Games and Economic Behavior 68 (2010) 1–14

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

A sequential selection game with vetoes [☆]

Steve Alpern [a], Shmuel Gal [b], Eilon Solan [c,*]

[a] Mathematics Department, London School of Economics, London WC2A 2AE, UK
[b] Statistics Department, Haifa University, Haifa 31905, Israel
[c] The School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel

ARTICLE INFO

Article history:
Received 10 July 2008
Available online 23 May 2009

JEL classification:
C72
C73
D71

ABSTRACT

We study a selection game between two committee members (the players). They interview candidates sequentially and have to decide, after each interview, whether to hire the candidate or to interview the next candidate. Each player can either accept or reject the candidate, and if he rejects the candidate while the other accepts her, he can cast a veto. The candidate is hired if accepted by at least one player and not vetoed. The total number of vetoes available for each player are fixed in advance.

We prove the existence of a subgame perfect equilibrium if there are a finite number of candidates types. For a general candidate distribution we prove the existence of a subgame perfect ε-equilibrium. We exhibit situations in which a player prefers that the other player would have an extra veto, and even prefers to give one of his vetoes to the other player.

1. Introduction

Candidates for various positions, both in the private and public sectors, are often chosen by committees. Sometimes the committee interviews the candidates sequentially ('on-line') and has to decide, after each interview, whether the current candidate is acceptable or not; in the latter case, the candidate "disappears," and cannot be selected in the future if it turns out that she has been the best candidate. Such a case occurs, e.g., in the selection of juries in the American legal system, in which both sides have the right to veto candidates, or when a couple looks for an apartment to rent when demand far exceeds supply, and apartments are rented within hours.

In the present paper we study an on-line selection game between two committee members, who interview candidates sequentially and have to decide, after each interview, whether to accept the current candidate or to reject her. We model the process by which the committee decides on each candidate as follows. First one of the committee members, the *leader*, decides whether to accept or reject the candidate, and then the second member, the *follower*, observing the decision of the leader, decides whether to accept or reject the candidate. If both members rejected the candidate then she is not hired; if both members accepted the candidate then she is hired; if one of the players accepted the candidate while the other rejected her, the member who rejected the candidate has the right to cast a veto, in which case the candidate is not hired. If no veto was cast, the candidate is hired. We assume that the number of vetoes of each member is given in advance, as occurs, e.g., in jury selection (see Brams and Davis, 1976, 1978; DeGroot and Kadane, 1980). In many cases the number of vetoes that a member has is not fixed explicitly, yet, successive vetoes often raise pressure from other committee members against more vetoes. Thus, the number of vetoes of a member measures his relative strength, both his mental strength and his position within the organization.

☆ The research of Alpern and Gal was supported by NATO Collaborative Research Grant CLG 983583.

* Corresponding author.
E-mail addresses: s.alpern@lse.ac.uk (S. Alpern), sgal@univ.haifa.ac.il (S. Gal), eilons@post.tau.ac.il (E. Solan).

0899-8256/$ – see front matter
doi:10.1016/j.geb.2009.05.006

Each candidate is characterized by two numbers, her utility to each committee member, which are termed together the *type* of the candidate. The full type is known to both players. The goal of each player is to maximize his own expected utility from the selected candidate.

Our model is a game theoretic extension of the well-known "secretary problem" (see, e.g., Ferguson, 1989, 2001; Eriksson et al., 2007), in which quality, rather than rank of quality, is optimized, and each player's utility is his evaluation of the hired candidate. A different game variant of the "secretary problem," in which various employers compete on employees who are observed sequentially, was studied by, e.g., Enns and Ferenstein (1987), Ramsey and Szajowski (2005) and Sakaguchi (2005).

Our model is also related to the selection problem studied in DeGroot and Kadane (1980). The main difference is that whereas they study jury selection problems, and therefore allow only two actions (veto or no veto), we study an employee selection problem, and therefore allow three actions (accept, reject with veto, and reject without veto). As we show below, the two models exhibit quite a different behavior.

Our model was introduced in Alpern and Gal (2008), who proved that if a player's veto rights increase (while the other's remains fixed), then his equilibrium payoff cannot decrease. Alpern and Gal (2008) also presented a detailed analysis of this game for uniformly distributed candidate types, and established the existence of equilibria in threshold strategies for the game with no vetoes. Additionally, they study a discounted version of the game, and in this case they determined the optimal number of vetoes that a social planner should give the players, if the social planner's goal is to maximize the expected sum of utilities of both players.

In this paper we show that when there are finitely many types of candidates, a subgame perfect equilibrium in Markovian strategies always exists, while when there are infinitely many types of candidates, an equilibrium may fail to exist, yet a subgame perfect ε-equilibrium in Markovian strategies exists for every $\varepsilon > 0$. We then study the structure of Markovian equilibria, distinguishing between two types of equilibria: *progressive equilibria* and *helpful equilibria*. In a progressive equilibrium each player prefers that the other player would have fewer vetoes, and even accepts bad (to him) candidates in order to force the other player to "waste" a veto. In a helpful equilibrium players prefer the other to have more vetoes and hence reject candidates that the other player would otherwise have to veto.

We also present surprising phenomena that may occur in our games. DeGroot and Kadane (1980) present an example of selecting a jury of two participants, in which one of the committee members prefers to give his veto right to the other member. This phenomenon is called the "paradox of redistribution" by Fisher and Schotter (1978) or the "donation paradox" (see Felsenthal and Machover, 1995; Kadane et al., 1999; Holler and Napel, 2004). DeGroot and Kadane (1980, Theorem 3) prove that if the committee has to choose a single candidate, then the donation paradox cannot occur. Surprisingly, as we show by an example, a donation paradox may occur in our game. We also provide an example that shows that a worse population (good candidates are rarer) may yield a better outcome.

The paper is organized as follows. Section 2 describes the model. Section 3 proves the existence of a subgame perfect ε-equilibrium (SPE), in mixed Markovian strategies. Section 4 describes the structure of the equilibria in the game. Section 5 describes some counter-intuitive behavior of the game with no vetoes, including the donation paradox.

2. The model

We consider a sequential two-player selection game $G(i, j)$, that is given by two non-negative integers i and j (the number of vetoes of the two players) and a distribution F_2 on the unit square (the distribution of candidate). At each stage t, if the game has not terminated before, the players play the extensive form game that appears in Fig. 1: first Nature chooses a candidate $(x, y) \in [0, 1]^2$ according to the distribution F_2 independent of past play. Then player I decides whether to Accept or to Reject the candidate. If player I decides to accept the candidate, and if $j > 0$, player II decides whether to Veto the candidate, and then from the next stage on the game $G(i, j-1)$ is played, or whether Not to Veto the candidate, in which case the game terminates with a terminal payoff x to player I and y to player II. If player I decides to Reject the candidate, player II decides whether to Accept or to Reject the candidate. If player II decides to reject the candidate, the stage ends, and the game continues to the next stage. If player II decides to accept the candidate, player I decides whether to Veto the candidate, and then from the next stage on the game $G(i-1, j)$ is played, or whether Not to Veto the candidate, in which case the game terminates with a terminal payoff (x, y). If no candidate is ever hired, the payoff to both players is 0. In Fig. 1, the actions Accept, Reject, Veto and No Veto are abbreviated to a, r, v and n respectively.

A *strategy* for a player is a rule that indicates how to play at each stage given any past history. A strategy is *Markovian* if the mixed action to be played after every given history depends only on the type of the current candidate and on the number of vetoes left for both players.

A special class of strategies is the class of *threshold* strategies: accept a candidate if her worth for you is at least some threshold α, and reject her otherwise. In general the threshold α may be history dependent. If the player mixes when the worth of the candidate (for him) is equal to α, and accepts (respectively rejects) the candidate if her worth for him is higher (respectively lower) than α, we say that the strategy is a *weak threshold strategy*.

We study Markovian strategies, so the continuation payoffs in $G(i, j-1)$ and $G(i-1, j)$ are independent of past play. We denote the continuation payoff in $G(i, j-1)$ by $(u+, v-)$, and the continuation payoff in $G(i-1, j)$ by $(u-, v+)$.

Every pair of strategies (σ, τ) induces an (undiscounted) expected payoff for each player $k \in \{\mathrm{I}, \mathrm{II}\}$, which we denote by $\gamma^k(\sigma, \tau)$. For every $\varepsilon \geqslant 0$, an *ε-equilibrium* (σ, τ) is a pair of strategies such that no player can profit more than ε by

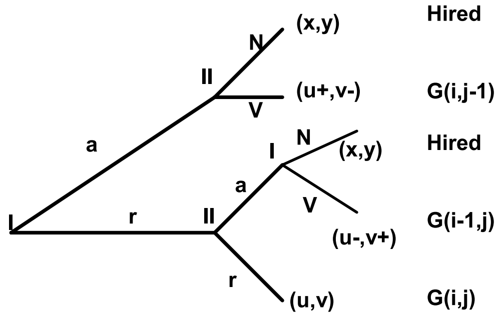

Fig. 1. Reduced tree for $G(i, j)$.

deviating. A pair of strategies is a *subgame-perfect ε-equilibrium* if for every finite history, the pair of strategies, restricted to the continuation game after this history occurs, is an ε-equilibrium.

Since the game is a game of perfect information (that is, there are no simultaneous moves, and all past moves are observed), it follows from Mertens (1987) that for every $\varepsilon > 0$ an ε-equilibrium exists. The construction of Mertens (1987) uses threats of punishment, and therefore it does not yield a subgame-perfect ε-equilibrium.

As we show below, the game can be reduced to a stopping game. Recently Mashiach-Yakovi (2008) proved that every stopping game admits a subgame-perfect ε-equilibrium. However, the subgame-perfect ε-equilibrium that is constructed by Mashiach-Yakovi (2008) need not be Markovian.

3. Existence of a subgame perfect equilibrium

Let $W \subseteq [0,1]^2$ be the support of F_2, that is, the set of all possible types (x, y) of the candidates. When (x, y) is the type of a candidate, x (respectively y) is the *worth* of the candidate for player I (respectively II). The following example shows that in general a subgame perfect 0-equilibrium may not exist.

Example 1. Suppose F_2 is the uniform distribution on the diagonal of the unit square; that is, the interval whose extreme points are $(0, 0)$ and $(1, 1)$. We claim that the game $G(0, 0)$ has no subgame perfect 0-equilibrium (σ, τ). Indeed, assume to the contrary that such a pair exists. Since F_2 is symmetric, the continuation payoff to both players is the same, say $\alpha \in [0, 1]$.

Since the players have no vetoes, the behavior of player II after a candidate was accepted by player I does not affect the outcome. Therefore, when player II decides whether to accept a candidate or to reject her, he compares his continuation payoff α to the candidate's type. Hence τ is a weak threshold strategy with threshold α.

When player I decides whether to accept a candidate or reject her, and knowing τ, he will compare his continuation payoffs: if he accepts the candidate his payoff will be the worth of the candidate (to either player), whereas if he rejects the candidate his payoff will be the maximum between α and the candidate's worth. Hence σ must be a weak threshold strategy with threshold α as well.

If both players use a weak threshold strategy with threshold α, and if $\alpha < 1$, then the expected payoff is $\frac{1+\alpha}{2}$. Since the unique solution of $\frac{1+\alpha}{2} = \alpha$ is $\alpha = 1$, there cannot be a SPE in weak threshold strategies with threshold lower than 1. If both players use a weak threshold strategy with threshold 1, since F_2 is the uniform distribution over the diagonal of the unit square, with probability 1 no candidate will be hired, and the expected payoff is 0. This shows that the game indeed does not have any subgame-perfect 0-equilibrium.

Observe that in this example there is a subgame perfect ε-equilibrium for every $\varepsilon > 0$: player I (respectively player II) accepts a candidate (x, y) if and only if $x \geqslant 1 - \varepsilon$ (respectively $y \geqslant 1 - \varepsilon$).

Remark 2. The distribution F_2 in Example 1 is not absolutely continuous with respect to the Lebesgue measure. One can obtain the same result by "fattening" the support of F_2; for example, by letting F_2 be the uniform distribution over the polytope whose extreme points are $(0, 0), (0, c), (c, 0), (1, 1)$, where $0 < c < \frac{1}{2}$. The calculations are tedious and not inspiring and therefore omitted.

In the present section we prove the following two results.

Theorem 3. *If F_2 has finite support, the game $G(i, j)$ admits a subgame-perfect* 0*-equilibrium in Markovian strategies.*

Theorem 4. *For every distribution F_2, the game $G(i, j)$ admits a subgame-perfect ε-equilibrium in Markovian strategies, for every $\varepsilon > 0$.*

Before proving the two theorems, we show that there may be equilibria which are not in Markovian strategies, and equilibria which need not be in threshold strategies.

Example 5. Let F_2 be the uniform distribution over a set of four atoms: $(\frac{1}{2}, \frac{1}{2})$, $(\frac{1}{2}, 0)$, $(0, \frac{1}{2})$, and $(1, 1)$. Consider the following strategy σ_p, where $p \in [0, 1]$: accept a candidate if her type is $(1, 1)$ or $(\frac{1}{2}, 0)$; accept her with probability p if her type is $(\frac{1}{2}, \frac{1}{2})$. Let τ_q be the symmetric strategy for player 2. For every $p, q \in [0, 1]$ the strategy pair (σ_p, τ_q) is a subgame perfect 0-equilibrium with expected payoff $(\frac{1}{2}, \frac{1}{2})$. However, only for $p = q = 1$ these strategies can be called threshold strategies. Now, suppose that at every stage, if no candidate was chosen so far, the players choose p and q according to some (non-stationary) rule, and use σ_p and τ_q at that stage. Then the resulting strategy pair is a subgame perfect 0-equilibrium, but it is not Markovian.

Proof of Theorem 3. The proof is by induction over $i + j$. We assume that $G(i - 1, j)$ and $G(i, j - 1)$ admit subgame-perfect 0-equilibria in Markovian strategies with expected payoff $(u-, v+)$ and $(u+, v-)$ respectively, and we study the part of the game $G(i, j)$, until either a candidate is hired (and then the game terminates), or some player casts a veto (and then the game terminates, with a terminal payoff that is either $(u-, v+)$ or $(u+, v-)$, depending on the player who cast the veto). Denote this part of the game by $\Gamma(i, j)$. To accommodate the case in which a player has no vetoes, if $i = 0$ we set $u- = -1$, if $j = 0$ we set $v- = -1$, and if $i = j = 0$ we also set $u+ = v+ = -1$.

The game $\Gamma(i, j)$ can be described as follows: before a candidate is chosen, each player chooses, for each type (x, y), whether he will accept such a type (A), reject such a type and cast a veto if needed (RV), or reject such a candidate and not cast a veto (RN). Thus, we can describe the game as a repeated game with absorbing entries, where the set of pure strategies of each player is the set of functions from the finite support W of F_2 to {A, RV, RN}. The payoff, as a function of the type of the chosen candidate, is described in Fig. 2(left). An empty entry in this figure indicates that the game continues if the entry is chosen.

	A	RV	RN
A	(x, y)	$(u+, v-)$	(x, y)
RV	$(u-, v+)$		
RN	(x, y)		

	A	R
A	$a(x, y)$	$a(x, y)$
R	$b(x, y)$	

Fig. 2. The game $\Gamma(i, j)$ (left) and its simplified version (right).

The game $\Gamma(i, j)$ is terminated once at least one player accepts a candidate. In equilibrium, the decision whether to cast a veto is dictated solely by the comparison between the worth of the candidate to the player and the continuation payoff: it affects the terminal payoff but not the termination of the game.

Let $a(x, y)$ be the terminal payoff if player I accepts (x, y); it is either (x, y) or $(u+, v-)$, depending on which one player II prefers:

$$a(x, y) = \begin{cases} (x, y), & y \geqslant v-, \\ (u+, v-), & y < v- . \end{cases} \tag{1}$$

Let $b(x, y)$ be the terminal payoff if player I rejects a candidate (x, y) and player II accepts her; it is either (x, y) or $(u-, v+)$, depending on which one player I prefers:

$$b(x, y) = \begin{cases} (x, y), & x \geqslant u-, \\ (u-, v+), & x < u- . \end{cases} \tag{2}$$

For the search of an equilibrium, we can use the functions a and b to simplify the game $\Gamma(i, j)$ as follows. We assume that at each round, each of the players chooses, for each candidate, whether to accept her or reject her, and the payoff (as a function of the candidate's type) is given in Fig. 2(right).

Note that the only (possible) discontinuities of $a(x, y)$ and $b(x, y)$ are on the horizontal interval $I_x = \{(x, v-), 0 \leqslant x \leqslant 1\}$ and the vertical interval $I_y = \{(u-, y), 0 \leqslant y \leqslant 1\}$.

To prove that the game $G(i, j)$ admits a subgame perfect Markovian equilibrium, it is sufficient to prove that the game $\Gamma(i, j)$ admits a stationary equilibrium. Indeed, by concatenating a stationary equilibrium in $\Gamma(i, j)$ with subgame perfect equilibria in $G(i - 1, j)$ and in $G(i, j - 1)$, which exists by induction, we obtain a Markovian subgame perfect equilibrium in $G(i, j)$.

The game $\Gamma(i, j)$ is a recursive absorbing game (see, e.g., Flesch et al., 1996); that is, a stochastic game with one non-absorbing state, such that the payoff in non-absorbing states is 0. Flesch et al. (1996) proved that every two-player recursive absorbing game admits a stationary ε-equilibrium, yet, since our game has sequential moves, this result does not imply the existence of a SPE in $\Gamma(i, j)$. Our proof uses ideas from Vrieze and Thuijsman (1989) and Solan and Vieille (2002).

A mixed action for player I is a probability distribution over pure actions. By Kuhn's theorem it is equivalent to a function $\sigma : W \to [0, 1]$ that assigns to each type (x, y) the probability that it will be accepted, which is denoted by $\sigma(x, y) \in [0, 1]$.

Mixed actions τ for player II are defined analogously. Every mixed action is identified with the stationary strategy that plays it at every stage.

For every pair (σ, τ) of stationary strategies, the probability of per-stage absorption is

$$p(\sigma, \tau) := \sum_{(x,y)\in W} F_2(x, y)\big(\sigma(x, y) + \big(1 - \sigma(x, y)\big)\tau(x, y)\big).$$

The pair (σ, τ) is called *absorbing* if $p(\sigma, \tau) > 0$, and *non-absorbing* otherwise. Denote by $\gamma(\sigma, \tau)$ the expected payoff under the pair of strategies (σ, τ). Then

$$\gamma(\sigma, \tau) = \begin{cases} (0, 0), & (\sigma, \tau) \text{ is non-absorbing}, \\ \frac{\sum_{(x,y)\in W} F_2(x,y)(\sigma(x,y)a(x,y)+(1-\sigma(x,y))\tau(x,y)b(x,y))}{\sum_{(x,y)\in W} F_2(x,y)(\sigma(x,y)+(1-\sigma(x,y))\tau(x,y))}, & \text{otherwise}. \end{cases} \tag{3}$$

Since the payoff function γ is the ratio of two polynomials in σ and τ, and since the denominator is positive whenever (σ, τ) is absorbing, we deduce that γ is continuous at every absorbing pair (σ, τ).

Let $\mathcal{S}$ and $\mathcal{T}$ be the sets of all mixed actions for the two players respectively. Let $\mathcal{S}_\varepsilon$ be the set of all mixed actions σ of player I that satisfy the following condition:

- $\sigma(x, y) \in [\varepsilon, 1 - \varepsilon]$ for each $(x, y) \in W$: player I accepts every type with probability at least ε, and rejects every type with probability at least ε.

The condition that $\sigma(x, y) \geqslant \varepsilon$ ensures that whatever player II plays, the game is bound to be absorbed, and therefore the payoff is continuous over $\mathcal{S}_\varepsilon \times \mathcal{T}$; the condition that $\sigma(x, y) \leqslant 1 - \varepsilon$ ensures that each candidate is not hired with positive probability, and so every decision point of player II is reached with positive probability: a best response of player II to a strategy in $\mathcal{S}_\varepsilon$ cannot use punishments that hurt player II.

Observe that $\mathcal{S}_\varepsilon$ is convex, compact, and non-empty (assuming $\varepsilon < \frac{1}{2}$). Let $\Gamma_\varepsilon(i, j)$ be the recursive absorbing game $\Gamma(i, j)$, in which the mixed-action space of player I is $\mathcal{S}_\varepsilon$, and the mixed-action space of player II is $\mathcal{T}$ (only player I is restricted).

For every $(\sigma, \tau) \in \mathcal{S}_\varepsilon \times \mathcal{T}$, the game is bound to be absorbed. Since $\mathcal{S}_\varepsilon$ and $\mathcal{T}$ are convex and compact, the standard argument of Shapley (1953) implies that the game $\Gamma_\varepsilon(i, j)$ has a 0-equilibrium in *stationary* strategies.

Lemma 6. *Let $(\sigma_\varepsilon, \tau_\varepsilon)$ be a stationary* 0*-equilibrium in $\Gamma_\varepsilon(i, j)$. For every type (x, y),*

1. *$\tau_\varepsilon(x, y)$ is a best reply (for player II) given the two possible continuation payoffs ($b(x, y)$ and $\gamma(\sigma_\varepsilon, \tau_\varepsilon)$): it chooses an action that attains the maximum* $\max\{b^{\mathrm{II}}(x, y), \gamma^{\mathrm{II}}(\sigma_\varepsilon, \tau_\varepsilon)\}$.
2. *$\sigma_\varepsilon(x, y)$ is the best reply (for player I, in the range $[\varepsilon, 1 - \varepsilon]$) between the two continuation payoffs $a(x, y)$ and $\tau_\varepsilon(x, y)b(x, y) + (1 - \tau_\varepsilon(x, y))\gamma(\sigma_\varepsilon, \tau_\varepsilon)$.*

Proof. We only prove the first statement. An analogous argument proves the second statement. Since $\sigma_\varepsilon(x, y) \leqslant 1 - \varepsilon$, the type (x, y) has a positive probability to be realized and accepted by player I. If τ_ε does not choose the higher continuation payoff, player II could gain by deviating. □

By compactness, there is a sequence $(\varepsilon_k)_{k\in\mathbf{N}}$ that converges to 0 such that (a) the sequence $(\sigma_{\varepsilon_k}, \tau_{\varepsilon_k})_{k\in\mathbf{N}}$ converges to a limit (σ_*, τ_*), (b) the set of types (x, y) for which $\tau_{\varepsilon_k}(x, y)$ is positive is independent of k, and (c) the corresponding payoffs $(\gamma(\sigma_{\varepsilon_k}, \tau_{\varepsilon_k}))_{k\in\mathbf{N}}$ converge to a limit γ_*. Eq. (3) implies that γ_* is a weighted average of $(a(x, y), b(x, y))_{(x,y)\in W}$:

$$\gamma_* = \sum_{(x,y)\in W} \big(\alpha(x, y)a(x, y) + \beta(x, y)b(x, y)\big), \tag{4}$$

where, for every $(x, y) \in W$,

$$\alpha(x, y) = \lim_{k\to\infty} \frac{F_2(x, y)\sigma_{\varepsilon_k}(x, y)}{\sum_{(x',y')\in W} F_2(x', y')(\sigma_{\varepsilon_k}(x', y') + (1 - \sigma_{\varepsilon_k}(x', y'))\tau_{\varepsilon_k}(x', y'))},$$

$$\beta(x, y) = \lim_{k\to\infty} \frac{F_2(x, y)(1 - \sigma_{\varepsilon_k}(x, y))\tau_{\varepsilon_k}(x, y)}{\sum_{(x',y')\in W} F_2(x', y')(\sigma_{\varepsilon_k}(x', y') + (1 - \sigma_{\varepsilon_k}(x', y'))\tau_{\varepsilon_k}(x', y'))},$$

and these limits exist without loss of generality.

We now prove that the game $\Gamma(i, j)$ admits a stationary equilibrium. We first deal with the case that (σ_*, τ_*) is absorbing.

Lemma 7. *If (σ_*, τ_*) is absorbing, then it is a stationary equilibrium in $\Gamma(i, j)$.*

Proof. We only prove that player I cannot profit by deviating; the proof that player II cannot profit by deviating is analogous.

Let σ be an arbitrary strategy of player I in $\Gamma(i, j)$. Assume first that according to τ_* there is a type that is accepted with positive probability. In particular, (σ, τ_*) is absorbing. We approximate σ by strategies in $\mathcal{S}_{\varepsilon_k}$. Define for every $k \in \mathbf{N}$ a strategy $\widehat{\sigma}_k$ as follows:

$$\widehat{\sigma}_k(x, y) = \begin{cases} \varepsilon_k, & \sigma(x, y) < \varepsilon_k, \\ \sigma(x, y), & \sigma(x, y) \in [\varepsilon_k, 1 - \varepsilon_k], \\ 1 - \varepsilon_k, & \sigma(x, y) \geqslant 1 - \varepsilon_k. \end{cases}$$

Note that $\sigma = \lim_{k\to\infty} \widehat{\sigma}_k$. By the continuity of γ at absorbing strategy pairs we deduce that

$$\gamma^{\mathrm{I}}(\sigma, \tau_*) = \lim_{k\to\infty} \gamma^{\mathrm{I}}(\widehat{\sigma}_k, \tau_{\varepsilon_k}). \tag{5}$$

From Eq. (5), since $(\sigma_{\varepsilon_k}, \tau_{\varepsilon_k})$ is an equilibrium in $\Gamma_\varepsilon(i, j)$ for every ε, and since (σ_*, τ_*) is absorbing, we obtain:

$$\gamma^{\mathrm{I}}(\sigma, \tau_*) = \lim_{k\to\infty} \gamma^{\mathrm{I}}(\widehat{\sigma}_k, \tau_{\varepsilon_k}) \leqslant \lim_{k\to\infty} \gamma^{\mathrm{I}}(\sigma_{\varepsilon_k}, \tau_{\varepsilon_k}) = \gamma_*^{\mathrm{I}} = \gamma^{\mathrm{I}}(\sigma_*, \tau_*), \tag{6}$$

so that player I does not profit by deviating to σ.

Suppose now that according to τ_* all types are rejected. If (σ, τ_*) is absorbing, then once again by continuity arguments it follows that $\gamma^1(\sigma, \tau_*) \leqslant \gamma_*^{\mathrm{I}} = \gamma(\sigma_*, \tau_*)$. If, on the other hand, (σ, τ_*) is non-absorbing, then since payoffs are non-negative we obtain:

$$\gamma^{\mathrm{I}}(\sigma, \tau_*) = 0 \leqslant \gamma_*^{\mathrm{I}} = \gamma(\sigma_*, \tau_*),$$

so that in this case as well player I does not profit by deviating to σ. $\square$

We now deal with the case that (σ_*, τ_*) is not absorbing.

Lemma 8. *If (σ_*, τ_*) is non-absorbing, then there is a type (x, y) and a stationary equilibrium in $\Gamma(i, j)$, such that on the equilibrium path both players accept only (x, y).*

Proof. We first show that for every type $(x, y) \in W$ we have

$$a^{\mathrm{I}}(x, y) \leqslant \gamma_*^{\mathrm{I}}. \tag{7}$$

Indeed, since player I accepts the type (x, y) with a probability that goes to 0, by Lemma 6 we have for k large enough:

$$a^{\mathrm{I}}(x, y) \leqslant \tau_{\varepsilon_k}(x, y) b^{\mathrm{I}}(x, y) + \bigl(1 - \tau_{\varepsilon_k}(x, y)\bigr) \gamma^{\mathrm{I}}(\sigma_{\varepsilon_k}, \tau_{\varepsilon_k}). \tag{8}$$

Since (σ_*, τ_*) is non-absorbing, τ_* rejects all candidates, so that $\lim_{k\to\infty} \tau_{\varepsilon_k}(x, y) = 0$ for every type (x, y). Taking the limit as k goes to infinity in (8), we obtain

$$a^{\mathrm{I}}(x, y) \leqslant \lim_{k\to\infty} \gamma^{\mathrm{I}}(\sigma_{\varepsilon_k}, \tau_{\varepsilon_k}) = \gamma_*^{\mathrm{I}}, \tag{9}$$

and therefore (7) holds.

We now argue that for every type $(x, y) \in W$ we have

$$b^{\mathrm{II}}(x, y) \leqslant \gamma_*^{\mathrm{II}}, \tag{10}$$

with an equality whenever $\tau_{\varepsilon_k}(x, y) > 0$ for every k.[1] Since (σ_*, τ_*) is non-absorbing, σ_* rejects all candidates, so that $\lim_{k\to\infty} \sigma_{\varepsilon_k}(x, y) = 0$ for every type (x, y). By Lemma 6 we obtain:

$$b^{\mathrm{II}}(x, y) \leqslant \gamma^{\mathrm{II}}(\sigma_{\varepsilon_k}, \tau_{\varepsilon_k}). \tag{11}$$

By taking the limit as k goes to infinity we obtain (10).

Suppose that there exists a type $(x_0, y_0) \in W$ such that $b^{\mathrm{I}}(x_0, y_0) \geqslant \gamma_*^{\mathrm{I}}$ and $b^{\mathrm{II}}(x_0, y_0) = \gamma_*^{\mathrm{II}}$. Consider the stationary strategy pair $(\widehat{\sigma}, \widehat{\tau})$ in $\Gamma(i, j)$ in which player I accepts the type (x_0, y_0) (and any type (x, y) with $x > x_0$, which may arrive with probability 0) and rejects any other type, and player II accepts any type (x, y) such that $y = y_0$ (and any type (x, y) with $y > y_0$) and rejects any other type. It follows that $\gamma(\widehat{\sigma}, \widehat{\tau}) = (x, y)$, and using (7) and (10) one can verify that $(\widehat{\sigma}, \widehat{\tau})$ is a stationary equilibrium in $\Gamma(i, j)$.

If there is a type $(x_0, y_0) \in W$ such that $a^{\mathrm{I}}(x_0, y_0) = \gamma_*^{\mathrm{I}}$ and $a^{\mathrm{II}}(x_0, y_0) \geqslant \gamma_*^{\mathrm{II}}$, then the following stationary strategy pair is a stationary equilibrium in $\Gamma(i, j)$: player I accepts the type (x_0, y_0) (and any type (x, y) with $x > x_0$, which may arrive with probability 0) and rejects any other type, and player II accepts any type (x, y) such that $y = y_0$ (and any type (x, y) with $y > y_0$) and rejects any other type.

[1] Recall that by the definition of $(\varepsilon_k)_{k\in\mathbf{N}}$, the set of types (x, y) for which $\tau_{\varepsilon_k}(x, y) > 0$ is independent of k.

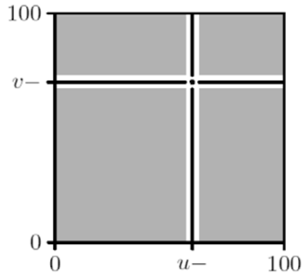

Fig. 3. The partition of $[0,1]^2$ into nine parts.

To conclude the proof, we suppose now that the conditions given in the previous two paragraphs do not hold, and we will derive a contradiction. Since the conditions given in the previous two paragraphs do not hold, for every type (x,y) at least one of the following conditions hold:

C1. $b^{\mathrm{I}}(x,y) < \gamma_*^{\mathrm{I}}$ or $b^{\mathrm{II}}(x,y) < \gamma_*^{\mathrm{II}}$;
C2. $a^{\mathrm{I}}(x,y) < \gamma_*^{\mathrm{I}}$ or $a^{\mathrm{II}}(x,y) < \gamma_*^{\mathrm{II}}$.

By Lemma 6, if $b^{\mathrm{II}}(x,y) < \gamma_*^{\mathrm{II}}$ then $\tau_{\varepsilon_k}(x,y) = 0$ for every k sufficiently large, whereas if $a^{\mathrm{I}}(x,y) < \gamma_*^{\mathrm{I}}$ then $\sigma_{\varepsilon_k}(x,y) = 0$ for every k sufficiently large. In other words, if $b^{\mathrm{II}}(x,y) < \gamma_*^{\mathrm{II}}$ then under τ_* the candidate (x,y) is not accepted by player II, so that $\beta(x,y) = 0$. Similarly, if $a^{\mathrm{I}}(x,y) < \gamma_*^{\mathrm{I}}$ then $\alpha(x,y) = 0$.

Finally, by Eq. (4), γ_* is a weighted average of $\{a(x,y), b(x,y), (x,y) \in W\}$ with weights $\{\alpha(x,y), \beta(x,y), (x,y) \in W\}$. The discussion above implies that if $\alpha(x,y) > 0$ then $a^{\mathrm{I}}(x,y) = \gamma_*^{\mathrm{I}}$ and $a^{\mathrm{II}}(x,y) < \gamma_*^{\mathrm{II}}$, while if $\beta(x,y) > 0$ then $b^{\mathrm{I}}(x,y) < \gamma_*^{\mathrm{I}}$ and $b^{\mathrm{II}}(x,y) = \gamma_*^{\mathrm{II}}$. However, any weighted average that satisfies these condition must yield to at least one of the players a payoff which is strictly lower than his payoff in γ_*, a contradiction. This contradiction implies that there exists a stationary equilibrium in which a single type is accepted by both players, as desired. □

So Theorem 3 is proved. □

We now turn to prove Theorem 4.

Proof of Theorem 4. Along the proof we use the same notations as in the proof of Theorem 3. We prove the theorem by induction on $i+j$. We assume that $G(i-1,j)$ and $G(i,j-1)$ admit a subgame perfect $\frac{\varepsilon}{6}$-equilibrium in Markovian strategies for every $\varepsilon > 0$, and we denote by $(u-, v+)$ and $(u+, v-)$ corresponding $\frac{\varepsilon}{6}$-equilibrium payoffs in Markovian strategies in these games. As in the proof of Theorem 3, if $i = 0$ we set $u- = -1$, if $j = 0$ we set $v- = -1$, and if $i = j = 0$ we also set $u+ = v+ = -1$.

Consider now the game $\Gamma(i,j)$ that was defined in the proof of Theorem 3. The idea of the proof is to approximate the game $\Gamma(i,j)$ (with a distribution F_2 that need not be finite), by a game with a distribution $\widehat{F}_2$ that has finite support, to argue that by Theorem 3 the approximating game admits a subgame perfect 0-equilibrium, and finally to prove that every subgame perfect 0-equilibrium of the approximating game is a subgame perfect ε-equilibrium of $\Gamma(i,j)$.

We now turn to the formal proof. We start by defining an $\frac{\varepsilon}{6}$-discretization $\widehat{F}_2$ of F_2. The construction is a little involved since we need to preserve continuity of payoffs, while the functions a and b have discontinuities. Split $[0,1]^2$ into the following nine parts (see Fig. 3):

$$\big\{(x,y), u- < x \leqslant 1, v- < y \leqslant 1\big\}, \quad \big\{(x,y), u- < x \leqslant 1, 0 \leqslant y < v-\big\}, \quad \big\{(x,y), 0 \leqslant x < u-, v- < y \leqslant 1\big\},$$
$$\big\{(x,y), 0 \leqslant x < u-, 0 \leqslant y < v-\big\}, \quad \big\{(x,v-), 0 \leqslant x < u-\big\}, \quad \big\{(x,v-), u- < x \leqslant 1\big\},$$
$$\big\{(u-,y), 0 \leqslant y < v-\big\}, \quad \big\{(u-,y), v- < y \leqslant 1\big\} \quad \text{and} \quad (u-,v-).$$

We partition each part (except the last part which is a single point) into finitely many sets such that each of the overall L resulting sets $A_1, A_2, \ldots, A_L$, has diameter at most $\frac{\varepsilon}{6}$:

$$\sup_{(x,y),(x',y') \in A_l} d\big((x,y),(x',y')\big) \leqslant \frac{\varepsilon}{6}, \quad \forall l,$$

and for each l, $1 \leqslant l \leqslant L$, there is $(x_l, y_l) \in A_l$ such that $\widehat{F}_2(x_l, y_l) = F_2(A_l)$. (That is, let $\widehat{F}_2$ be the atomic distribution that is defined as follows: for each $l = 1, 2, \ldots, L$, choose $(x_l, y_l) \in A_l$ and set $\widehat{F}_2(x_l, y_l) = F_2(A_l)$.) Observe that the functions a and b are continuous over each A_l.

We are going to consider three games, G_1, G_2, and G_3. These are recursive games, in which at every stage a candidate is chosen according to some distribution, then player I, upon observing the candidate's type, decides whether to accept the candidate or reject her, and then player II, after observing the candidate's type and player I's decision, decides whether to accept her or reject her. The game continues once at least one player accepts a candidate (there is no opportunity to cast vetoes), and the games only differ in the underlying distribution according to which candidates are chosen and by their absorbing payoffs.

- In the game G_1 the candidates' distribution is F_2 and the absorbing payoff is $a(x,y)$ (respectively $b(x,y)$) if (x,y) is accepted by player I (respectively player II). Thus, G_1 is equal to $\Gamma(i,j)$.
- In the game G_2 the candidates' distribution is F_2. If $(x,y)\in A_l$ is accepted by player I (respectively player II), the absorbing payoffs is $a(x_l,y_l)$ (respectively $b(x_l,y_l)$).
- In the game G_3 the candidates' distribution is $\widehat{F}_2$ and the absorbing payoffs is $a(x_l,y_l)$ (respectively $b(x_l,y_l)$) if $(x,y)\in A_l$ is accepted by player I (respectively player II).

As mentioned in the proof of Theorem 3, a $\frac{5\varepsilon}{6}$-equilibrium in the game G_1 can be turned into an ε-equilibrium in $G(i,j)$, by supporting it with $\frac{\varepsilon}{6}$-equilibria in $G(i-1,j)$ and $G(i,j-1)$ that yield payoff $(u-,v+)$ and $(u+,v-)$ respectively if one of the players casts a veto. Denote by $\gamma_1(\sigma,\tau)$, $\gamma_2(\sigma,\tau)$ and $\gamma_3(\sigma,\tau)$ the payoff functions in these three games. Since the absorbing payoffs in G_1 and in G_2 differ by at most $\frac{\varepsilon}{6}$, we have

$$\left\|\gamma_1(\sigma,\tau)-\gamma_2(\sigma,\tau)\right\|\leqslant\frac{\varepsilon}{6},\quad\forall\sigma,\forall\tau.\tag{12}$$

Fix a strategy pair (σ,τ). Let σ_3 be the strategy that is defined as follows:

$$\sigma_3(x,y)=\frac{\int_{A_l}\sigma(x',y')\,dF_2(x',y')}{F_2(A_l)}\quad\text{if }(x,y)\in A_l.\tag{13}$$

This is the average probability to accept (x,y) under σ, provided it is in A_l. We define this strategy only for $(x,y)\in A_l$ such that $F_2(A_l)>0$. Let τ_3 be the analogous strategy for player II:

$$\tau_3(x,y)=\frac{\int_{A_l}(1-\sigma(x',y'))\tau(x',y')\,dF_2(x',y')}{\int_{A_l}(1-\sigma(x',y'))\,dF_2(x',y')}\quad\text{if }(x,y)\in A_l.\tag{14}$$

This is the probability that player II accepts (x,y) if player I rejected it, conditional on $(x,y)\in A_l$, and it is defined only for $(x,y)\in A_l$ such that the denominator is not zero. If the denominator is zero, then I will aways accept a candidate it A_l, so the definition of τ_3 does not matter. Suppose also that under σ_3 player I vetoes a candidate accepted by player II with the probability which is the conditional probability he would have vetoed her under σ (conditional on A_l), and similarly for player II. Since the diameter of each A_l is at most $\frac{\varepsilon}{6}$,

$$\left|\gamma_2(\sigma,\tau)-\gamma_3(\sigma_3,\tau_3)\right|\leqslant\frac{\varepsilon}{6},\quad\forall\sigma,\forall\tau.\tag{15}$$

This holds only if τ or σ (or both) depend on (x,y) only through the atom that contains it. Let now (σ_3,τ_3) be a subgame-perfect 0-equilibrium in G_3. Since $\widehat{F}_2$ has a finite support, as we exhibited in the proof of Theorem 3 such an equilibrium exists. Let σ^* and τ^* be the strategies that are defined as follows: as long as no player casts a veto, σ^* follow σ_3 and τ^* follows τ_3. Once player I (respectively player II) casts a veto, (σ^*,τ^*) coincides with a subgame-perfect $\frac{\varepsilon}{6}$-equilibrium in Markovian strategies in $G(i-1,j)$ (respectively in $G(i,j-1)$) with corresponding payoff $(u-,v+)$ (respectively $(u+,v-)$).

The strategies σ^* and τ^* are Markovian. We now argue that (σ^*,τ^*) is a subgame-perfect ε-equilibrium in $G(i,j)$. We only prove that player I cannot gain more than ε by deviating. Fixing τ^*, player I's maximization problem reduces to a Markov decision problem, and therefore he has an optimal strategy σ which is Markovian. Let σ_1 be the strategy that σ induces on G_1, and let σ_3 be the corresponding discretized strategy defined above. Let σ_1^* and τ_1^* be the strategies that σ^* and τ^* induce in G_1, that is, $\sigma_1^*(x,y)=\sigma_1^*(x_1,y_1)$ whenever $(x,y)\in A_l$. Let τ_3^* be the strategy defined in (14) with respect to τ^*. Since σ^* plays an $\frac{\varepsilon}{6}$-equilibrium once one of the players casts a veto, by (12), (15), and since (σ_3,τ_3) is a subgame perfect 0-equilibrium in G_3:

$$\gamma^{\mathrm{I}}(\sigma,\tau^*)\leqslant\gamma_1^{\mathrm{I}}(\sigma_1,\tau_1^*)+\frac{\varepsilon}{6}\leqslant\gamma_2^{\mathrm{I}}(\sigma_1,\tau_1^*)+\frac{2\varepsilon}{6}\tag{16}$$

$$\leqslant\gamma_3^{\mathrm{I}}(\sigma_1,\tau_1^*)+\frac{3\varepsilon}{6}\leqslant\gamma_3^{\mathrm{I}}(\sigma_1^*,\tau_1^*)+\frac{3\varepsilon}{6}\tag{17}$$

$$\leqslant\gamma_2^{\mathrm{I}}(\sigma_1^*,\tau_1^*)+\frac{4\varepsilon}{6}\leqslant\gamma_1^{\mathrm{I}}(\sigma_1^*,\tau_1^*)+\frac{5\varepsilon}{6}\tag{18}$$

$$\leqslant\gamma^{\mathrm{I}}(\sigma^*,\tau^*)+\varepsilon,\tag{19}$$

so that indeed player I cannot gain more than ε by deviating. □

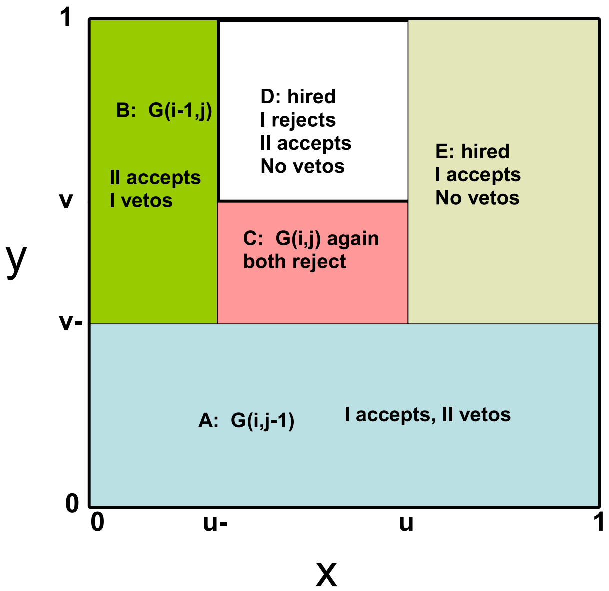

Fig. 4. Progressive equilibrium for $G(i, j)$.

4. Behavior in equilibrium

In general, for every continuation payoffs $(u+, v-)$ and $(u-, v+)$ there may be many equilibria in $G(i, j)$. In this section we single out two types of equilibria, progressive equilibria and helpful equilibria. These equilibria represent the typical behavior of the players (see Remark 11 in this section). They also help us explain why one of the players may prefer that the other player would become more powerful in some situations, such as the donation paradox that we present in Section 5.1 below.

Definition 9. Let (σ, τ) be a Markovian subgame perfect equilibrium in $G(i, j)$ with continuation payoff $(u+, v-)$ (respectively $(u-, v+)$) in $G(i, j-1)$ (respectively in $G(i-1, j)$). (σ, τ) is called *progressive* if $u \leqslant u+$, $v \leqslant v+$, and (in case of equality) players choose to play the game in which the other player loses a veto rather than repeat the game $G(i, j)$.

We emphasize that whether or not an equilibrium is progressive depends on the continuation payoffs, and therefore on the continuation strategies in $G(i-1, j)$ and in $G(i, j-1)$.

In a progressive equilibrium a player does not profit from the fact that the other player has one extra veto, and therefore he tries to make the other player use one veto.

Given the values $u-, u, v-, v$ it is easy to derive the SPE strategies as a function of the candidate values x and y, depending on which of the five partitioning rectangles A, B, C, D, E in Fig. 4 the point (x, y) belongs to.[2] (To simplify the discussion of ties, we assume these rectangles contain their left and bottom sides.)

Consider Fig. 4, and suppose $y < v-$, that is, $(x, y) \in A$. Such a candidate will never be hired, because player II will cast a veto if necessary. The best player I can do is force player II to cast a veto. He does this by accepting the candidate. Similarly, if $(x, y) \in B$, then player I will cast a veto if necessary. The best player II can do is to accept the candidate, thereby forcing player I to cast a veto. If $(x, y) \in C$, neither player would cast a veto, and both prefer rejecting the candidate and replaying the game. If $(x, y) \in D \cup E$, the candidate will be hired, as at least one player prefers this to replaying $G(i, j)$, and neither player is willing to cast a veto. Denote by (u, v) the equilibrium payoff in $G(i, j)$. Then

$$(u, v) = T(u, v),$$

$$\text{where } T(u, v) \equiv F_2(A) \cdot (u+, v-) + F_2(B) \cdot (u-, v+) + F_2(C) \cdot (u, v) + \int_{D \cup E} \big(x(z), y(z)\big)\, dF_2(z).$$

Here, for every subset $H \subseteq [0, 1]^2$, $F_2(H)$ is the probability that $(x, y) \in H$. If F_2 is the uniform distribution, then Lemma 16 of Alpern and Gal (2008) shows that, for any game $G(i, j)$, T is a contraction map of the unit square, which implies the existence of a unique progressive equilibrium (provided $u+$, $u-$, $v+$ and $v-$ are given).

[2] Recall that the inequalities $u- \leqslant u$ and $v- \leqslant v$ were established by Alpern and Gal (2008). Also note that the placement of the candidate (x, y) with respect to $(u+, v+)$ does not affect decisions.

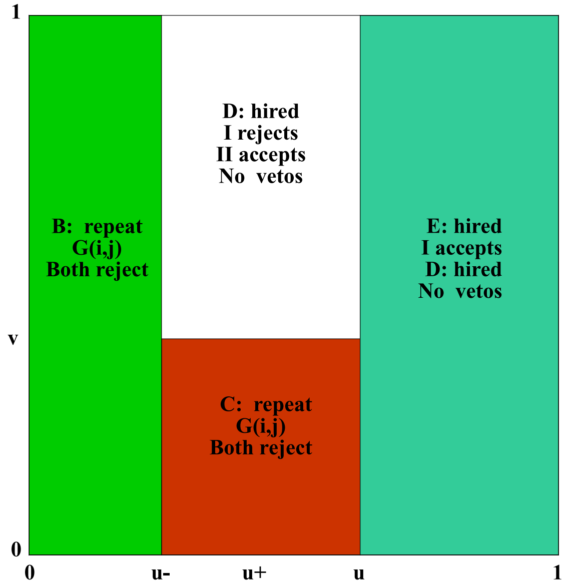

Fig. 5. A helpful equilibrium for $G(1,0)$.

Note that if $u < u+$ and $v < v+$, then the equilibrium actions played in the interior of each of the sets A, B, C, D and E in Fig. 4 are strictly dominant.

In order to illustrate the other types of equilibrium we first analyze the game $G(1,0)$ for a symmetric F_2 (this includes the case in which x and y are i.i.d.). The game $G(0,0)$ has been analyzed in Alpern and Gal (2008). They proved that if F_2 is symmetric, then there exists a symmetric equilibrium. Denote the expected payoff to each player in $G(0,0)$ by u_0. We now consider the possibility of an equilibrium for $G(1,0)$ with equilibrium values (u,v) with $u \geqslant u_0$ and $v \geqslant u_0$. Previously we had assumed an extra veto for player I would hurt player II, with $v < u_0$. If $v > u_0$, player II would rather replay $G(1,0)$ and get v than go to $G(0,0)$ and get u_0. So player I will never use his veto, as in such a case player II would previously have rejected the candidate. We call such an equilibrium *helpful*. (To simplify the discussion of ties, we assume the rectangles in Fig. 5 contain their left and bottom sides.)

We now consider the general situation.

Definition 10. Let (σ,τ) be a Markovian subgame perfect equilibrium in $G(i,j)$ with continuation payoff $(u+, v-)$ (respectively $(u-, v+)$) in $G(i, j-1)$ (respectively in $G(i-1, j)$). (σ,τ) is called *helpful* if $u > u+$ and $v > v+$.

As in the case of progressive equilibrium, the property of being helpful depends on the continuation payoffs in $G(i-1,j)$ and in $G(i,j-1)$, and therefore on the continuation strategies in these games.

An equilibrium is helpful if each player loses when the other player casts a veto. In such a case, in equilibrium a player will not accept a candidate that the other player does not like, so that the other player will not be forced to use his veto power.

A *helpful* equilibrium has the profile depicted in Fig. 6. (To simplify the discussion of ties, we assume these rectangles contain their left and bottom sides. Also, player II accepts a candidate in E if $y \geqslant v$ and rejects her if $y < v$.)

Indeed, consider a candidate (x,y). If $y < v-$ player II will cast a veto. Since $u < u+$ player I prefers that player II would not cast a veto, hence he will reject such a candidate (region A in Fig. 6). For an analog reason both players reject the candidate if $(x,y) \in B$. The behavior of the players in regions C, D and E is as in a progressive equilibrium. Thus, the payoff (u,v) that corresponds to an helpful equilibrium is a fixed point

$$(u,v) = T(u,v),$$

where

$$T(u,v) = F_2(A \cup B \cup C) \cdot (u,v) + \int_{D \cup E} \big(x(z), y(z)\big)\, dF_2(z).$$

Note that if $u > u+$ and $v > v+$, then the actions in each region in Fig. 6 are strictly dominant.

There can also be a mixed equilibrium, in which $u < u_+$ but $v > v_+$ or $u > u_+$ but $v < v_+$.

Remark 11. Note that if a 'degeneracy,' $u = u+$ or $v = v+$, does not occur, then, for any game $G(i,j)$, $i, j \geqslant 0$, with an absolutely continuous F_2, there exist only pure equilibria (possibly more than one) which are either progressive or helpful

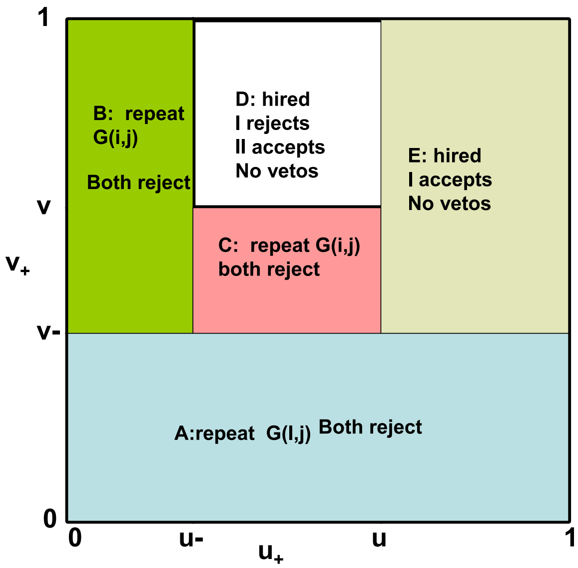

Fig. 6. A helpful equilibrium for $G(i, j)$.

or mixed. For example, assume that we found an equilibrium payoff (u, v) (either by a numerical scheme based on Section 3, or by some theoretical analysis), and it turns out that $u < u+$ and $v < v+$. Then the only possible equilibrium strategies, are given by the progressive equilibrium in Fig. 4, because the actions in this figure are strictly dominant. Similarly, if $u > u+$ and $v > v+$, then the only possible equilibria are given by helpful equilibria as in Fig. 6, because the actions in this figure are strictly dominant. The mixed cases can be handled in a similar way.

5. Unexpected behavior of equilibria

If F_2 is symmetric, e.g., if x and y are independent with the same distribution F, then the game $G(0,0)$ always has a symmetric SPE in threshold strategies (u_0, u_0), where $0 < u_0 < 1$ is the threshold of acceptance: a candidate is accepted by a player if and only if her worth for that player is at least u_0 (see Alpern and Gal, 2008). Simple as it looks, the game $G(0,0)$ displays some unexpected behavior. We first show that the number of equilibria is not bounded.

Example 12. Let n be a natural number. We construct a game $G(0,0)$ with n threshold equilibria. Set

$$z_0 = 0, \qquad z_i = 1/(2+\varepsilon)^{n-i}, \quad i = 1, 2 \ldots, n. \tag{20}$$

Let x and y be i.i.d. random variables with a discrete distribution: their support is $\{z_0, z_1, \ldots, z_n\}$. The probabilities $p(z_i) = P(x = z_i) = P(y = z_i)$ satisfy

$$p(z_i) \ll p(z_{i-1}), \quad i = 1, 2 \ldots, n. \tag{21}$$

That is, the first positive score is almost certainly z_1, if not, then almost certainly z_2, etc. We show that all positive z_i, $i = 1, 2, \ldots, n$ can be threshold of symmetric equilibria.

Let $1 \leqslant i \leqslant n$. If one player, say player I, uses a threshold strategy with threshold z_i, and accepts a candidate if and only if his worth is at least z_i, then the payoff of each player is close to $z_i/2$, and player II's best response is to use a threshold strategy with threshold z_i as well. Indeed, if player II uses a higher threshold, then player I will almost surely stop before player II, and player II's payoff will be close to 0. If player II uses a lower threshold, then he would stop before player I, and his payoff would be close to z_{i-1}, which is smaller than $z_i/2$.

It seems natural that the outcome is monotonic in the distribution of the population: if, say, one population is better than a second population, that is, if the distribution of the first population stochastically dominates the distribution of the second population, then the equilibrium payoff will be higher in the first population. The next example shows that this is not the case. This phenomenon corresponds to the well-known non-monotonicity of the equilibrium correspondence in one-shot games.

Example 13. Let x, y be i.i.d. r.v.s. with the following three atom distribution, where ε is a small positive real number:

$$z = \begin{cases} 1 & \text{with probability } \varepsilon, \\ d & \text{with probability } \varepsilon, \text{ where } d \in (0,1) \text{ and} \\ 0 & \text{with probability } 1 - 2\varepsilon. \end{cases}$$

We consider only SPE in threshold strategies, which always exists for $G(0,0)$ with a symmetric F_2, as mentioned above. Two thresholds may produce a SPE, either 1 or d.

For threshold 1 to be a Nash equilibrium (NE) it is necessary and sufficient that $\frac{1}{2} \geqslant \frac{1+d}{3} + O(\varepsilon)$ and for this NE to be subgame perfect (SP) it is necessary and sufficient that $d \leqslant \frac{1}{2} + O(\varepsilon)$ so threshold 1 is SPE if and only if $0 < d \leqslant \frac{1}{2} + O(\varepsilon)$.

For threshold d to be a NE it is necessary and sufficient that $\frac{1+d}{4} + O(\varepsilon) \geqslant \frac{1}{3}$ and for this NE to be SP it is necessary and sufficient that $d \geqslant \frac{1+d}{4} + O(\varepsilon)$ so threshold d is SPE if and only if $\frac{1}{3} + O(\varepsilon) \leqslant d < 1$.

Thus, for $0 < d < 1/3 + O(\varepsilon)$ threshold 1 is the unique SPE with payoff to each player $\frac{1}{2} + O(\varepsilon)$, while for $\frac{1}{2} + O(\varepsilon) < d < 1$ threshold d is the unique SPE with payoff to each player $\frac{1+d}{4} + O(\varepsilon) < \frac{1}{2}$. Since any population with $d > \frac{1}{2}$ stochastically dominates a population with $d < \frac{1}{3}$ we obtain a non-monotonicity of the equilibrium payoffs.

5.1. Donation paradox

The donation paradox was introduced by Fisher and Schotter (1978), who called it the "paradox of redistribution," in the context of power indices of simple voting games. This paradox occurs when a player gives an apparently valuable prerogative to another player, but 'does better,' according to some criterion. Specifically, when voting weights are reallocated, it may be observed that the voting power of some members, as measured by the Shapley-Shubik and Banzhaf power indices, increases while their voting weight decreases. (See also Felsenthal and Machover, 1995; Kadane et al., 1999; Holler and Napel, 2004.) DeGroot and Kadane (1980) present an example of selecting a jury of two participants, in which one of the players prefers to give his veto right to the other player, because such a move yields a better outcome for himself. However, Theorem 3 in DeGroot and Kadane (1980) implies that if the (two) players have to choose only one candidate, a donation paradox cannot occur: giving away one of your vetoes to the other player lowers the payoff of a player. It should be noted, though, that the game presented by DeGroot and Kadane (1980) has a finite length, due to the rule of always accepting a candidate that was not vetoed. This property simplifies the situation in their model. On the other hand, in our model, it is possible that both players reject a candidate and the next candidate is called for. This introduces the inherent complexity of infinitely repeated games into our model.

As we show now, a donation paradox may occur in our model. The number of vetoes each player has at the beginning of the game measures, in a sense, its relative strength. Indeed, Alpern and Gal (2008) proved that in a Markovian equilibrium, player I's (respectively player II's) expected payoff in $G(i,j)$ is always at least as much as his expected payoff after he casts a veto. We present an example in which player I *prefers* $G(0,1)$ to $G(1,0)$.[3]

Example 14. F_2 has 3 types of candidates, as follows:

Type	Probability
(b,b)	$1-2\varepsilon$
$(1,0)$	ε
$(1-\varepsilon, 1-\varepsilon)$	ε

where $\frac{1}{2} < b < 1$ and $\varepsilon \ll 1$.

We argue that player I prefers $G(0,1)$ to $G(1,0)$. The intuition is as follows. There are three types of candidates: (1) a type (b,b) which is mediocre for both players, (2) a type $(1-\varepsilon, 1-\varepsilon)$ which is excellent for both players, and (3) a type $(1,0)$ which is perfect for player I and bad for player II. The mediocre type arrives frequently, whereas the other two types arrive rarely.

If player II cannot prevent player I from accepting the perfect-to-I-bad-for-II candidate, then he has to accept a mediocre candidate. Therefore the SPE payoff in both $G(0,0)$ and $G(1,0)$ is close to the payoff from the mediocre type. In $G(0,1)$, on the other hand, player II can threaten player I, that if player I accepts the perfect-to-I-bad-for-II candidate type, then he will cast a veto, and the outcome will be mediocre to both players. Therefore, player I will not accept this type, and the players can end up with the excellent candidate. The formal calculations follow.

In $G(0,0)$ no player has a veto. In a subgame perfect equilibrium each player accepts his most preferred type: player I accept the type $(1,0)$, while player II accepts the type $(1-\varepsilon, 1-\varepsilon)$. Since the average of these two vectors gives player II a payoff $\frac{1-\varepsilon}{2} < b$, player II accepts the type (b,b) as well. Therefore the unique equilibrium payoff, $\gamma_{0,0}$, is:

$$\gamma_{0,0}^{\mathrm{I}} = (1-2\varepsilon)b + \varepsilon(1-\varepsilon) + \varepsilon = b + \varepsilon(2-2b-\varepsilon) > b, \tag{22}$$

$$\gamma_{0,0}^{\mathrm{II}} = (1-2\varepsilon)b + \varepsilon(1-\varepsilon) = b + \varepsilon(1-2b-\varepsilon) < b. \tag{23}$$

[3] Player II usually prefers $G(0,1)$ to $G(1,0)$. For example assume that there are just two types of candidates, $(1,0)$ and $(0,1)$, each with probability $\frac{1}{2}$. If no player has a veto, then each of them accepts his preferred type, so that the unique SPE payoff in $G(0,0)$ is $(\frac{1}{2}, \frac{1}{2})$. If one player has a single veto, while the other has no veto, then in a SPE each player will accept his preferred type, and the player with veto will veto once the candidate preferred by the other player. Therefore the unique SPE payoff in $G(1,0)$ (respectively in $G(0,1)$) is $(\frac{3}{4}, \frac{1}{4})$ (respectively $(\frac{1}{4}, \frac{3}{4})$).

Consider a SPE in $G(1,0)$. As in $G(0,0)$, player I accepts the type $(1,0)$. Since $\gamma^{\mathrm{I}}_{0,0} > b$ player I will veto (b,b). We claim that in a SPE player II will accept the type (b,b). Indeed, if he rejects this type, then the only types that can be hired are $(1,0)$ and $(1-\varepsilon, 1-\varepsilon)$, which arrive at the same frequency. Player I accepts the type $(1,0)$, and player II will then accept the type $(1-\varepsilon, 1-\varepsilon)$, since this type gives him the highest payoff. The payoff of player II would then be less than $\frac{1}{2}$, yet $\gamma^{\mathrm{II}}_{0,0} \geqslant \frac{1}{2}$, so that player II prefers to accept the type (b,b) (who will be rejected by player I) and continue to $G(0,0)$ than to reject her.

Since type (b,b) arrives frequently, the SPE payoff is close to (b,b), and therefore player I will accept the type $(1-\varepsilon, 1-\varepsilon)$. This implies that the SPE payoff in $G(1,0)$ is close to (b,b) for both players: with high probability the first candidate has type (b,b); if player I casts a veto, the game moves to $G(0,0)$, where, by (22) and (23), the unique SPE payoff is close to (b,b); if player I does not cast a veto then player II will accept this candidate, since otherwise. Thus, in the unique SPE, player I rejects the type (b,b), the other two types are hired, and the payoff $\gamma_{1,0}$ is:

$$\gamma^{\mathrm{I}}_{1,0} = (1-2\varepsilon)\gamma^{\mathrm{I}}_{0,0} + \varepsilon(1-\varepsilon) + \varepsilon = b + \varepsilon\big(4 - 2b + \varepsilon(-6 + 2\varepsilon + 4b)\big) > \gamma^{\mathrm{I}}_{0,0},$$
$$\gamma^{\mathrm{II}}_{1,0} = (1-2\varepsilon)\gamma^{\mathrm{II}}_{0,0} + \varepsilon(1-\varepsilon) = b + \varepsilon\big(2 - 4b + \varepsilon(-4 + 2\varepsilon + 4b)\big) < \gamma^{\mathrm{II}}_{0,0}.$$

Consider now an equilibrium in $G(0,1)$. As in $G(0,0)$ player II will accept the excellent type $(1-\varepsilon, 1-\varepsilon)$. Since $\gamma^{\mathrm{II}}_{0,0} < b$, player II will not veto (b,b); Since b is the minimal payoff to player I, it follows that player I will reject (b,b). Since $\gamma^{\mathrm{II}}_{0,0} > 0$ player II will veto $(1,0)$. Finally, player I will reject $(1,0)$, since by accepting it he will get $\gamma^{\mathrm{I}}_{0,0}$, whereas by rejecting it he will get $1-\varepsilon$. Therefore the unique equilibrium is that both players accept only $(1-\varepsilon, 1-\varepsilon)$, and the corresponding equilibrium payoff is $\gamma_{0,1} = (1-\varepsilon, 1-\varepsilon)$.

As we see, player I's expected payoff in $G(0,1)$ is higher than his expected payoff in $G(1,0)$, so that a donation paradox occurs: player I prefers to give his veto to player II.

Note that in $G(1,0)$ the unique Markovian SPE is progressive ($v < v+$, so that player II accepts the type (b,b) and forces player I to use his veto). On the other hand, in $G(0,1)$ the unique Markovian SPE is helpful ($u > u+$, so that player I does not accept the type $(1,0)$ that will make player II use his veto). This explains why player I prefers $G(0,1)$ over $G(1,0)$.

The donation paradox presented above cannot occur when F_2 is symmetric, and, in particular, when x and y are i.i.d.

Theorem 15. *Assume that F_2 is symmetric. If the SPE of $G(1,0)$ is unique, then player I never prefers $G(0,1)$ to $G(1,0)$.*

Proof. Denote the unique equilibrium payoff of $G(1,0)$ by $(\gamma^{\mathrm{I}}_{1,0}, \gamma^{\mathrm{II}}_{1,0})$. By Alpern and Gal (2008, Theorem 9), there is a symmetric threshold equilibrium in $G(0,0)$. As we have seen in the proof of Theorem 3, every equilibrium payoff in $G(0,0)$ can be extended into a Markovian SPE in $G(1,0)$. Since there is a unique equilibrium payoff in $G(1,0)$, it follows that we can assume that in the equilibrium in $G(1,0)$ the continuation payoff is $G(0,0)$ is symmetric.

One can verify that when one of the players does not have any veto rights, the order in which players play is irrelevant: by properly re-defining the strategies, an equilibrium in $G(0,1)$ can be turned into an equilibrium in $G(1,0)$. It follows that the unique equilibrium payoff in $G(0,1)$ is $(\gamma^{\mathrm{II}}_{1,0}, \gamma^{\mathrm{I}}_{1,0})$. We now show that $\gamma^{\mathrm{I}}_{1,0} \geqslant \gamma^{\mathrm{II}}_{1,0}$.

Assume to the contrary that $\gamma^{\mathrm{I}}_{1,0} < \gamma^{\mathrm{II}}_{1,0}$. We prove that if a candidate (x,y) with $x < y$ is hired with positive probability, then the candidate (y,x) is hired with probability 1. Since F_2 is symmetric, and the continuation payoff if player I casts a veto is symmetric, this will imply that $\gamma^{\mathrm{I}}_{1,0} \geqslant \gamma^{\mathrm{II}}_{1,0}$, a contradiction. Now, if such a candidate (x,y) is accepted with positive probability by player I then

$$y > x \geqslant \gamma^{\mathrm{I}}_{1,0},$$

so that player I accepts (y,x) with probability 1. If the candidate is accepted with positive probability by player II (and not vetoed by player I) then

$$y \geqslant \gamma^{\mathrm{II}}_{1,0} > \gamma^{\mathrm{I}}_{1,0},$$

so that (y,x) is accepted with probability 1 by player I. Since player II has no vetoes in $G(1,0)$ the result follows. □

Remark 16. We end the paper by commenting on the monotonicity of the equilibrium payoff as a function of the number of vetoes that the players have. Assume that x, y are i.i.d. with any non-constant distribution. Assume without loss of generality that $(1,1)$ is in the support of (x,y). Since one available strategy is to accept all candidates, all equilibrium payoffs in $G(0,0)$ are at least the expectation $E(x)$ (for both players), and since x and y are non-constant, one can verify that all equilibrium in $G(0,0)$ are higher than $E(x)$ (for both players). Also, in all equilibria in $G(\infty,0)$ player I obtains 1, since he can reject all candidates except those whose worth for him is 1. Since player I is indifferent among all candidates whose worth for him is 1, in $G(\infty,0)$ there may be multiple equilibria that differ in the payoff to player II. As we mentioned before, it follows from Alpern and Gal (2008) that the payoff to player I in $G(n,0)$ is monotonic increasing as a function of n. On the other hand, if helpful equilibrium exists for some n, then the payoff to player II is not monotonic in n.

References

Alpern, S., Gal, S., 2008. Analysis and design of selection committees: A game theoretic secretary problem. Int. J. Game Theory, in press, doi:10.1007/s00182-009-0159-4.

Brams, S.J., Davis, M.D., 1976. A game-theory approach to jury selection. Trial 12, 47–49.

Brams, S.J., Davis, M.D., 1978. Optimal jury selection: A game theoretic model for the exercise of peremptory challenges. Oper. Res. 26, 966–991.

DeGroot, M.H., Kadane, J.B., 1980. Optimal challenges for selection. Operations Res. 28, 952–968.

Enns, E.G., Ferenstein, E.Z., 1987. On a multi-person time-sequential game with priorities. In: Sequential Analysis, Design Methods and Applications. Taylor & Francis.

Eriksson, K., Sjöstrand, J., Strimling, P., 2007. Optimal expected rank in a two-sided secretary problem. Oper. Res. 55, 921–938.

Felsenthal, D., Machover, M., 1995. Postulates and paradoxes of relative voting power—A critical re-appraisal. Theory Dec. 38, 195–229.

Ferguson, T., 1989. Who solved the secretary problem? Statist. Sci. 4, 282–296.

Ferguson, T., 2001. Selection by Committee. eScholarship Repository, University of California. Available at http://repositories.edlib.org/uclastat/papers/2001010108.

Fisher, D., Schotter, A., 1978. The inevitability of the "paradox of redistribution" in the allocation of voting weights. Public Choice 33, 49–67.

Flesch, J., Thuijsman, F., Vrieze, K., 1996. Recursive repeated games with absorbing states. Math. Oper. Res. 21, 1016–1022.

Holler, M.J., Napel, S., 2004. Monotonicity of power and power measures. Theory Dec. 56, 93–111.

Kadane, J.B., Stone, C.A., Wallstrom, G., 1999. The donation paradox for peremptory challenges. Theory Dec. 47, 139–151.

Mashiach-Yakovi A., 2008. Subgame perfect equilibria in stopping games. Submitted for publication.

Mertens, J.F., 1987. Repeated games. In: Proceedings of the International Congress of Mathematicians. Berkeley, CA, pp. 1528–1577.

Ramsey, D., Szajowski, K., 2005. Bilateral approach to the secretary problem. In: Nowak, A.S., Szajowski, K. (Eds.), Advances in Dynamic Games: Applications to Economics, pp. 271–284.

Sakaguchi, M., 2005. Optimal stopping games where players have weighted privilege. In: Nowak, A.S., Szajowski, K. (Eds.), Advances in Dynamic Games: Applications to Economics, pp. 285–294.

Shapley, L.S., 1953. Stochastic games. Proc. Natl. Acad. Sci. U.S.A., 1095–1100.

Solan, E., Vieille, N., 2002. Correlated equilibrium in stochastic games. Games Econ. Behav. 38, 362–399.

Vrieze, O.J., Thuijsman, F., 1989. On equilibria in repeated games with absorbing states. Int. J. Game Theory 18, 293–310.