(s_0, \dots, s_t) = \begin{cases} 1 & \text{if } [t = 0 \text{ and } s_0 = x_2] \text{ or } [t \geq 1 \text{ and } s_t \in \{x_2, x_3\}] \\ 0 & \text{otherwise,} \end{cases}$$

Consider now the tuple

$$(f_0, f_{x_1}, f_{x_2}, f_{x_3}) = (g, h, g, g)$$

and let f be the induced simple collective strategy. Some of the values of f are given by Table 4. To derive these values, we reason as follows. Each alternative is accepted according to the strategy $f_0 = g$, hence $f(x_1) = 1$. If x_1 is rejected instead, we switch to the strategy $f_{x_1} = h$. According to h only x_2 is initially accepted, so $f(x_1, x_3) = 0$. After the first round, only x_2 and x_3 are accepted according to h , therefore $f(x_1, x_3, x_1) = 0$ and $f(x_1, x_3, x_1, x_2) = 1$. If x_2 is rejected, we switch to the strategy $f_{x_2} = g$ according to which each alternative is accepted. Hence $f(x_1, x_3, x_1, x_2, x_1) = 1$. If x_1 is rejected we switch to $f_{x_1} = h$. Since according to h only x_2 is initially accepted we have $f(x_1, x_3, x_1, x_2, x_1, x_1) = 0$. \square

A *game-plan* is a part of the strategy f that is to be followed until the first deviation from f occurs. More precisely, a game-plan of the strategy f , denoted by $P(f)$ is a set of sequences $(s_0, \dots, s_t) \in S$ such that there is no $k \in \{0, \dots, t-1\}$ with $f(s_0, \dots, s_k) = 1$. Notice that a game-plan of a strategy is always a tree and that it always contains all sequences of length 1. In the preceding example, the game-plan $P(g)$ consists only of sequences of length 1, namely x_1 , x_2 , and x_3 . The game-plan $P(h)$ is depicted in Figure 1.

Notice that the payoffs on a strategy only depend on the game-plan. Formally if f and f' are collective strategies with $P(f) = P(f')$ then $V_x(f) = V_x(f')$ for each $x \in X$.

The following lemma is obvious from the definition of an induced strategy. It will be helpful for computing payoffs on a simple strategy. Here $K(s)$ is defined as above.




Figure 1: The game plan $P(h)$ for h as in Example 7.1

LEMMA 7.2: Suppose the collective strategy f is induced by the tuple $(f_j : j \in \{0\} \cup X)$. Consider $s = (s_0, \dots, s_t) \in S$. Then

- [1] $P(f) = P(f_0)$.
- [2] If $f(s) = 1$ then $P(f[s]) = P(f_{s_t})$.
- [3] If $f(s) = 0$ and $K(s) = \emptyset$ then $P(f[s]) = P(f_0[s])$.
- [4] If $f(s) = 0$ and $K(s) \neq \emptyset$ and $k = \max K(s)$ then $P(f[s]) = P(f_{s_k}[s_{k+1}, \dots, s_t])$.

In particular, the strategy f is completely characterized by the collection of game-plans $(P(f_j) : j \in \{0\} \cup X)$. More formally:

COROLLARY 7.3: Let $\mathcal{F} = (f_j : j \in \{0\} \cup X)$ and $\mathcal{G} = (g_j : j \in \{0\} \cup X)$. Suppose that $P(f_j) = P(g_j)$ for each $j \in \{0\} \cup X$. Then the strategy induced by \mathcal{F} equals that induced by \mathcal{G} .

In view of this observation we shall say that the strategy f is induced by a collection of game-plans, rather than the collection of strategies. The game-plan $P(f_0) = P(f)$ is said to be the main game-plan of the strategy f . It is the plan of actions to be followed until the first deviation from f occurs. An unlawful rejection of an alternative x triggers the punishment game-plan $P(f_x)$.

Notwithstanding the completely different context under consideration, our notion of simplicity bears some resemblance to the one of Abreu (1988), Definition 1. Abreu (1988)

defines a simple strategy in a repeated game by a collection of paths, the main path and n punishments paths, one for every player. A path is a sequence of actions in a stage game and corresponds to our concept of a game-plan. A simple strategy is then defined as follows: the main path is to be followed until the first unilateral deviation, and a unilateral deviation by player i at any point in the game triggers the corresponding punishment path of play. By comparison, in our definition of simplicity, the punishment game-plan does not depend on the identity of the deviating player, but rather on what alternative has been unlawfully rejected.

8 Two-step game-plans

In this section we consider collective strategies that are induced by two-step game-plans. Before we proceed with the definition of a two-step game-plan, we explore the simpler, and more restrictive, condition that the collective strategy be induced by stationary game-plans.

DEFINITION 8.1: For every $Y \subset X$, the collective strategy f_Y is defined by setting, for every $s = (s_0, \dots, s_t) \in S$,

$$f_Y(s) = \begin{cases} 1 & \text{if } s_t \in Y \\ 0 & \text{otherwise.} \end{cases}$$

A collective strategy $f \in F$ is *stationary* if $f = f_Y$ for some $Y \subset X$.

The definition of a stationary collective strategy f in Definition 8.1 is consistent with the definition of a stationary strategy σ in Definition 4.1 in the sense that the pure action-independent joint strategy σ derived from a stationary collective strategy f by means of (6.4) is stationary, and a stationary strategy σ induces a stationary collective strategy f_σ . A stationary collective strategy is easily seen to be simple.

DEFINITION 8.2: A collective strategy f is *induced by stationary game-plans* if it is induced by a tuple of strategies $(f_0, f_x, : x \in X)$ that are all stationary.

Notice that a collective strategy induced by stationary game-plans need not be stationary. A stationary collective strategy is clearly induced by stationary game-plans. We show that the game of Example 4.6 has a collective equilibrium induced by stationary game-plans.

EXAMPLE 8.3: We study Example 4.6. When $\delta \leq 1/2$, we have already derived that there is a stationary pure strategy equilibrium where all alternatives are accepted. This equilibrium is stationary.

We consider next the case where $\delta > 1/2$. Consider the collective strategy f induced by the tuple of game-plans

$$(f_0, f_{x_0}, f_{x_1}, f_{x_2}, f_{x_3}) = (f_X, f_X, f_{\{x_0, x_1, x_2\}}, f_{\{x_0, x_2, x_3\}}, f_{\{x_0, x_1, x_3\}}).$$

Under the collective strategy f , every alternative is accepted in period 0 and the break-down alternative is accepted in every period. Suppose, for some $i = 1, 2, 3$, alternative x_i is unlawfully rejected. The punishment strategy f_{x_i} is such that the most attractive alternative for player i is no longer accepted.

According to Definition 6.2, $f \in F$ is a collective equilibrium in the game of Example 4.6 if and only if for every sequence $s = (s_0, \dots, s_t) \in S$, where $s_t = x_i$, it holds that $f(s) = 1$ if $i = 0$, and for $i = 1, 2, 3$ it holds that

$$u_i(x_i) > V_i(f[s]) \quad \text{implies} \quad f(s) = 1 \quad (8.1)$$

$$u_i(x_i) < V_i(f[s]) \quad \text{implies} \quad f(s) = 0. \quad (8.2)$$

To verify Condition (8.1) take an $s = (s_0, \dots, s_t)$ with $f(s) = 1$. If $s_t = x_i$ then $V_i(f[s]) = V_i(f_{x_i}) = \delta/(3 - \delta) < 1 = u_i(x_i)$, where the first equation uses part [2] of Lemma 7.2. To verify Condition (8.2) suppose $f(s) = 0$. Suppose for concreteness that $s_t = x_1$. This means that the punishment strategy f_{x_2} is being followed. Since this strategy is stationary we find that $V_1(f[s]) = V_1(f_{x_2}) = 5\delta/(3 - \delta) > 1 = u_1(x_1)$, where the first equality uses part [4] of Lemma 7.2 and the inequality uses the fact that $\delta > 1/2$. The cases where $s_t = x_2$ and $s_t = x_3$ are similar.

Example 4.6 demonstrates that the unique stationary equilibrium is in mixed strategies, has payoffs $(1, 1, 1)$ irrespective of the value of δ , and leads to inefficiency because of delay, which tends to infinity. The collective equilibrium constructed here has immediate acceptance and Pareto efficient payoffs $V(f) = (2\delta, 2\delta, 2\delta)$. \square

We show that the game of Example 4.4 admits no collective equilibrium induced by stationary game-plans. Let the *carrier* of f , denoted by $\text{car}(f)$ be the set of $x \in X$ for which there exists a sequence $s = (s_0, \dots, s_t) \in S$ such that $s_t = x$ and $f(s) = 1$.

EXAMPLE 8.4: According to Definition 6.2, $f \in F$ is a collective equilibrium in the game of Example 4.4 if and only if for every sequence $s = (s_0, \dots, s_t) \in S$, where $s_t = x_i$, it holds

that

$$u_i(x_i) > V_i(f[s]) \quad \text{implies} \quad f(s) = 1 \quad (8.3)$$

$$u_i(x_i) < V_i(f[s]) \quad \text{implies} \quad f(s) = 0. \quad (8.4)$$

Let f be a collective equilibrium. Then $V_i(f) > 0$ for and each player i . For suppose that $V_i(f) = 0$ for some i . Since $u_i(x_i) > 0$ we must then have $f(x_i) = 0$ and $V_i(f[x_i]) = 0$, which contradicts condition (8.3). Now for each $s \in S$ the collective strategy $f[s]$ is also a collective equilibrium. Hence $V_i(f[s]) > 0$ for each player i and each $s \in S$.

The fact that $V_i(f) > 0$ for each player i implies that $\text{car}(f)$ contains at least two distinct policies. We now argue that $x_1 \in \text{car}(f)$. Suppose on the contrary. Then $\text{car}(f) = \{x_2, x_3\}$. Take any sequence $s = (s_0, \dots, s_t) \in S$ such that $s_t = x_3$ and $f(s) = 1$. Condition (8.4) implies that $u_3(x_3) \geq V_3(f[s])$. On the other hand since $\text{car}(f[s]) \subset \text{car}(f) = \{x_2, x_3\}$, it holds that $V(f[s])$ is a convex combination of the vectors $u(x_2)$ and $u(x_3)$. Since $u_3(x_3) < u_3(x_2)$ we necessarily have $V(f[s]) = u(x_3)$. But this contradicts a conclusion of the preceding paragraph since $u_2(x_3) = 0$.

Now suppose f is as in Definition 8.2, and let f_{x_1} be given by f_Y for some $Y \subset X$. Take any sequence $s = (s_0, \dots, s_t) \in S$ such that $s_t = x_1$ and $f(s) = 1$. We have $V(f[s]) = V(f_Y)$. Since by the result of the second paragraph $V_i(f_Y) > 0$ for each player i the set Y contains at least two alternatives. By condition (8.4) we have $u_1(x_1) \geq V_1(f[s]) = V_1(f_Y)$. This inequality, together with the fact that Y contains 2 or 3 alternatives implies that $Y = \{x_1, x_2\}$. In particular, $f(s \oplus x_3) = 0$. We obtain a contradiction since $V_3(f[s \oplus x_3]) = V_3(f_Y) = 7/2 < 4 = u_3(x_3)$, so by (8.3) it should hold that $f(s \oplus x_3) = 1$.

We conclude that there is no strategy f satisfying (8.3) and (8.4) that is induced by stationary game-plans. \square

Since the game of Example 4.4 has no collective equilibrium that is induced by stationary game-plans, we now shift our attention to the bigger set of collective strategies induced by two-step game-plans.

DEFINITION 8.5: For every $X^1 \subset X^2 \subset X$ and for every non-negative integer m , the collective strategy f_{X^1, m, X^2} is defined by setting, for every $s = (s_0, \dots, s_t) \in S$,

$$f(s) = \begin{cases} 1 & \text{if } [t < m \text{ and } s_t \in X^1] \text{ or } [t \geq m \text{ and } s_t \in X^2] \\ 0 & \text{otherwise.} \end{cases}$$

A collective strategy $f \in F$ is *two-step* if $f = f_{X^1, m, X^2}$ for some $X^1 \subset X^2 \subset X$ and non-negative integer m . The *threshold* of f is zero if $X^1 = X^2$ and is equal to the integer m otherwise.

According to a two-step strategy, the players wait for nature to choose an alternative from X^1 in the first m periods; as soon as such an alternative is chosen by nature, it is accepted. If no alternative from X^1 is chosen in the first m periods, then the players wait for an alternative from the bigger set X^2 .

DEFINITION 8.6: A collective strategy f is said to be *induced by two-step game-plans* if it is induced by a tuple of strategies $(f_j : j \in \{0\} \cup X)$ that are all two-step.

EXAMPLE 8.7: The strategy g in Example 7.1 is stationary, $g = f_X$, while the strategy h is two-step with the threshold of 1, $h = f_{\{x_2\}, 1, \{x_2, x_3\}}$. The collective strategy f is therefore induced by two-step game-plans. We now show that f is a collective equilibrium in the game Γ defined in Example 4.4. We verify that f satisfies Conditions (8.3) and (8.4).

Consider $s = (s_0, \dots, s_t) \in S$ such that $f(s) = 1$. Assume first that $s_t \in \{x_2, x_3\}$. After the rejection of s_t , the play of the game continues in accordance with the strategy g , hence $V(f[s]) = V(g) = (4/3, 11/3, 11/3)$. For players $i \in \{2, 3\}$, we have that $u_i(x_i) = 4 > 11/3 = V_i(f[s])$, so Condition (8.4) is satisfied. Now assume $s_t = x_1$. After the rejection of x_1 , the play of the game continues in accordance with the strategy h which results in a payoff of 1 to player 1, thus $u_1(x_1) = 1 = V_1(h) = V_1(f[s])$. This shows that Condition (8.4) is satisfied.

Now take some $s = (s_0, \dots, s_t) \in S$ such that $f(s) = 0$. Notice that rejections are only prescribed by the punishment strategy h . It follows that either we are in the first round of h and $s_t \in \{x_1, x_3\}$ or we are at least in the second round of h and $s_t = x_1$. In either case, the continuation play after s prescribes the acceptance of alternatives x_2 and x_3 and the rejection of x_1 , so $V(f[s]) = (3/2, 2, 11/2)$. We see that $u_1(x_1) = 1 < 3/2 = V_1(f[s])$ and $u_3(x_3) = 4 < 11/2 = V_3(f[s])$. Hence f satisfies Condition (8.3). \square

9 Existence of equilibria

The collection of all two-step game-plans is a countable set. Our next step is to provide an upper bound on the threshold of two-step game-plans that is sufficient to demonstrate the existence of a collective equilibrium. This restriction of the threshold leads to a finite set of two-step game-plans.

LEMMA 9.1: *There exists a natural number M with the following property: For all sets X^1 and X^2 if $\emptyset \neq X^1 \subset X^2 \subset X$ then $\text{SD}(f_{X^1}) \subset \text{SD}(f_{X^1, M, X^2})$.*

PROOF: We explicitly define the number M with the desired property. Let

$$\bar{u} = \max\{u_i(x) \mid i \in N, x \in X\}$$

be the maximum utility level reached by any alternative. By our assumption that there is at least one alternative with a strictly positive payoff for some player, we have $\bar{u} > 0$. Define

$$\epsilon = \min \left\{ u_i(x) - V_{x,i}(f_Y) \mid \begin{array}{l} (x, i, Y) \in X \times N \times 2^X \text{ such that} \\ u_i(x) - V_{x,i}(f_Y) > 0 \end{array} \right\}.$$

Notice that ϵ is well-defined and is positive since the sets N and X are finite.

For $x \in X$, let M_x be the least natural number such that for every alternative $x \in X$ there is probability greater than $1 - \epsilon/\bar{u}$ that x is selected at least once in the next $M - 1$ rounds if the current alternative is x . The irreducibility of μ implies that such an M_x exists. We define $M = \max_{x \in X} M_x$.

Since f_{X^1} and f_{X^1, M, X^2} coincide in periods $0, \dots, M-1$, and the probability that period M is reached conditional on a rejection of x in period 0 is less than ϵ/\bar{u} under the collective strategy f_{X^1} , we have that

$$|V_{x,i}(f_{X^1, M, X^2}) - V_{x,i}(f_{X^1})| < (1 - \frac{\epsilon}{\bar{u}}) 0 + \frac{\epsilon}{\bar{u}} \bar{u} = \epsilon.$$

Now take an $x \in \text{SD}(f_{X^1})$ and let $C = \{i \in N \mid u_i(x) > V_{x,i}(f_{X^1})\}$. Then we have $u_i(x) - V_{x,i}(f_{X^1}) \geq \epsilon$ for every $i \in C$ by definition of ϵ . Since $V_{x,i}(f_{X^1}) > V_{x,i}(f_{X^1, M, X^2}) - \epsilon$, we conclude that $u_i(x) > V_{x,i}(f_{X^1, M, X^2})$ for every $i \in C$, so $x \in \text{SD}(f_{X^1, M, X^2})$, as desired. \square

We are now in a position to state the main result of the paper.

THEOREM 9.2: *The game Γ has a collective equilibrium f induced by the collection of strategies $(f_0, f_x : x \in X)$ such that f_0 is a stationary strategy and for every $x \in X$ the strategy f_x is two-step with a threshold at most M .*

The proof of the theorem consists of two parts. The first part can be thought of as iterated elimination of unacceptable alternatives: We inductively reduce the set of alternatives by eliminating those alternatives that cannot be accepted in any collective equilibrium that is induced by two-step game-plans with a threshold of at most M . The reduction of the set of acceptable alternatives in turn results in the elimination of collective strategies that can be used as punishment game-plans. At the end of this process we are left with a set of acceptable alternatives, and, corresponding to each surviving alternative, a two-step game-plan that is a suitable punishment game-plan following an unlawful rejection of that alternative.

In the second part of the proof we use these building blocks to construct a collective equilibrium with the desired properties. Set $Y_0 = X$. Define F_0 to be a subset of F consisting of strategies f such that

[1] f is two-step with a threshold at most M , and

[2] f satisfies Condition (6.2): for each (s_0, \dots, s_t) in S : if $s_t \in \text{SD}(f[s_0, \dots, s_t])$ then $f(s_0, \dots, s_t) = 1$.

Notice that while a strategy in F_0 satisfies Condition (6.2) of collective equilibrium, it may well violate Condition (6.3). The set F_0 is non-empty as f_X is an element of it. We now give a simple criterion to check whether a two-step strategy satisfies Condition (6.2).

LEMMA 9.3: Consider some $X^1, X^2 \subset X$ with $X^1 \subset X^2$. For $m = 0, 1, \dots$ define the collective strategy $g_m = f_{X^1, m, X^2}$.

1. The collective strategy g_0 satisfies Condition (6.2) if and only if $\text{SD}(g_0) \subset X^2$.
2. Consider some non-negative integer m . The collective strategy g_{m+1} satisfies Condition (6.2) if and only if g_m satisfies Condition (6.2) and $\text{SD}(g_m) \subset X^1$.

PROOF: To prove the first claim observe that $g_0 = f_{X^2}$. This means that $g_0(s_0, \dots, s_t) = 1$ if and only if $s_t \in X^2$. Moreover $g_0[s_0, \dots, s_t] = g_0$. The result follows at once.

To prove the second claim we notice that for each $x \in X$ we have $g_{m+1}[x] = g_m$. Moreover for each sequence (s_0, \dots, s_t) with $t \geq 1$ we have $g_{m+1}(s_0, s_1, \dots, s_t) = g_m(s_1, \dots, s_t)$ and $g_{m+1}[s_0, s_1, \dots, s_t] = g_m[s_1, \dots, s_t]$.

To prove the "only if" part suppose g_{m+1} satisfies condition (6.2). For a sequence (s_1, \dots, s_t) if $s_t \in \text{SD}(g_m[s_1, \dots, s_t])$ then $s_t \in \text{SD}(g_{m+1}[s_0, s_1, \dots, s_t])$, where s_0 is any element of X . Hence $g_{m+1}(s_0, s_1, \dots, s_t) = 1$, therefore $g_m(s_1, \dots, s_t) = 1$. We conclude that g_m also satisfies (6.2). If $x \in \text{SD}(g_m)$ then $x \in \text{SD}(g_{m+1}[x])$ so $g_{m+1}(x) = 1$ and therefore $x \in X^1$.

To prove the "if" part suppose that $s_t \in \text{SD}(g_{m+1}[s_0, \dots, s_t])$. If $t = 0$ then $s_0 \in \text{SD}(g_m)$, hence $s_0 \in X^1$, and therefore $g_{m+1}(s_0) = 1$. If $t \geq 1$ then $s_t \in \text{SD}(g_m[s_1, \dots, s_t])$. Since g_m satisfies (6.2) we have $g_m(s_1, \dots, s_t) = 1$ and therefore $g_{m+1}(s_0, s_1, \dots, s_t) = 1$. \square

For every $k \in \mathbb{N}$, we inductively define

$$Y_k = \bigcup_{f \in F_{k-1}} \text{WD}(f)$$

$$F_k = \{f \in F_0 \mid \text{car}(f) \subset Y_k\}.$$

The alternatives for which there is a suitable punishment game-plan in F_{k-1} , i.e. those alternatives y for which there is $f \in F_{k-1}$ such that $y \in \text{WD}(f)$, are collected in the set Y_k . Next the set F_k is defined as those collective strategies f_{X^1, m, X^2} in F_0 where alternatives

outside Y_k are never accepted, so both X^1 and X^2 are subsets of Y_k . From there one defines the set Y_{k+1} , and so on.

LEMMA 9.4: *For every $k \in \mathbb{N}$ it holds that $Y_k \subset Y_{k-1}$ and $F_k \subset F_{k-1}$.*

PROOF: The proof is by induction on k . It is clear that $Y_1 \subset Y_0$ and $F_1 \subset F_0$.

Assume for some $k \in \mathbb{N}$ we have shown that $Y_k \subset Y_{k-1}$ and $F_k \subset F_{k-1}$. We complete the proof by showing that $Y_{k+1} \subset Y_k$ and $F_{k+1} \subset F_k$.

If $y \in Y_{k+1}$, then there is $f \in F_k$ such that $y \in \text{WD}(f)$. In view of the induction hypothesis, we have $f \in F_{k-1}$ and hence $y \in Y_k$. This proves that $Y_{k+1} \subset Y_k$. If $f \in F_{k+1}$, then $\text{car}(f) \subset Y_{k+1}$. By the previous step, we have $\text{car}(f) \subset Y_k$, so $f \in F_k$. This shows that $F_{k+1} \subset F_k$. \square

We define the sets

$$Y_* = \bigcap_{k=0}^{\infty} Y_k \text{ and } F_* = \bigcap_{k=0}^{\infty} F_k.$$

Since Y_0, Y_1, \dots is a decreasing sequence of finite sets, we have $Y_k = Y_{k+1} = Y_*$ for k sufficiently large, and therefore $F_k = F_{k+1} = F_*$. It follows that

$$\begin{aligned} Y_* &= \bigcup_{f \in F_*} \text{WD}(f) \\ F_* &= \{f \in F_0 \mid \text{car}(f) \subset Y_*\}. \end{aligned}$$

THEOREM 9.5: *Suppose the stationary collective strategy f_Y is a collective equilibrium. Then $f_Y \in F_*$ and $Y \subset Y_*$.*

PROOF: By assumption f_Y satisfies Conditions (6.2) and (6.3). Hence f_Y is an element of F_0 . Suppose we have shown that f_Y is an element of F_k . For each $x \in Y$ we have $f_Y(x) = 1$, so by condition (6.2) $x \in \text{WD}(f_Y[x])$. By stationarity we have $f_Y[x] = f_Y$, so $x \in Y_{k+1}$. We have thus shown that $Y \subset Y_{k+1}$. Since $\text{car}(f_Y) = Y$ we conclude that $f_Y \in F_{k+1}$. The result follows. \square

Now we have seen earlier that if the policy x has a core property then $f_{\{x\}}$ is a stationary collective equilibrium. We thus obtain the following.

COROLLARY 9.6: *The core is a subset of Y_* .*

Our next aim is to show the non-emptiness of Y_* .

THEOREM 9.7: *For $k = 0, 1, \dots$, the set Y_k is non-empty and $f_{Y_k} \in F_k$.*

PROOF: The proof is by induction on k .

It clearly holds that the set $Y_0 = X$ is non-empty and $f_{Y_0} \in F_0$. Assume that for some non-negative integer k we have shown that $Y_k \neq \emptyset$ and $f_{Y_k} \in F_k$.

We prove first that $Y_{k+1} \neq \emptyset$. Suppose on the contrary that $Y_{k+1} = \emptyset$. Let x be an element of Y_k and let $g_m = f_{\{x\}, m, Y_k}$.

We claim that g_0, \dots, g_M are elements of F_k . By the induction hypothesis it holds that $g_0 = f_{Y_k} \in F_k$. Assume that, for some $m \in \{0, \dots, M-1\}$, we have shown that $g_m \in F_k$. Then $\text{SD}(g_m) \subset \text{WD}(g_m) \subset Y_{k+1} = \emptyset$. Applying Lemma 9.3 we conclude that g_{m+1} satisfies Condition (6.2), so is an element of F_0 . Clearly g_{m+1} is carried by the set Y_k , hence is an element of F_k .

The alternative x does not have a core property, for otherwise it would be an element of Y_* (by Theorem 9.6) and hence an element of Y_{k+1} which is empty by supposition. Hence the set $\text{SD}(x)$ is non-empty. Recall that $\text{SD}(x) = \text{SD}(f_{\{x\}})$. We have inclusions

$$\text{SD}(f_{\{x\}}) \subset \text{SD}(f_{\{x\}, M, Y_k}) = \text{SD}(g_M) \subset \text{WD}(g_M) \subset Y_{k+1}$$

where the leftmost inclusion holds by Lemma 9.1 and the rightmost inclusion because g_M has been shown to be an element of F_k . But this contradicts our supposition that Y_{k+1} is empty.

We have thus proven that Y_{k+1} is non-empty.

We show next that $f_{Y_{k+1}} \in F_{k+1}$. Let $g_m = f_{Y_{k+1}, m, Y_k}$.

We claim that g_0, \dots, g_M are elements of F_k . Indeed, $g_0 = f_{Y_k} \in F_k$ by the induction hypothesis. Assume that, for some $m \in \{0, \dots, M-1\}$, we have shown that $g_m \in F_k$. Then $\text{SD}(g_m) \subset \text{WD}(g_m) \subset Y_{k+1}$. Applying Lemma 9.3.2, we conclude that $g_{m+1} \in F_0$. Clearly g_{m+1} is carried by the set Y_k , so belongs to F_k .

Now we have the inclusions

$$\text{SD}(f_{Y_{k+1}}) \subset \text{SD}(f_{Y_{k+1}, M, Y_k}) = \text{SD}(g_M) \subset \text{WD}(g_M) \subset Y_{k+1}$$

where the leftmost inclusion is by Lemma 9.1, and the rightmost inclusion follows since g_M has been shown to be an element of F_k . By Lemma 9.3.1 it follows that $f_{Y_{k+1}} \in F_{k+1}$, as desired.

This completes the induction step. \square

Corollary 9.8 follows immediately from Theorem 9.7.

COROLLARY 9.8: *The set Y_* is non-empty and $f_{Y_*} \in F_*$.*

We now construct a collective equilibrium that satisfies the properties required in Theorem 9.2. For each $x \in Y_*$, choose some $f_x \in F_*$ with $x \in \text{WD}(f_x)$, and for $x \in X \setminus Y_*$ we define $f_x = f_{Y_*}$. Let f_0 be any element of F_* . As a special case we can take f_0 to be the stationary strategy f_{Y_*} . The collective strategy induced by the tuple $(f_0, f_x : x \in X)$ is denoted by f_* .

THEOREM 9.9: *The strategy f_* is a collective equilibrium.*

PROOF: We use Lemma 7.2.

To prove that f_* satisfies Condition (6.3) take an $s = (s_0, \dots, s_t) \in S$ with $f_*(s) = 1$. Then $P(f_*[s]) = P(f_{s_t})$ and therefore $V_x(f_*[s]) = V_x(f_{s_t})$ for every $x \in X$. Hence $\text{WD}(f_*[s]) = \text{WD}(f_{s_t})$. Since s_t is an element of $\text{WD}(f_{s_t})$ by the choice of f_{s_t} we conclude that Condition (6.3) is satisfied.

To prove that f_* satisfies (6.2) take an $s = (s_0, \dots, s_t) \in S$ with $f_*(s) = 0$. Assume first that $f_*(s_0, \dots, s_k) = 0$ for every $k = 0, \dots, t-1$. We have $0 = f_*(s) = f_0(s)$. Moreover, $P(f_*[s]) = P(f_0[s])$. Consequently $V_x(f_*[s]) = V_x(f_0[s])$ for each $x \in X$ so $\text{SD}(f_*[s]) = \text{SD}(f_0[s])$. Since f_0 satisfies Condition (6.2), s_t is not an element of $\text{SD}(f_0[s])$, hence also not of $\text{SD}(f_*[s])$, as desired.

Assume next that $f_*(s_0, \dots, s_k) = 1$ for some $k = 0, \dots, t-1$. Let k^* be the largest k with this property. Then $0 = f_*(s) = f_{s_{k^*}}(s_{k^*+1}, \dots, s_t)$. Moreover, $P(f_*[s]) = P(f_{s_{k^*}}[s_{k^*+1}, \dots, s_t])$. Consequently $V_x(f_*[s]) = V_x(f_{s_{k^*}}[s_{k^*+1}, \dots, s_t])$ for each $x \in X$ so $\text{SD}(f_*[s]) = \text{SD}(f_{s_{k^*}}[s_{k^*+1}, \dots, s_t])$. Since $f_{s_{k^*}}$ satisfies Condition (6.2), s_t is not an element of $\text{SD}(f_{s_{k^*}}[s_{k^*+1}, \dots, s_t])$, hence also not of $\text{SD}(f_*[s])$, as desired. \square

10 Three alternatives

In this section we consider the special case of the model with three alternatives, where the alternatives give rise to a Condorcet cycle: $\text{SD}(x_2) = \{x_1\}$, $\text{SD}(x_3) = \{x_2\}$, and $\text{SD}(x_1) = \{x_3\}$. We do not make any further assumption regarding the number of players, the collection of decisive coalitions, or the Markov process by which alternatives are selected. Example 4.4 provides an illustration for the case of a quitting game with three players and time and history independent probabilities by which alternatives are selected, and Example 4.5 for the case when decision making takes place by means of majority voting.

The result claims that the set Y_* as constructed in the previous section consists of all three alternatives. The equilibrium of Theorem 9.9 therefore has the immediate acceptance

property: whatever alternative is chosen by nature in period zero is accepted.

THEOREM 10.1: *Let $X = \{x_1, x_2, x_3\}$ and assume that $\text{SD}(x_2) = \{x_1\}$, $\text{SD}(x_3) = \{x_2\}$, and $\text{SD}(x_1) = \{x_3\}$. Then $Y_* = X$.*

PROOF: It holds by definition that $Y_0 = X$. We show that $Y_1 = X$ as well, from which it follows that $Y_k = X$ for every $k = 0, 1, \dots$.

We prove first that $x_1 \in Y_1$. Since $Y_0 = X$, it clearly holds that $f_X \in F_0$.

Assume $x_1 \in \text{WD}(f_X)$. We then obviously have $x_1 \in Y_1$.

Assume $x_1 \notin \text{WD}(f_X)$. For $m = 0, \dots, M$, we define the collective strategy $g_m = f_{\{x_2\}, m, \{x_2, x_3\}}$. It holds that $x_1 \in \text{SD}(x_2) = \text{SD}(f_{\{x_2\}})$. By the choice of the threshold M we know that $\text{SD}(f_{\{x_2\}}) \subset \text{SD}(g_M)$ and hence $x_1 \in \text{SD}(g_M)$. Now let k be the smallest number in $\{0, \dots, M\}$ such that $x_1 \in \text{WD}(g_k)$.

We claim that g_0, \dots, g_k belong to F_0 . To prove this we only need to show that g_0, \dots, g_k all satisfy Condition (6.2).

We show that $g_0 = f_{\{x_2, x_3\}}$ satisfies Condition (6.2). Thus we have to prove that $x_1 \notin \text{WD}(f_{\{x_2, x_3\}})$. Indeed,

$$V_{x_1}(f_{\{x_2, x_3\}}) = \frac{\mu(x_2 | x_1)}{1 - \mu(x_1 | x_1)}u(x_2) + \frac{\mu(x_3 | x_1)}{1 - \mu(x_1 | x_1)}u(x_3),$$

while

$$V_{x_1}(f_X) = \mu(x_1 | x_1)u(x_1) + \mu(x_2 | x_1)u(x_2) + \mu(x_3 | x_1)u(x_3).$$

Rearranging the terms we find that if $u_i(x_1) \geq V_{x_1, i}(f_{\{x_2, x_3\}})$ then $u_i(x_1) \geq V_{x_1, i}(f_X)$. Since $x_1 \notin \text{WD}(f_X)$, we have $x_1 \notin \text{WD}(f_{\{x_2, x_3\}})$, as desired.

Assume g_m satisfies Condition (6.2) for some integer $m \in \{0, \dots, k-1\}$. We know that $x_1 \notin \text{SD}(g_m)$ by the choice of k and since $m < k$. We claim that also $x_3 \notin \text{SD}(g_m)$. Indeed, $V_{x_3}(g_m)$ is a convex combination of $u(x_2)$ and $u(x_3)$. Hence if $u_i(x_3) > V_{x_3, i}(g_m)$ then necessarily $u_i(x_3) > u_i(x_2)$. Since $x_3 \notin \text{SD}(x_2)$ we conclude that $x_3 \notin \text{SD}(g_m)$. We have shown that $\text{SD}(g_m) \subset \{x_2\}$. By Lemma 9.3.2 it follows that g_{m+1} satisfies Condition (6.2).

We have thus shown that $g_k \in F_0$. Since k is chosen such that $x_1 \in \text{WD}(g_k)$, it follows that $x_1 \in Y_1$. It follows by symmetry that x_2 and x_3 belong to Y_1 . \square

When we apply the construction in the proof of Theorem 10.1 to Example 4.4, we find the collective strategy as defined in Example 7.1. The main game-plan is determined by the stationary collective strategy $f_0 = f_X$ so all alternatives are accepted in period 0. Since $\text{WD}(f_X) = \{x_2, x_3\}$, the punishment plays following unlawful rejections of x_2 and x_3 are given by $f_{x_2} = f_{x_3} = f_X$. If we define $g_m = f_{\{x_2\}, m, \{x_2, x_3\}}$, then the smallest value for m for which $x_1 \in \text{WD}(g_m)$ is equal to 1. This leads to the choice $f_{x_1} = g_1$ in Example 7.1.




Figure 2: The set X_*

11 Three player allocation problem

In this section we consider a situation where three players have to allocate K identical units of an indivisible good, using a simple majority voting rule. Thus the set of alternatives is

$$X = \{x \in \mathbb{N}^3 : x_1 + x_2 + x_3 = K\},$$

where \mathbb{N} is the set of natural numbers including zero, and the utility functions are given by $u_i(x) = x_i$. The set $\mathcal{C}(x)$ consists of all sets containing at least two players, for every policy x . We assume that each policy arrives with equal probability, thus $\mu(x|y) = 1/|X|$ for all pairs of policies (x, y) . We refer to this situation as a K -unit allocation problem.

Let X_* denote the set of allocations where at least two players each receive at least $K/3$ units of the good, that is

$$X_* = \{x \in X : |\{i \in N : x_i \geq K/3\}| \geq 2\}.$$

The set X_* is depicted in Figure 2. It is non-empty if $K \geq 2$. The set X_* is symmetric so that $V(f_{X_*}) = (K/3, K/3, K/3)$. Furthermore $X_* = \text{WD}(f_{X_*})$. This establishes the following result.

LEMMA 11.1: *If $K \geq 2$ then f_{X_*} is a collective equilibrium.*

This implies that $X_* \subset Y_*$. On the other hand, the set Y_* does not include allocations where one player gets everything.

LEMMA 11.2: *Let $K \geq 3$. Then the points $(K, 0, 0)$, $(0, K, 0)$, and $(0, 0, K)$ are not elements of Y_* .*

PROOF: To obtain a contradiction suppose that $(K, 0, 0)$ is an element of Y_* . Then there is a function $f \in F_*$ such that $V_2(f) = 0$ or $V_3(f) = 0$. Let $f = f_{X^1, m, X^2}$ and without loss of generality suppose $V_3(f) = 0$. But then $x_3 = 0$ for each $x \in X^2$. In particular, the allocation $(K-1, 0, 1)$ is not in X^2 , and hence not an element of $SD(V(f_{X^2}))$. Since $V_3(f_{X^2}) = 0$ this implies that $K-1 \leq V_1(f_{X^2})$. A similar argument shows that $K-1 \leq V_2(f_{X^2})$. Adding these latter inequalities up yields $K < 2(K-1) \leq V_1(f_{X^2}) + V_2(f_{X^2})$, a contradiction. \square

Let α denote a player's worst payoff on a collective strategy from the set F_* :

$$\alpha = \min\{V_i(f) : f \in F_*\}. \quad (11.1)$$

This expression does not depend on i by the symmetry of the game. We know that $\alpha \leq K/3$ since $V_i(f_{X_*}) = K/3$. We have the inclusion

$$Y_* \subset \{x \in X : |\{i \in N : x_i \geq \alpha\}| \geq 2\}. \quad (11.2)$$

To see this notice that if $y \in Y_*$ then there is an $f \in F_*$ such that $y \in \text{WD}(f)$. Hence there are at least two players i such that $y_i \geq V_i(f) \geq \alpha$. We do not know whether in general the equality in (11.2) holds. The following result implies that for K large, all equilibria satisfying the conditions of Theorem 9.2 give a payoff of approximately $K/3$.

THEOREM 11.3: *Let $\alpha(K)$ be defined by (11.1) in the K -unit allocation problem. Then*

$$\lim_{K \rightarrow \infty} \frac{\alpha(K)}{K} = \frac{1}{3}.$$

Let Z denote the set of non-negative vectors $z = (z_1, z_2, z_3)$ with $z_1 + z_2 + z_3 = K$. For each $z \in Z$ and $\alpha \in [0, K/3]$ let

$$\phi(z, \alpha) = \frac{1}{|X|} \sum_{x \in X} \phi_x(z, \alpha)$$

where

$$\phi_x(z, \alpha) = \begin{cases} z_3 & \text{if } x_1 < \alpha \text{ and } x_3 < \alpha, \\ z_2 & \text{if } x_2 < \alpha \text{ and } x_3 < \alpha, \\ z_3 & \text{if } x_1 \leq z_1 \text{ and } x_2 \leq z_2, \\ x_3 & \text{otherwise.} \end{cases}$$

LEMMA 11.4: *There exists a $z \in Z$ such that*

$$\alpha \leq z_1, \alpha \leq z_2, \alpha \leq z_3,$$

$$\phi(z, \alpha) \leq \alpha.$$

PROOF: Let $f_{X^1, m, X^2} \in F_*$ be such that $\alpha = V_3(f_{X^1, m, X^2})$. Let $z = V(f_{X^1, m-1, X^2})$. If $m = 0$ we set $X^1 = X^2$ and we let $f_{X^1, m-1, X^2}$ denote the strategy f_{X^2} . The inequalities $\alpha \leq z_i$ hold because the strategies $f_{X^1, m-1, X^2}$ and f_{X^2} are elements of F_* .

Write

$$V_3(f_{X^1, m, X^2}) = \frac{1}{|X|} \sum_{x \in X} \gamma_x,$$

where

$$\gamma_x = \begin{cases} x_3 & \text{if } x \in X^1 \\ z_3 & \text{otherwise.} \end{cases}$$

We prove that $\phi_x(z, \alpha) \leq \gamma_x$ for every $x \in X$. Take an $x \in X$. If $x_3 = z_3$ then $\phi_x(z, \alpha) = \gamma_x$. Thus assume that $x_3 \neq z_3$. Consider two cases:

Case 1: $\phi_x(z, \alpha) = z_3$ and $\gamma_x = x_3$. Then $x \in X^1$. Since $X^1 \subset Y_*$ we know that the inequalities $\alpha \leq x_i$ hold for at least two players i . Since also $\phi_x(z, \alpha) = z_3$ we must have $x_1 \leq z_1$ and $x_2 \leq z_2$. This implies that $z_3 \leq x_3$, as desired.

Case 2: $\phi_x(z, \alpha) = x_3$ and $\gamma_x = z_3$. We know that x is not an element of X^1 . Since f_{X^1, m, X^2} is an element of F_* it holds that $\text{SD}(f_{X^1, m-1, X^2}) \subset X^1$. We conclude that x does not strictly dominate $f_{X^1, m-1, X^2}$. That is the inequality $z_i < x_i$ holds for one player i at most. On the other hand since $\phi_x(z, \alpha) = x_3$ we must have $z_1 < x_1$ or $z_2 < x_2$. Therefore $x_3 \leq z_3$, as desired. \square

LEMMA 11.5: *There exists a $z \in Z$ such that*

$$\alpha \leq z_1, \alpha \leq z_2, \alpha = z_3,$$

$$\phi(z, \alpha) \leq \alpha.$$

Notice that if the minimum in (11.1) is attained at a stationary strategy, then this result follows at once from our proof of the preceding lemma. The proof given below is needed if no stationary strategy in F_* gives payoff α .

PROOF: Let $z \in Z$ be as in the preceding lemma and let $\bar{z} = (z_1 + z_3 - \alpha, z_2, \alpha, \alpha)$. To prove the lemma it is sufficient to show that $\phi(\bar{z}, \alpha) \leq \phi(z, \alpha)$. To prove the latter inequality it is sufficient to show that $\phi_x(\bar{z}, \alpha) \leq \phi_x(z, \alpha)$ for each $x \in X$.

Fix an $x \in X$ and consider two cases:

Case 1: $[x_1 < \alpha$ and $x_3 < \alpha]$, or $[x_2 < \alpha$ and $x_3 < \alpha]$, or $[x_1 \leq z_1$ and $x_2 \leq z_2]$. Then $\phi_x(z, \alpha) = z_3$ and $\phi_x(\bar{z}, \alpha) = \bar{z}_3$, and the result follows since $\bar{z}_3 = \alpha \leq z_3$.

Case 2: otherwise. Then $\phi_x(z, \alpha) = x_3$. If $\phi_x(\bar{z}, \alpha) = \bar{z}_3$ we must have that $x_1 \leq \bar{z}_1$ and $x_2 \leq \bar{z}_2$. Hence $\bar{z}_3 \leq x_3$, as desired. \square

LEMMA 11.6: *It holds that*

$$\frac{K}{3} + \frac{2}{(K+1)(K+2)} \left(\alpha^3 - \left(\frac{K-\alpha}{2} + 1 \right)^3 \right) \leq \alpha.$$

PROOF: Take a z as in the preceding lemma. We estimate the value of $\phi(z, \alpha)$ from below. For a real number a such that $k < a < k+1$ where k is an integer let $\lfloor a \rfloor = k$ and $\lceil a \rceil = k+1$. For an integer k set $\lfloor k \rfloor = \lceil k \rceil = k$. Define the sets

$$\begin{aligned} C^0 &= \{x \in X : x_1 < \alpha, x_3 < \alpha\} &= \{x \in X : x_1 \leq \lceil \alpha \rceil - 1, x_3 \leq \lceil \alpha \rceil - 1\}, \\ C^1 &= \{x \in X : x_2 < \alpha, x_3 < \alpha\} &= \{x \in X : x_2 \leq \lceil \alpha \rceil - 1, x_3 \leq \lceil \alpha \rceil - 1\}, \\ C^2 &= \{x \in X : x_1 \leq z_1, x_2 \leq z_2\} &= \{x \in X : x_1 \leq \lfloor z_1 \rfloor, x_2 \leq \lfloor z_2 \rfloor\}, \\ C &= C^0 \cup C^1 \cup C^2. \end{aligned}$$

We have

$$\begin{aligned} |C^0| &= \lceil \alpha \rceil^2, & V_3(f_{C^0}) &= \frac{\lceil \alpha \rceil - 1}{2} \\ |C^1| &= \lceil \alpha \rceil^2, & V_3(f_{C^1}) &= \frac{\lceil \alpha \rceil - 1}{2} \\ |C^2| &= (\lfloor z_1 \rfloor + 1)(\lfloor z_2 \rfloor + 1), & V_3(f_{C^2}) &= \frac{K + K - \lfloor z_1 \rfloor - \lfloor z_2 \rfloor}{2}. \end{aligned}$$

Consider

$$V(f_C) = \sum_{i=0,1,2} \frac{|C^i|}{|C|} V(f_{C^i}).$$

We notice that

$$\frac{K}{3} = V_3(f_X) = \frac{|X \setminus C|}{|X|} V_3(f_{X \setminus C}) + \frac{|C|}{|X|} V_3(f_C).$$

We compute

$$\begin{aligned}
\phi(z, \alpha) &= \frac{|X \setminus C|}{|X|} V_3(f_{X \setminus C}) + \frac{|C|}{|X|} z_3 \\
&= \frac{K}{3} - \frac{|C|}{|X|} V_3(f_C) + \frac{|C|}{|X|} z_3 \\
&= \frac{K}{3} + \frac{1}{|X|} \sum_{i=0,1,2} |C^i| (z_3 - V_3(f_{C^i})).
\end{aligned}$$

In the remainder of the proof we estimate the summands in the latter expression from below. First consider the terms corresponding to $i = 0, 1$. Using the fact that $\lceil \alpha \rceil \geq \alpha \geq \lceil \alpha \rceil - 1$ and $z_3 = \alpha$ we obtain

$$\sum_{i=0,1} |C^i| (z_3 - V_3(f_{C^i})) = 2 \lceil \alpha \rceil^2 \left(\alpha - \frac{\lceil \alpha \rceil - 1}{2} \right) \geq \alpha^3.$$

Now consider the term corresponding to $i = 2$. Using the fact that $\lfloor z_i \rfloor \geq z_i - 1$ and that $z_1 + z_2 + \alpha = K$ we obtain

$$(z_3 - V_3(f_{C^2})) = z_3 - K + \frac{\lfloor z_1 \rfloor + \lfloor z_2 \rfloor}{2} \geq \alpha - K + \frac{z_1 + z_2 - 2}{2} = - \left(\frac{K - \alpha}{2} + 1 \right),$$

(notice that the expression on the right-hand side is negative). Using the fact that $ab \leq ((a+b)/2)^2$ and that $\lfloor z_i \rfloor \leq z_i$ we derive

$$|C^2| \leq \left(\frac{\lfloor z_1 \rfloor + 1 + \lfloor z_2 \rfloor + 1}{2} \right)^2 \leq \left(\frac{z_1 + z_2 + 2}{2} \right)^2 = \left(\frac{K - \alpha}{2} + 1 \right)^2.$$

Therefore,

$$|C^2| (z_3 - V_3(f_{C^2})) \geq - \left(\frac{K - \alpha}{2} + 1 \right)^3$$

At last noticing that $|X| = (K+1)(K+2)/2$ yields the desired estimate. \square

To complete the proof of the theorem suppose without loss of generality that the sequence $\alpha(K)/K$ converges, and let β be the limit. It is easy to see from the above lemma that the limit β satisfies

$$g(\beta) = \frac{1}{3} + 2\beta^3 - 2 \left(\frac{1 - \beta}{2} \right)^3 \leq \beta.$$

The function g carries the interval $[0, \frac{1}{3}]$ into itself and satisfies $0 < g'(\beta) < 1$ whenever $0 \leq \beta < \frac{1}{3}$ and $g'(\frac{1}{3}) = 1$, and $g(\frac{1}{3}) = \frac{1}{3}$. These properties imply that $\frac{1}{3}$ is the only solution element of $[0, \frac{1}{3}]$ satisfying the inequality $g(\beta) \leq \beta$.

12 Robustness of collective equilibria

Examples 4.4 and 4.5 show that collective stopping games may not have stationary equilibria. In the presence of breakdown, stationary equilibria are guaranteed to exist, but may be in mixed strategies as demonstrated by Example 4.6. It can be hard to compute the exact probabilities in a mixed strategy that are needed to make the players indifferent between accepting and rejecting a particular alternative, and, moreover, these probabilities change when the utilities of the alternatives change. We argue next that collective equilibria generated by two-step game-plans are robust to small changes in utilities and in Markov transition probabilities that govern the choices of the alternative.

To do so, we fix the set of players, alternatives, and the decision rules used, and parameterize a game $\Gamma(u, \mu)$ by the profile of utility functions u , and the the Markov process μ that governs the choice of alternatives. The profile of utility functions u is identified with a point in a non-negative orthant of $\mathbb{R}^{N \times X}$, denoted \mathcal{U} . In this section we drop the assumption that the the Markov process μ be irreducible. We allow μ to be any Markov process and identify it with the corresponding stochastic matrix of transition probabilities. The set of stochastic matrices is denoted by \mathcal{M} .

Consider a collective equilibrium f of the game $\Gamma(u_0, \mu_0)$. The strategy f is said to be *robust* if it is a collective equilibrium of the game $\Gamma(u, \mu)$ for each (u, μ) in some open neighborhood of (u_0, μ_0) . Let T denote the set of collective strategies induced by two-step game-plans.

THEOREM 12.1: *There is a set $\mathcal{G}_* \subset \mathcal{U} \times \mathcal{M}$ such that*

- [1] *For each $\mu \in \mathcal{M}$ the μ -section $\{u \in \mathcal{U} : (u, \mu) \in \mathcal{G}_*\}$ of the set \mathcal{G}_* is negligible as a subset of \mathcal{U} .*
- [2] *For each $(u_0, \mu_0) \in \mathcal{G}_*$, each collective equilibrium $f \in T$ of the game $\Gamma(u_0, \mu_0)$ is robust.*

A subset of \mathcal{U} is said to be negligible if its complement is nowhere dense and has Lebesgues measure zero.

One important implication of the theorem is that collective equilibria in T are robust to discounting. Recall that our setup accommodates discounted payoffs as a special case in which there is a breakdown alternative that is selected with probability ϵ and that yields a payoff of zero to every player. Consider the game $\Gamma(u)$ parameterized by the payoff function u only. Starting from the original game $\Gamma(u)$, define a new game $\Gamma_\epsilon(u)$ that has one additional alternative d . We let $u_i(d) = 0$ for each i , and $\mathcal{C}(d) = 2^N$. Define the

Markov process on $X \cup \{d\}$ by letting $\mu_\epsilon(x|y) = (1 - \epsilon)\mu(x|y)$ for each $(x, y) \in X$ and $\mu(d|y) = \epsilon$ for each $y \in X$. The game $\Gamma_\epsilon(u)$ can be thought of as the original game $\Gamma(u)$ with payoffs discounted by $1 - \epsilon$. The game $\Gamma_0(u)$ is equivalent to the original game $\Gamma(u)$.

The above theorem implies that for almost every $u \in \mathcal{U}$, given a collective equilibrium $f \in T$ of $\Gamma(u)$ there exists an $\epsilon_* \in (0, 1]$ such that f remains a collective equilibrium of $\Gamma_\epsilon(u)$ for each $\epsilon \in [0, \epsilon_*)$. This result provides an important justification for the approach taken in this paper: While we do not insist on discounted payoffs, we are able to always offer a solution for a discounted version of our game.

The key property of the strategies induced by two-step game plans needed for the proof of the theorem is the following.

LEMMA 12.2: *For each $f \in T$ the set $\{f[s] : s \in S\}$ is finite.*

PROOF: The result is true if f is a two-step strategy. Indeed, in this case the given set has cardinality of at most $m + 1$, where m is the threshold of f .

Now suppose f is induced by a tuple $\mathcal{F} = (f_0, f_x : x \in X)$ of two-step strategies. Then for each $s \in S$ the strategy $f[s]$ is induced by the tuple $(g, f_x : x \in X)$ where the strategy g is of the form $g = h[r]$ for h a member of the collection \mathcal{F} and some $r \in \{\emptyset\} \cup S$. By the observation in the previous paragraph there are at most finitely many of such strategies g , and the result follows. \square

We let $V_{x,i}^{u,\mu}(f)$ denote the payoff to player i in the game $\Gamma_x(u, \mu)$ if the play follows f .

LEMMA 12.3: *The function $V_{x,i}^{u,\mu}(f)$ is continuous in (u, μ) .*

PROOF: We have

$$V_{x,i}^{u,\mu}(f) = \sum_{y \in X} u_i(y) \nu^\mu(f, x; y), \quad (12.1)$$

where $\nu^\mu(x, f; y)$ denotes the probability that, in the game Γ_x the play according to f eventually leads to the acceptance of the alternative y . This probability can be written as

$$\nu^\mu(f, x; y) = \sum_{t=0}^{\infty} \nu^\mu(f, x; t, y)$$

where $\nu^\mu(x, f; t, y)$ is the probability that the play of the game Γ_x ends in period t with the acceptance of the alternative y , provided that the collective strategy f is being played. It is easy to see that $\nu^\mu(f, x; t, y)$ is continuous as a function of μ . The result follows. \square

For each $\mu \in \mathcal{M}$ we define the μ -section of the set \mathcal{G}_* , denoted by $\mathcal{U}(\mu)$, as follows. Let $T(\mu)$ denote the subset of T consisting of strategies $f \in T$ such that there is an $x \in X$ with $\nu^\mu(f, x; x) = 1$. If $f \in T(\mu)$ then clearly $V_{x,i}^{u,\mu}(f) = u_i(x)$ for all $u \in \mathcal{U}$ and all $i \in N$. For each $f \in T$ define the subset $\mathcal{U}_f(\mu)$ of \mathcal{U} as follows: If f is an element of $T(\mu)$ we set

$$\mathcal{U}_f(\mu) = \{u \in \mathcal{U} : u_i(x) \neq u_i(y) \text{ for all } (i, x, y) \in N \times X \times X \text{ with } x \neq y\}.$$

If f is not an element of $T(\mu)$ we set

$$\mathcal{U}_f(\mu) = \{u \in \mathcal{U} : u_i(x) \neq V_{x,i}^{u,\mu}(f) \text{ for all } (i, x) \in N \times X.\}$$

We define

$$\mathcal{U}(\mu) = \bigcap_{f \in T} \mathcal{U}_f(\mu) \text{ and } \mathcal{G}_* = \bigcup_{\mu \in \mathcal{M}} \mathcal{U}(\mu) \times \{\mu\}.$$

We complete the proof by showing that the set \mathcal{G}_* thus defined has properties [1] and [2] required by the theorem.

We show that the complement of $\mathcal{U}(\mu)$ is a negligible set. Since the set T is countable, it suffices to show that for each $f \in T$ the complement of $\mathcal{U}_f(\mu)$ is negligible.

First consider the case $f \in T(\mu)$. In this case the complement of $\mathcal{U}_f(\mu)$ is covered by finitely many sets of the form $\{u \in \mathcal{U} : u_i(x) = u_i(y)\}$ where $x \neq y$. Each such set is a lower-dimensional subspace of \mathcal{U} and hence a negligible set.

Suppose f is not in $T(\mu)$. The complement of $\mathcal{U}_f(\mu)$ is then covered by finitely many sets of the form $C = \{u \in \mathcal{U} : u_i(x) = V_{x,i}^{u,\mu}(f)\}$. Since the payoff $V_{x,i}^{u,\mu}(f)$ is linear in u (see Equation (12.1)), the set C is a linear subspace of \mathcal{U} . It is a proper subset of \mathcal{U} because $\nu^\mu(f, x; x) < 1$. We conclude that C is a lower-dimensional subspace of \mathcal{U} and is hence negligible.

Take a game $(\bar{u}, \bar{\mu}) \in \mathcal{G}_*$. For each $f \in T$ define an open neighborhood of $(\bar{u}, \bar{\mu})$ by letting

$$\mathcal{O}(f) = \bigcap_{i \in N} \bigcap_{x \in X} \left\{ (u, \mu) \in \mathcal{U} \times \mathcal{M} : \begin{array}{l} \bar{u}_i(x) > V_{x,i}^{\bar{u}, \bar{\mu}}(f) \text{ implies } u_i(x) > V_{x,i}^{u,\mu}(f) \\ \bar{u}_i(x) < V_{x,i}^{\bar{u}, \bar{\mu}}(f) \text{ implies } u_i(x) < V_{x,i}^{u,\mu}(f) \end{array} \right\}.$$

Invoking the definition of $\mathcal{U}(\bar{\mu})$ it is easy to see that for every game $(u, \mu) \in \mathcal{O}(f)$ the alternative $x \in X$ strictly (weakly) dominates f in the game $\Gamma(u, \mu)$ if and only if it does so in the game $\Gamma(\bar{u}, \bar{\mu})$.

Let $\bar{f} \in T$ be a collective equilibrium of the game $(\bar{u}, \bar{\mu})$. We argue that \bar{f} is robust. Define $\mathcal{O} = \bigcap \{\mathcal{O}(\bar{f}[s]) : s \in S\}$, the intersection being finite by Lemma 12.2. Take any

game $(u, \mu) \in \mathcal{O}$. By the remark in the previous paragraph, for every $s \in S$, the alternative x strictly (weakly) dominates $\bar{f}[s]$ in the game $\Gamma(u, \mu)$ if and only if it does so in the game $\Gamma(\bar{u}, \bar{\mu})$. It follows directly from Definition 6.2 that \bar{f} is a collective equilibrium of the game $\Gamma(u, \mu)$, as desired.

13 Incomplete information

In this section, we generalize our model by allowing the voting order to be probabilistic and history-dependent, and information to be incomplete. We continue to assume that players observe the moves by nature, but we do no longer require that a player observes the actions of other players.

We therefore consider a game Γ' defined as follows. In each period $t = 0, 1, \dots$ nature draws an alternative x from a set X according to the Markov process μ on X , where the selection in period 0 occurs according to the probability distribution μ^0 on X . Then nature selects the first player to vote on x , say player $i_1 \in N$. After Player i_1 casts a “y” or an “n” vote, nature chooses the next player to cast a vote, say player $i_2 \in N \setminus \{i_1\}$, who casts a “y” or an “n” vote. Next nature chooses player $i_3 \in N \setminus \{i_1, i_2\}$, and so on until all players have voted. The choice of a responder i_k may be probabilistic and may depend on the entire history of play. If the set of players who voted “y” is decisive, the game ends and player $i \in N$ receives utility $u_i(x)$. Otherwise the next period begins. Perpetual disagreement results in utility zero for all players.

Let H_i denote the set of finite histories after which it is player i 's turn to cast a vote. The description of the game Γ' is completed by specifying the information sets for every player: \mathcal{H}_i denotes the information partition of H_i for player i . Given a history $h \in H_i$, t_h denotes the current period at h , and corresponds to the number of alternatives that have been voted down so far. For $t = 0, \dots, t_h$, the alternative chosen by nature in period t is denoted by $x_t(h)$ and for $t = 0, \dots, t_h - 1$, we write $a_t^i(h)$ for the action taken by player i in period t . With respect to the information structure of the game we make the following assumption.

[A5] If h, h' belong to the same information set I_i in \mathcal{H}_i , then

1. $t_h = t_{h'}$,
2. for every $t = 0, \dots, t_h$, $x_t(h) = x_t(h')$,
3. for every $t = 0, \dots, t_h - 1$, $a_t^i(h) = a_t^i(h')$.

Since we allow for incomplete information, besides the case with sequential voting, we now also cover cases where some or all of the players cast their votes simultaneously. For example, we allow for the case where all players vote simultaneously and the votes are never revealed, as well as the case where all players vote simultaneously and all votes cast in a given period are revealed at the the end of that period.

A strategy of player i is a map $\sigma_i : \mathcal{H}_i \rightarrow [0, 1]$. A sequential equilibrium of Γ' specifies a joint strategy as well as the beliefs that each player has regarding the relative likelihood of different histories in the same information set. A belief system of player i is a collection of probability distributions $\beta_i(I_i)$ on the set I_i for each $I_i \in \mathcal{H}_i$. A sequential equilibrium is defined as a pair (σ, β) such that σ is a joint strategy and β a belief system such that σ is sequentially rational given β , and β is consistent with σ .

A joint strategy σ is said to be *sequentially rational* given the belief system β if for every player $i \in N$ and each of player i 's information sets $I_i \in \mathcal{H}_i$, the strategy σ_i is a best response against the joint strategy σ conditional on reaching information set I_i given the belief system $\beta_i(I_i)$. The condition of consistency encompasses the requirement that β be derived from σ using Bayes' rule and it also imposes some restrictions on how $\beta_i(I_i)$ should be defined for those information sets I_i that are reached, under σ , with probability zero. We deliberately avoid stating a more precise definition of consistency since the result below is valid for any definition of consistency satisfying the extremely mild requirement that for every joint strategy σ there is at least one system of beliefs consistent with σ . Hence we simply assume that some definition of consistency satisfying this requirement is used and we write $L(\sigma)$ to denote the set of beliefs consistent with the joint strategy σ .

Notice that each action-independent strategy σ_i of the game Γ naturally induces a strategy σ'_i in the game Γ' . For $I_i \in \mathcal{H}_i$, choose $h \in I_i$, define $s \in S$ by $s_t = x_t(h)$ for $t = 0, \dots, t_h$, and let $\sigma'_i(I_i)$ be equal to $\sigma_i(s)$. We can now state the following result:

THEOREM 13.1: *Let σ be a subgame perfect equilibrium of Γ in pure action-independent strategies. Then (σ', β) is a sequential equilibrium of Γ' whenever $\beta \in L(\sigma')$.*

The intuition for the result is straightforward: If all players use an action-independent strategy, then any additional information contained in the information partition \mathcal{H}_i is irrelevant: a strategy is a sequentially rational (or not) irrespective of the beliefs. We spell out the proof at the risk of proving the obvious. Consider an information set $I_i \in \mathcal{H}_i$, $h \in I_i$, and define $s \in S$ by $s_t = x_t(h)$ for $t = 0, \dots, t_h$. Since σ is a subgame perfect equilibrium of Γ , the strategy σ_i is a best response for player i in each subgame of Γ that follows the rejection of the alternatives s_0, \dots, s_t . It follows that σ_i is a best response in Γ' conditional on reaching I_i when player i puts probability 1 on any given information node

$h \in I_i$. It then follows that σ_i is a best response under any probability distribution that player i might have on I_i .

14 Conclusion

Collective stopping games form a rich class of games that include collective choice games, quitting games, and coalition formation games as special cases. Particular examples are of collective stopping games involve recruitment decisions by committees, business financing decisions, and oligopoly exit decisions. The standard equilibrium concept to analyze such games, stationary equilibrium in subgame perfect strategies, is problematic as for collective stopping games it may not exist at all or only exist in mixed strategies. Moreover, the existing literature has only demonstrated the existence of subgame perfect ϵ -equilibria. This triggers the question whether even subgame perfect equilibria may not always exist for such games.

In this paper we show not only that subgame perfect equilibria always exist in collective stopping games, but also that we can define a finite set of strategies which is guaranteed to contain a subgame perfect equilibrium. We require strategies to be pure and action-independent, where the latter property has the advantage that we can incorporate various degrees of incomplete information as far as the voting behavior of the players is concerned. Action-independence also allows us to summarize the play of the game by collective strategies. We require collective strategies to be simple and to be generated by two-step game-plans. Under a two-step game-plan, the game ends when as soon as nature draws an alternative from some set X_1 in the first m periods, or when nature draws an alternative from a bigger set X_2 in a later period. Since we can formulate an upper bound on m , this leads to the desired finite set of strategies which is shown to contain a subgame perfect equilibrium.

A particular example is the case of three alternatives, where utilities of the players and decision rules for the acceptance of an alternative are such that a Condorcet cycle results. Such examples can easily occur in the context of collective choice or in coalition formation processes. In such an example stationary equilibria may not exist at all. Subgame perfect equilibria in pure and action-independent strategies do exist, where the induced collective strategy is such that any alternative is accepted in period 0 and the unlawful rejection of an alternative is followed by a particular two-step punishment game-plan. We show that stationary punishment game-plans are not sufficient for equilibrium existence.

References

- ABREU, D. (1988), “On the Theory of Infinitely Repeated Games with Discounting,” *Econometrica*, 56, 383–396.
- BANKS, J., AND J. DUGGAN (2000), “A Bargaining Model of Collective Choice,” *American Political Science Review*, 94, 73–88.
- BLACKWELL, D., AND T. FERGUSON (1968), “The Big Match,” *Annals of Mathematical Statistics*, 39, 159–163.
- BLOCH, F. (1996), “Sequential Formation of Coalitions in Games with Externalities and Fixed Payoff Division,” *Games and Economic Behavior*, 14, 90–123.
- BLOCH, F., AND E. DIAMANTOUDI (2011), “Noncooperative Formation of Coalitions in Hedonic Games,” *International Journal of Game Theory*, 40, 263–280.
- BRETON, M. LE (1987), “On the Core of Voting Games,” *Social Choice and Welfare*, 4, 295–305.
- COMPTE, O., AND P. JEHIEL (2010A), “Bargaining and Majority Rules: A Collective Search Perspective,” *Journal of Political Economy*, 118, 189–221.
- COMPTE, O., AND P. JEHIEL (2010B), “The Coalitional Nash Bargaining Solution,” *Econometrica*, 78, 1593–1623.
- COX, G.W. (1984), “Non-Collegial Simple Games and the Nowhere Denseness of the Set of Preference Profiles Having a Core,” *Social Choice and Welfare*, 1, 159–164.
- DUGGAN, J., AND T. KALANDRAKIS (2012), “Dynamic Legislative Policy Making,” *Journal of Economic Theory*, 147, 1653–1688.
- FINK, A.M. (1964), “Equilibrium in a Stochastic n -Person Game,” *Journal of Science of Hiroshima University Series A-I*, 28, 89–93.
- FLESCH, J., J. KUIPERS, G. SCHOENMAKERS, AND K. VRIEZE (2010), “Subgame Perfection in Positive Recursive Games with Perfect Information,” *Mathematics of Operations Research*, 35, 193–207.
- FUDENBERG, D., AND J. TIROLE (1991), *Game Theory*, MIT Press, Cambridge, Massachusetts.
- HERINGS, P.J.J., AND H. HOUBA (2010), “The Condorcet Paradox Revisited,” METEOR Research Memorandum 10/09, Maastricht University, Maastricht, pp. 1–45.
- MASHIAH-YAAKOVI, A. (2009), “Periodic Stopping Games,” *International Journal of Game Theory*, 38, 169–181.
- MASKIN, E.S., AND J. TIROLE (2001), “Markov Perfect Equilibrium, I. Observable Actions,” *Journal of Economic Theory*, 100, 191–219.
- MAYNARD-SMITH, J. (1974), “The Theory of Games and the Evolution of Animal Conflicts,” *Journal of Theoretical Biology*, 47, 209–221.
- OKADA, A. (1996), “A Noncooperative Coalitional Bargaining Game with Random Proposers,” *Games and Economic Behavior*, 16, 97–108.

- PENN, E.M. (2009), "A Model of Farsighted Voting," *American Journal of Political Science*, 53, 36–54.
- PLOTT, C.R. (1967), "A Notion of Equilibrium and its Possibility Under Majority Rule," *American Economic Review*, 57, 787–806.
- RAY, D., AND R. VOHRA (1999), "A Theory of Endogenous Coalition Structures," *Games and Economic Behavior*, 26, 286–336.
- ROBERTS, K. (2007), "Condorcet Cycles? A Model of Intertemporal Voting," *Social Choice and Welfare*, 29, 383–404.
- RUBINSTEIN, A. (1979), "A Note about the "Nowhere Denseness" of Societies Having an Equilibrium under Majority Rule," *Econometrica*, 47, 511–514.
- SCHOFIELD, N. (1983), "Generic Instability of Majority Rule," *Review of Economic Studies*, 50, 695–705.
- SOBEL, M.J. (1971), "Noncooperative Stochastic Games," *The Annals of Mathematical Studies*, 42, 1930–1935.
- SOLAN, E. (2005), "Subgame-perfection in Quitting Games with Perfect Information and Differential Equations," *Mathematics of Operations Research*, 30, 51–72.
- SOLAN, E., AND N. VIEILLE (2001), "Quitting Games," *Mathematics of Operations Research*, 26, 265–285.
- TAKAHASHI, M. (1964), "Equilibrium Points of Stochastic Non-cooperative n -Person Games," *Journal of Science of Hiroshima University Series A-I*, 28, 95–99.