

The value of a draw*

Casilda Lasso de la Vega [†]

University of the Basque Country

Oscar Volij [‡]

Ben-Gurion University of the Negev

July 23, 2018

Abstract

We model a match as a recursive zero-sum game with three possible outcomes: player 1 wins, player 2 wins or there is a draw. We focus on matches whose point games also have three possible outcomes: player 1 scores the point, player 2 scores the point, or the point is drawn in which case the point game is repeated. We show that a value of a draw can be attached to each state so that an easily-computed stationary equilibrium exists in which players' strategies can be described as minimax behavior in the point games induced by these values.

Journal of Economic Literature Classification Numbers: C72, C73.

Keywords: Matches, stochastic games, recursive games.

*The authors thank the Spanish Ministerio de Economía y Competitividad (project ECO2015-67519-P) and the Gobierno Vasco (project IT568-13) for research support. Oscar Volij also thanks the Department of Foundations of Economics Analysis I at the University of the Basque Country for its kind hospitality. This paper was partly written during his stay there.

[†]casilda.lassodelavega@ehu.es

[‡]ovolij@bgu.ac.il

1 Introduction

A match is a recursive zero-sum game with three possible outcomes: player 1 wins, player 2 wins, or the game never ends. Play proceeds by steps from state to state. In each state players play a “point” and move to the next state according to transition probabilities jointly determined by their actions. Examples of matches include tennis, penalty shootouts and, you will forgive the repetition, chess matches.¹ In a chess match, two players play a sequence of chess games until some prespecified score is reached. For instance, the Alekhine–Capablanca match played in 1927 took the format known as first-to-6 wins, according to which the winner is the first player to win six games. Some matches are finite horizon games. As an example we have a best-of-seven playoff series. Indeed, this match will necessarily end in at most seven stages. A penalty shootout, on the other hand, is an infinite horizon game. It will never end if, for instance, every penalty kick is scored. Similarly, a first-to-6-wins chess match is also an infinite horizon game.² Matches can further be classified into binary and non-binary games. A penalty shootout is an example of the former and a chess match of the latter. The reason is that while each penalty kick has only two outcomes, either the goal is scored or it is not scored, a chess game may also end in a draw.

Matches have been the object of several empirical studies. For instance, Walker and Wooders [14] and Gauriot, Page and Wooders [4], using data on tennis, and Palacios-Huerta [7] using data on penalty kicks, show that players’ behavior is broadly consistent with the minimax hypothesis. Specifically, they show that professional players regard each point game as a one-shot zero-sum game and that their play is consistent with its equilibrium. On the other hand, Apesteguia and Palacios-Huerta [1] observe a first-kicker advantage in penalty shootouts and Gonzalez-Díaz and Palacios-Huerta [5] find a

¹The game of chess itself is also a match. In fact, in the first article to appear on game theory, Zermelo [16] models chess as a zero-sum recursive game.

²The 1984 Karpov-Kasparov match lasted five months and was aborted after 48 games when the partial score was 5-3. Coincidentally, the longest penalty shootout to date also had 48 kicks.

similar anomaly in chess matches. This last paper also offers a brief theoretical analysis of a particular finite chess match.

A theoretical foundation of Walker and Wooders [14] appears in Walker, Wooders and Amir [15]. They define the class of binary Markov games and model tennis as one such game. They show that under a certain monotonicity condition, minimax behavior in each of the point games constitutes an equilibrium of the whole match. Namely, by maximizing the lowest probability of his scoring each point, each player is best responding to the other player's also maximizing the lowest probability of his scoring each point. This result implies that as long as the monotonicity condition holds, binary Markov games have stationary equilibria that dictate behavior that depends only on the current point game and therefore is independent of the structure of the match.

Walker, Wooders and Amir [15] define a binary Markov game as a binary match in which never-ending play is the worst outcome for both players, a feature that renders their matches non-zero sum games. An alternative way to model matches is to view them as standard recursive games as defined by Everett [2]. A recursive game is a stochastic game where non-zero payoffs are obtained only at absorbing states. In other words, in a never-ending play players get a payoff of 0. Besides being a natural way to model matches,³ an advantage of this approach is that one can apply well-established results from the theory of zero-sum stochastic games.

In this paper we model matches as zero-sum recursive games and focus on those whose point games have three possible outcomes: player 1 scores the point, player 2 scores the point, or (something that happens with probability less than 1) the point is drawn, in which case the point game is repeated. We call these games *quasi-binary matches*. The results we obtain extend those of Walker, Wooders and Amir [15]. We show that a value of a draw can be attached to each state so that quasi-binary matches always have an easily-computed stationary equilibrium which prescribes minimax play

³In his article on the game of chess, Zermelo [16] states: "Such a possible endgame \mathbf{q} can find its natural end in a "checkmate" . . . , but could also – at least theoretically – go on forever, in which case the game would without doubt have to be called a draw . . ." (see Schwalbe and Walker [9]).

in the point games induced by these values. Moreover, the value of a draw attached to a given state depends only on the point played in it and as a result equilibrium behavior at that state is independent of the structure of the match. We can conclude then that Walker, Wooders and Amir's [15] results extend to quasi-binary games and thus are robust to the modeling choice concerning payoffs at infinite plays. Moreover, the extension applies to all quasi-binary games and not just those which satisfy a restrictive monotonicity condition.

For binary matches where minimax play induces a finite history with probability one, it is not surprising that both modeling choices yield the same result. After all, as Walker, Wooders and Amir [15] remark (see their footnote 8), in this case binary Markov games are zero-sum games with probability one. What we show is that the conclusions of Walker, Wooders and Amir's equilibrium and minimax theorems continue to hold even if minimax play induces infinite histories with positive probability, *as long as matches are modeled as zero-sum games*.

Quasi-binary games are a modest extension of binary games. Indeed, whereas they allow for a draw in their component point games, they impose that the point game be repeated after a draw. Also, for any pair of actions, the probability of a draw must be less than one. However, if we tried to further extend this class of games, our results would fail to be valid, as shown by two examples we offer in the paper. Additionally, a nice feature of quasi-binary games is that the concept of the value of a draw, which turns out to be helpful for the interpretation of the equilibrium, emerges naturally there.

In order to describe these equilibria, note that since in a quasi-binary game the probability of staying in the current state, say k , is less than one, players will eventually move to one of two different states. Label them $w(k)$ and $\ell(k)$. If they move to $w(k)$ we say that player 1 wins the point and if they move to state $\ell(k)$ we say that player 1 loses the point. Finally, if they stay in the current state we say that the point is drawn. Note that since there are two different states to which players can move from state k , there are two different ways to select a labeling. Even so, once a labeling is chosen, we can

define a simple zero-sum matrix game as follows. First we assign a value e^k to the draw in the current state and then we define the payoffs to player 1 as his expected earnings when winning the point is worth 1, losing the point is worth 0, and a draw is worth e^k . In this paper we show that there is a labeling $w(k), \ell(k)$ of the successors of each state k , and a value e^k of the draw in the respective point games such that minimax play in the above zero-sum matrix games constitutes an equilibrium of the match. We also show that if the game satisfies a mild monotonicity condition, namely that there is no state at which in equilibrium players are indifferent between the two states they can move to, every stationary equilibrium of the match prescribes minimax play in these zero-sum games.

It is worth noting that the conclusion of our main result is not valid for all of Walker, Wooders and Amir’s binary Markov games. Indeed, we provide an example at the end of the paper of a binary Markov game for which, no matter how we interpret the outcomes of the point games, minimax play does not necessarily lead to an equilibrium of the game.

The paper is organized as follows. Section 2 offers a motivating example and introduces the basic definitions. Section 3 defines the concept of the value of a draw and shows that it satisfies some interesting properties. In Section 4 we formulate and prove the results.

2 Matches

2.1 A motivating example

Before we introduce the formal model, we illustrate the main result by means of the following simple match. Two players play a sequence of 2×2 “simplified chess” games. Each game may end in a victory for either player or in a draw. The winner earns one point and the match ends as soon as the score difference is either 2 or -2. Formally, there are three non-absorbing states, 1, 0, and -1, corresponding to each partial score,

and two absorbing states, 2 and -2. Let's adopt the labeling according to which when player 1 wins the chess game played at state k , for $k = 1, 0, -1$, play moves to state $k + 1$, and when he loses it there is a transition to $k - 1$. When the partial score is 0 player 1 plays with the white pieces and the chess game is governed by the following matrix of probabilities:

$$P^W = \begin{pmatrix} (\frac{2}{3}, \frac{1}{3}, 0) & (\frac{8}{27}, \frac{1}{3}, \frac{10}{27}) \\ (0, \frac{1}{2}, \frac{1}{2}) & (\frac{2}{3}, \frac{1}{3}, 0) \end{pmatrix}.$$

Each entry displays the probabilities of player 1 winning, drawing or losing the point when the corresponding actions are chosen.⁴ For instance, when player 1 chooses his first action and player 2 chooses his second action, player 1 wins the point with probability $8/27$, loses the point with probability $10/27$, and there is a draw with probability $1/3$. As soon as one of the players wins the point and the partial score becomes 1 or -1, they go on to play a new chess game in which player 1 has the black pieces. Correspondingly, this new game is governed by the following matrix of probabilities:

$$P^B = \begin{pmatrix} (0, \frac{1}{3}, \frac{2}{3}) & (\frac{1}{2}, \frac{1}{2}, 0) \\ (\frac{10}{27}, \frac{1}{3}, \frac{8}{27}) & (0, \frac{1}{3}, \frac{2}{3}) \end{pmatrix}.$$

Here too, the entries are the probabilities that player 1 wins, draws or loses the point when the corresponding action pair is chosen. Players continue playing this game until one of them wins the point. If the player who has the score advantage wins the point the match ends. If the player with the score disadvantage wins the point, the partial score becomes 0 again and they go back to playing a chess game where player 1 has the white pieces.

Although matrices P^W and P^B represent the strategic interaction involved in each of the chess games, they themselves are not games. In order to transform them into games

⁴We are aware that in real chess, the outcome of a pair of strategies is deterministic. We hope chess enthusiasts will forgive this distortion.

we need to specify the proportion of the point at stake a draw represents. Consider for instance the matrix P^W . If a draw is worth $\varepsilon \in [0, 1]$ of a point, then by taking the expected value of the point earned by player 1, P^W can be transformed into the following matrix game:

$$P^W(\varepsilon) = \begin{pmatrix} \frac{2}{3} + \frac{1}{3}\varepsilon & \frac{8}{27} + \frac{1}{3}\varepsilon \\ \frac{1}{2}\varepsilon & \frac{2}{3} + \frac{1}{3}\varepsilon \end{pmatrix}.$$

Routine calculations show that the value of this matrix is $\frac{24+16\varepsilon-3\varepsilon^2}{56-9\varepsilon}$, and that in particular when $\varepsilon = 2/3$ the value of the matrix is also $2/3$. Namely, $2/3$ is a fixed point of the function that assigns to each $\varepsilon \in [0, 1]$ the value of $P^W(\varepsilon)$. We call this fixed point the value of the draw when player 1 plays with the white pieces, and we call the corresponding matrix $P^W(2/3)$ the associated point game. One can also check that the equilibrium strategies of this point game are $((3/5, 2/5), (2/5, 3/5))$.

Similarly, one can check that when the draw in the chess game governed by P^B is worth ε of a point, the associated matrix game is

$$P^B(\varepsilon) = \begin{pmatrix} \frac{2}{3}\varepsilon & \frac{1}{2} + \frac{1}{2}\varepsilon \\ \frac{10}{27} + \frac{1}{3}\varepsilon & \frac{1}{3}\varepsilon \end{pmatrix}$$

and that the value of this game when a draw is worth $1/3$ of a point is also $1/3$. In other words, the value of a draw when player 1 plays with black is $1/3$, and the associated point game is $P^B(1/3)$. Furthermore, equilibrium strategies of the associated point game $P^B(1/3)$ are $((2/5, 3/5), (3/5, 2/5))$.

Our main result will imply that choosing the mixed action $(3/5, 2/5)$ when playing with the white pieces, and choosing the mixed action $(2/5, 3/5)$ when playing with the black pieces is an optimal strategy for each of the players in the match. Furthermore, our second result shows that the corresponding pair of strategies is the only stationary equilibrium of the match. Notice that this equilibrium dictates that in each point game players should behave in a way that depends only on the chess game played. In particular,

since when the partial score is 1 or -1 the chess games played are the same, equilibrium behavior is also the same. Also notice that we have been able to compute the equilibrium actions in each state using only the matrix of probabilities that is relevant to that state.

Next section starts with the formal model, which extends the analysis of the foregoing example to all quasi-binary matches.

2.2 Basic definitions

Consider the following zero-sum stochastic game, which we call *a match*. There are two players, 1 and 2, and a set of states $S = \{0, 1, \dots, K + 1\}$. States 0 and $K + 1$ are absorbing states which if reached the match ends. In state $k \in S$, the actions available to players 1 and 2 are labeled by the integers $1, \dots, I_k$ and $1, \dots, J_k$, respectively. Without loss of generality we assume that for all k , $I_k = I$ and $J_k = J$ and denote the action sets of player 1 and 2 by \mathcal{I} and \mathcal{J} , respectively. Players are endowed with action sets in states 0 and $K + 1$ only for notational convenience. A mixed action for player 1 is a probability distribution over \mathcal{I} and a mixed action for player 2 is a probability distribution over \mathcal{J} . We denote the sets of mixed actions of player 1 and 2 by $\Delta_{\mathcal{I}}$ and $\Delta_{\mathcal{J}}$, respectively. For any $I \times J$ matrix game A , $\text{val}(A)$ denotes its value. A mixed action $x \in \Delta_{\mathcal{I}}$ is said to be *optimal* for player 1 in A if it guarantees that he gets a payoff of at least $\text{val}(A)$. Similarly, a mixed action $y \in \Delta_{\mathcal{J}}$ is said to be optimal for player 2 in A if it guarantees that player 1 gets a payoff of at most $\text{val}(A)$. Recall that for $A = (a_{ij} | i \in \mathcal{I}, j \in \mathcal{J})$ and $B = (b_{ij} | i \in \mathcal{I}, j \in \mathcal{J})$, $|\text{val}(A) - \text{val}(B)| \leq \max_{ij} |a_{ij} - b_{ij}|$ and that if $b_{ij} = \alpha a_{ij} + \beta$ for some $\alpha > 0$ and $\beta \in \mathbb{R}$ and for all $i \in \mathcal{I}$ and $j \in \mathcal{J}$, then $\text{val}(B) = \alpha \text{val}(A) + \beta$.

For each state $k \in S$ there is a matrix

$$P^k = (p_{ij}^k | i \in \mathcal{I}, j \in \mathcal{J})$$

of probability distributions on the set of states S . Namely, for each pair of actions i, j

of player 1 and 2, respectively, $p_{ij}^k = (p_{ij}^{kk'})_{k' \in S}$ where

$$p_{ij}^{kk'} \geq 0 \text{ and } \sum_{k' \in S} p_{ij}^{kk'} = 1.$$

Matrices P^0 and P^{K+1} are introduced for notational convenience; since states 0 and $K + 1$ are absorbing, $p_{ij}^{00} = p_{ij}^{K+1, K+1} = 1$ for all $i \in \mathcal{I}$ and $j \in \mathcal{J}$. We will henceforth refer to P^k as the *point matrix* at k .

The interpretation of the match is as follows. In state $k = 1, \dots, K$, after player 1 chooses an action $i \in \mathcal{I}$ and player 2 chooses an action $j \in \mathcal{J}$ they move to state $k' \in S$ with probability $p_{ij}^{kk'}$. If state 0 is reached the match ends and player 1 wins. If state $K + 1$ is reached, the match ends and player 2 wins. If neither state 0 nor $K + 1$ is ever reached, the match is drawn. The payoffs are undiscounted; the timing of a victory or a defeat does not affect preferences.

In order to define the match we need to specify the initial state and, for each player, his set of available strategies and his payoff function. But first we need some definitions. The set of histories of length $t = 0, 1, 2, \dots$ is denoted by $H_t = S \times (\mathcal{I} \times \mathcal{J} \times S)^t$. A typical history of length t is $h_t = (s_0, (i_1, j_1, s_1), \dots, (i_t, j_t, s_t)) \in H_t$. Here, the initial state is $s_0 \in S$ and at stage $\tau = 1, \dots, t$, players chose actions i_τ and j_τ as a result of which the state becomes s_τ . By the end of h_t , the state is s_t . The set of all finite histories is denoted by $H = \cup_{t \geq 0} H_t$.

A player's strategy is a specification of a mixed action for each stage conditional on the current state and on the history of play up to that stage. Formally, a strategy for player 1 is a map $\chi : H \rightarrow \Delta_{\mathcal{I}}$ that prescribes a mixed action $\chi(h_t) = (\chi_1(h_t), \dots, \chi_I(h_t))$ to be used by player 1 after every finite history h_t . Similarly, a strategy for player 2 is a map $\psi : H \rightarrow \Delta_{\mathcal{J}}$ that prescribes a mixed action $\psi(h_t) = (\psi_1(h_t), \dots, \psi_J(h_t))$ to be used by player 2 after every finite history h_t . *Stationary strategies* are strategies whose prescriptions depend only on the current state. Thus, a stationary strategy for player 1 can be represented by a vector $\vec{x} = (x^0, \dots, x^{K+1})$, where for each $k \in S$,

$x^k = (x_1^k, \dots, x_I^k)$ is a mixed action for player 1. Similarly, a stationary strategy for player 2 is a vector $\vec{y} = (y^0, \dots, y^{K+1})$ of mixed actions for player 2. We denote the sets of strategies for players 1 and 2 by X and Y , respectively, and their subsets of stationary strategies by \vec{X} and \vec{Y} . Given an initial state $k \in S$, a pair of strategies χ and ψ induces a probability distribution on the histories of length t as follows. For histories $h_0 \in H_0$, of length 0,

$$\pi_k^{\chi, \psi}(h_0) = \begin{cases} 1 & \text{if } h_0 = k \\ 0 & \text{otherwise.} \end{cases}$$

And for histories of length $t = 1, 2, \dots$ this probability distribution is defined inductively as follows. For $h_t = h_{t-1} \circ (i_t, j_t, s_t)$,

$$\pi_k^{\chi, \psi}(h_t) = \pi_k^{\chi, \psi}(h_{t-1}) \chi_{i_t}(h_{t-1}) \psi_{j_t}(h_{t-1}) p_{i_t j_t}^{s_{t-1} s_t}.$$

Consequently, given an initial state k and a pair of strategies χ and ψ the probability that at stage $t = 1, 2, \dots$, the current state is k' is given by

$$\mu_t^{k k'}(\chi, \psi) = \sum_{\{h_t \in H_t : s_t = k'\}} \pi_k^{\chi, \psi}(h_t). \quad (1)$$

Since states 0 and $K + 1$ are absorbing, the probability sequences $\{\mu_t^{k0}(\chi, \psi)\}_{t=1}^{\infty}$ and $\{\mu_t^{kK+1}(\chi, \psi)\}_{t=1}^{\infty}$ are non-decreasing and bounded. Therefore they have limits, which are denoted $\mu_{\infty}^{k0}(\chi, \psi)$ and $\mu_{\infty}^{kK+1}(\chi, \psi)$, respectively. Each of these limits is the probability that player 1 and player 2, respectively, eventually wins the match conditional on the initial state being k when they choose the strategy pair (χ, ψ) .

We can now define the (undiscounted) match Γ^k which starts at state $k \in S$. Formally, Γ^k is the zero-sum game where the sets of strategies of player 1 and 2 are X and Y , respectively, and player 1's payoff function $u^k : X \times Y \rightarrow [-1, 1]$ is defined by $u^k(\chi, \psi) = \mu_{\infty}^{k0}(\chi, \psi) - \mu_{\infty}^{kK+1}(\chi, \psi)$. Player 2's payoff function is consequently $-u^k(\chi, \psi)$. Note that Γ^0 and Γ^{K+1} are degenerate games with $u^0(\chi, \psi) = 1$ and $u^{K+1}(\chi, \psi) = -1$. Also note that players are indifferent among all pairs of strategies that induce equal

chances of winning and losing the match. We denote by Γ the collection of matches $\{\Gamma^k : k = 1, \dots, K\}$ and remark that Γ is fully determined by the set of states S and by the set of point matrices $(P^k)_{k=1}^K$.

We mention that if we endowed players 1 and 2 with the payoff functions $W_1^k(\chi, \psi) = \mu^{k0}(\chi, \psi)$ and $W_2^k(\chi, \psi) = \mu^{kK+1}(\chi, \psi)$, respectively, we would obtain a match as defined by Walker, Wooders and Amir [15]. Since it is not necessarily true that any pair of strategies leads to an absorbing state with probability 1, this match is not a constant-sum game. Also, for future reference we define, for $\lambda \in (0, 1)$, the discounted match Γ_λ^k which starts at state $k \in S$ as the zero-sum game where player 1's payoff function $u_\lambda^k : X \times Y \rightarrow [-1, 1]$ is defined by $u_\lambda^k(\chi, \psi) = (1-\lambda) \sum_{t=1}^{\infty} \lambda^{t-1} (\mu_t^{k0}(\chi, \psi) - \mu_t^{kK+1}(\chi, \psi))$ and player 2's payoff function is $-u_\lambda^k(\chi, \psi)$.

The number v^k is said to be the value of Γ^k if $\sup_{\chi \in X} \inf_{\psi \in Y} u^k(\chi, \psi) = v^k = \inf_{\psi \in Y} \sup_{\chi \in X} u^k(\chi, \psi)$. If v^k is the value of Γ^k for $k = 1, \dots, K$ we say that (v^1, \dots, v^K) is the value of Γ . A strategy pair $(\chi^*, \psi^*) \in X \times Y$ is an equilibrium of Γ^k if

$$u^k(\chi, \psi^*) \leq u^k(\chi^*, \psi^*) \leq u^k(\chi^*, \psi) \quad \text{for all } \chi \in X, \psi \in Y.$$

In this case, $u^k(\chi^*, \psi^*)$ is clearly the value of Γ^k . We say that $(\chi^*, \psi^*) \in X \times Y$ is an equilibrium of Γ if it is an equilibrium of Γ^k for all $k \in \{1, \dots, K\}$.

Matches are recursive games as defined by Everett [2]. Recursive games are a special case of stochastic games, which were earlier introduced by Shapley [10]. Everett [2] shows that recursive games have a value that can be approached with stationary strategies. Mertens and Neyman [6] prove more generally that when streams of payoffs are undiscounted all stochastic games with finite state and action spaces have a value. Thuijsman and Vrieze [12] provide an ingenious constructive proof of Everett's result, showing that ε -optimal strategies can be built by appropriately modifying optimal strategies in the discounted games. For further results in recursive games see Thuijsman [11] Flesch, Thuijsman and Vrieze [3], Vieille [13] and the references therein.

The point matrix P^k represents the point played at state k . Note that P^k is not a

game since its entries are probability distributions on S . However, it can be transformed into a zero-sum game by assigning values to the states and averaging them according to the entries of P^k . More specifically, for any $\alpha = (\alpha^1, \dots, \alpha^K) \in \mathbb{R}^K$ we can define the matrix game $A^k(\alpha)$ as follows:

$$A^k(\alpha) = (p_{ij}^{k0} + \sum_{k'=1}^K p_{ij}^{kk'} \alpha^{k'} - p_{ij}^{kK+1} \mid i \in \mathcal{I}, j \in \mathcal{J})$$

and the associated value mapping $\alpha \rightarrow (\text{val} A^k(\alpha))_{k=1}^K$. Everett [2] showed that every recursive game with bounded payoffs has a value, and that this value is a fixed point of the associated value mapping. Applied to the present setting this result yields the following observation which plays a fundamental role in our analysis.

Observation 1 For $k = 1, \dots, K$, Γ^k has a value v^k and this value satisfies $v^k = \text{val}(A^k(v^1, \dots, v^K))$.

Although Γ^k has a value, it may not have an equilibrium. See Everett's [2] Example 1, reproduced in Section 4.1 below.

2.3 Stationary strategies

Given an initial state $k \in S$, a pair of stationary strategies induce a Markov chain that allows us to compute the transition probabilities defined in (1) recursively. Specifically, a pair of stationary strategies (\vec{x}, \vec{y}) induces a Markov matrix $M(\vec{x}, \vec{y}) = (\mu^{ss'}(\vec{x}, \vec{y}) \mid s, s' \in S)$ whose transition probabilities are given by the probability of moving to state s' conditional on the current state being s :

$$\begin{aligned} \mu^{ss'}(\vec{x}, \vec{y}) &= \frac{\sum_{\{h_t: s_t=s\}} \pi_k^{\vec{x}, \vec{y}}(h_t) \sum_{i=1}^I \sum_{j=1}^J x_i^s y_j^s p_{ij}^{ss'}}{\sum_{\{h_t: s_t=s\}} \pi_k^{\vec{x}, \vec{y}}(h_t)} \\ &= \sum_{i=1}^I \sum_{j=1}^J x_i^s y_j^s p_{ij}^{ss'}. \end{aligned} \tag{2}$$

As is well known, this probability does not depend on the initial state k .

Note that $\mu_1^{kk'}(\vec{x}, \vec{y}) = \mu^{kk'}(\vec{x}, \vec{y})$ and that the probabilities $\mu_t^{kk'}(\vec{x}, \vec{y})$ defined in (1) satisfy the recursive relation

$$\mu_t^{kk'}(\vec{x}, \vec{y}) = \sum_{s \in S} \mu_{t-1}^{ks}(\vec{x}, \vec{y}) \mu^{sk'}(\vec{x}, \vec{y}) \quad k \in S.$$

In other words, they are none other than the entries of the t -th power of $M(\vec{x}, \vec{y})$.

3 Quasi-binary matches and the value of a draw

In this paper we restrict attention to a particular class of simple matches which we now define. Let Γ be a match characterized by the point matrices $P^k = (p_{ij}^k | i \in \mathcal{I}; j \in \mathcal{J})$, for $k = 1, \dots, K$. For each state k , define the set of its immediate successors, or simply *successors*, to be

$$S(k) = \{k' \in S : p_{ij}^{kk'} > 0, \text{ for some } (i, j) \in \mathcal{I} \times \mathcal{J}\}.$$

This set contains the states that can possibly be reached from state k in a single step. Successors of k that are not k itself are called *proper successors*.

Definition 1 A match is quasi-binary if for each state $k = 1, \dots, K$ the number of its proper successors is exactly two, and $p_{ij}^{kk} < 1$ for all $i \in \mathcal{I}, j \in \mathcal{J}$.

Although our results are stated for the class of quasi-binary matches, they still hold for the larger class that includes those matches where some state k has a single proper successor, even if $p_{ij}^{kk} = 1$ for some $i \in \mathcal{I}, j \in \mathcal{J}$. In this case, the proof treats k as its second proper successor. (We will provide more details in footnote 5 later). For the sake of brevity, however, we decided to drop these matches from the class of quasi-binary games.

In a quasi-binary match each state $k = 1, \dots, K$ has only two proper successors. We denote them by $w(k)$ and $\ell(k)$. If the game moves to state $w(k)$ we say that player 1 won the point played at k . If the game moves to state $\ell(k)$ we say that player 1 lost the point played at k . And if the game stays in state k we say that the point played at k ended in a draw. We denote by (w, ℓ) the labeling $(w(k), \ell(k))_{k=1}^K$.

We can take advantage of the labeling (w, ℓ) to transform the point matrix P^k into a matrix game as follows. We first award player 1 a payoff of 1 if he wins the point, a payoff of 0 if he loses the point and a payoff of ε if the point is drawn, and then replace the distribution p_{ij}^k in the ij th entry by the corresponding expected payoff $p_{ij}^{kw(k)} + p_{ij}^{kk}\varepsilon$. Formally, for each $\varepsilon \in [0, 1]$ we define the matrix game $P^k(\varepsilon)$ by letting its ij th entry be $p_{ij}^{kw(k)} + p_{ij}^{kk}\varepsilon$, namely the expected value of the point played at k when players choose the action pair (i, j) and a draw is valued at ε .⁵ Note that $P^k(\varepsilon)$ depends on the labeling choice $w(k), \ell(k)$. Consequently, all the ancillary definitions in this section depend on this choice.

The question we want to address is the following: Is there a labeling (w, ℓ) and an associated value of the draw e^k for each $k \in \{1, \dots, K\}$ so that two stationary strategies $\vec{x}^* = (x^0, \dots, x^{K+1})$ and $\vec{y}^* = (y^0, \dots, y^{K+1})$ constitute an equilibrium of Γ if for all $k \in \{1, \dots, K\}$, (x^k, y^k) is an equilibrium of $P^k(e^k)$? Our main theorem will give a positive answer to this question. Meanwhile, the next proposition singles out, given a labeling, a candidate for a suitable value of the draw.

Proposition 1 Let Γ be a quasi-binary match and let (w, ℓ) be a labeling. For $k = 1, \dots, K$, let $f^k : [0, 1] \rightarrow [0, 1]$ be the function defined by $f^k(\varepsilon) = \text{val}(P^k(\varepsilon))$. Then f^k has a unique fixed point.

Proof : Since the entries of $P^k(\varepsilon)$ are in $[0, 1]$ and are non-decreasing in ε , f^k is a

⁵ If a state k had only one proper successor we could treat k as the missing proper successor and denote these successors by $w(k)$ and $\ell(k)$. The matrix $P^k(\varepsilon)$ would then be defined as $\{p_{ij}^{kw(k)} | i \in \mathcal{I}, j \in \mathcal{J}\}$ and with this amended definition, the ensuing analysis would remain valid.

nondecreasing function that maps the interval $[0, 1]$ into itself. Therefore, by Tarski's fixed-point theorem f^k has a fixed point, which we denote e^k .

Assume that $\hat{\varepsilon}^k$ is another fixed point of f^k . Then,

$$\begin{aligned}
|\hat{\varepsilon}^k - e^k| &= |f^k(\hat{\varepsilon}^k) - f^k(e^k)| \\
&= |\text{val}(P^k(\hat{\varepsilon}^k)) - \text{val}(P^k(e^k))| \\
&\leq \max_{ij} |(p_{ij}^{kw(k)} + p_{ij}^{kk} \hat{\varepsilon}^k) - (p_{ij}^{kw(k)} + p_{ij}^{kk} e^k)| \\
&= |\hat{\varepsilon}^k - e^k| \max_{ij} p_{ij}^{kk} \\
&< |\hat{\varepsilon}^k - e^k|
\end{aligned}$$

where we have used the assumption that $p_{ij}^{kk} < 1$ for all $i \in \mathcal{I}$ and all $j \in \mathcal{J}$. But since the above inequality is absurd, we conclude that e^k is the only fixed point of f^k . \square

We denote by e^k the unique fixed point identified in the above proposition and call it the *value of the draw* in state k (with respect to (w, ℓ)). We also call $P^k(e^k)$ the *point game* played at k . Notice that in order to compute the value of the draw in state k only the point matrix P^k and the labeling (w, ℓ) are needed. In particular, no prior knowledge of the value of Γ is required. The next proposition shows, however, that when the labeling happens to be such that $v^{w(k)} > v^{\ell(k)}$, the value of the draw at k bears an interesting relationship with the values of the successors of k .

Proposition 2 Let Γ be a quasi-binary match, let (v^1, \dots, v^K) be its value and extend it so that $v^0 = 1$ and $v^{K+1} = 0$. Let (w, ℓ) be a labeling and let k be a state such that $v^{w(k)} > v^{\ell(k)}$. Also, let e^k be the unique fixed point identified in Proposition 1. Then,

$$e^k = \frac{v^k - v^{\ell(k)}}{v^{w(k)} - v^{\ell(k)}}.$$

Proof : Denote $\varepsilon^k = (v^k - v^{\ell(k)}) / (v^{w(k)} - v^{\ell(k)})$. By Proposition 1, the value of the

draw in state k is the unique fixed point of the function $f^k : [0, 1] \rightarrow [0, 1]$ given by $f^k(\varepsilon) = \text{val}(P^k(\varepsilon))$. Therefore, it is enough to show that ϵ^k is a fixed point of f^k . Recall that by Observation 1 $v^k = \text{val}(A^k(v^1, \dots, v^k))$ where $A^k(v) = (p_{ij}^{kw(k)} v^{w(k)} + p_{ij}^{kk} v^k + p_{ij}^{k\ell(k)} v^{\ell(k)})_{i \in \mathcal{I}, j \in \mathcal{J}}$. But note that $A^k(v)$ and $P^k(\epsilon^k)$ are strategically equivalent. Indeed, for $i \in \mathcal{I}$ and $j \in \mathcal{J}$ the ij th entry of the matrix $A(v)$ can be written

$$A_{ij}^k(v) = (p_{ij}^{kw(k)} + p_{ij}^{kk} \epsilon^k)(v^{w(k)} - v^{\ell(k)}) + v^{\ell(k)}$$

where $v^{w(k)} - v^{\ell(k)} > 0$. Therefore,

$$\text{val}(A^k(v)) = \text{val}(P^k(\epsilon^k))(v^{w(k)} - v^{\ell(k)}) + v^{\ell(k)}$$

and consequently,

$$\begin{aligned} \text{val}(P^k(\epsilon^k)) &= \frac{\text{val}(A^k(v)) - v^{\ell(k)}}{v^{w(k)} - v^{\ell(k)}} \\ &= \frac{v^k - v^{\ell(k)}}{v^{w(k)} - v^{\ell(k)}} = \epsilon^k. \end{aligned}$$

□

The foregoing proposition justifies calling ϵ^k the value of the draw in state k with respect to (w, ℓ) . To see this, notice that from state k , players will eventually move to one of its proper successors, $w(k)$ or $\ell(k)$, in which case player 1 will get (assuming optimal play) a payoff to $v^{w(k)}$, or $v^{\ell(k)}$, respectively. Therefore, since $v^{w(k)} > v^{\ell(k)}$, player 1 has a guaranteed expected payoff to $v^{\ell(k)}$ and hence what is really at stake in state k is $v^{w(k)} - v^{\ell(k)}$. When the point is drawn, the players remain in state k , in which case player 1 gets an expected payoff to v^k . Namely, he nets a proportion $\frac{v^k - v^{\ell(k)}}{v^{w(k)} - v^{\ell(k)}}$ of what is at stake. The above proposition shows that ϵ^k , the unique fixed point identified in Proposition 1, is precisely this proportion – hence its interpretation as the value of a draw.

The next definition identifies those stationary strategies which at every state dictate mixed actions that are optimal in the respective point games. According to these strategies, behavior in each state k depends only on the matrix P^k and, in particular, is independent of the structure of the match in all the other states.

Definition 2 Let Γ be a quasi-binary match, (w, ℓ) be a labeling, and for $k = 1, \dots, K$ let e^k be the value of the draw in k and $P^k(e^k)$ the point game played at k with respect to (w, ℓ) . Also, let $\vec{x} = (x^k)_{k=0}^{K+1} \in \vec{X}$ and $\vec{y} = (y^k)_{k=0}^{K+1} \in \vec{Y}$ be two stationary strategies, one for each player. We say that (\vec{x}, \vec{y}) is a minimax-stationary strategy pair with respect to (w, ℓ) if for all $k = 1, \dots, K$, (x^k, y^k) is an equilibrium of $P^k(e^k)$.

It follows from Proposition 1 that if (\vec{x}, \vec{y}) is a pair of minimax-stationary strategies then x^k guarantees that player 1 gets a payoff of at least e^k in $P^k(e^k)$ and y^k guarantees that player 1 gets at most e^k in $P^k(e^k)$. Notice that minimax-stationary strategies always exist.

The following observation states that when players behave according to a minimax-stationary strategy pair, the probability of player 1 eventually winning the point game played at k is precisely the value of the draw in state k .

Observation 2 Let Γ be a quasi-binary match, (w, ℓ) be a labeling and let (\vec{x}, \vec{y}) be a minimax-stationary strategy pair w.r.t (w, ℓ) . Then the value of the draw at k is the corresponding probability of eventually leaving k and transiting to $w(k)$. Formally, for $k = 1, \dots, K$

$$e^k = \frac{\mu^{kw(k)}(\vec{x}, \vec{y})}{1 - \mu^{kk}(\vec{x}, \vec{y})}.$$

Proof: Since $\vec{x} = (x^0, \dots, x^{K+1})$ and $\vec{y} = (y^0, \dots, y^{K+1})$ constitute a pair of minimax-stationary strategies, for $k = 1, \dots, K$, (x^k, y^k) is an equilibrium of $P^k(e^k)$, and $e^k = \text{val}(P^k(e^k))$,

$$e^k = \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} x_i^k y_j^k (p_{ij}^{kw(k)} + p_{ij}^{kk} e^k)$$

which, using equation (2) can be written as $e^k = \mu^{kw(k)}(\vec{x}, \vec{y}) + \mu^{kk}(\vec{x}, \vec{y})e^k$. Since $p_{ij}^{kk} < 1$ for all $i \in \mathcal{I}$ and all $j \in \mathcal{J}$, we have that $\mu^{kk}(\vec{x}, \vec{y}) < 1$. Therefore, solving for e^k we obtain the result. \square

4 Minimax-stationary strategies and equilibrium

We have seen that given a labeling (w, ℓ) we can associate to each state k a value of the draw e^k and a point game $P^k(e^k)$. Additionally, the point games $P^k(e^k)$ induce stationary strategies in Γ in a natural way: they prescribe that players choose at k mixed actions that are optimal in $P^k(e^k)$. In this section we will find a particular labeling all of whose induced minimax-stationary strategies constitute an equilibrium of the match. Specifically, we will show the following.

Theorem 1 Let Γ be a quasi-binary match. There exists a labeling such that any pair of minimax-stationary strategies with respect to it constitutes an equilibrium of Γ .

Before proving the theorem we discuss the result.

4.1 Discussion

a) *Interpretation of the result.* Not only does Theorem 1 show the existence of equilibrium, but it also identifies one with a particularly appealing interpretation. The labeling $w(k), \ell(k)$ identified in the theorem, along with the associated value of a draw e^k , suggests that moving to $w(k)$ can be seen as player 1 winning the point played at k , moving to $\ell(k)$ as player 2 winning the point, and drawing as if the point was shared in the proportions $(e^k, 1 - e^k)$. Theorem 1 identifies an equilibrium in which both players adopt this interpretation and aim at maximizing their respective expected shares of the point at stake. We will later show that not all stationary equilibria admit such an interpretation.

b) *Independence of the equilibrium behavior in state k .* Theorem 1 states that there exists an endogenously determined labeling such that any pair of minimax-stationary strategies with respect to it constitutes an equilibrium of Γ . Notice that these strategies dictate behavior in state k that depends on the point matrices in states different from k only to the extent that they affect the equilibrium labeling. Therefore, any modification in the structure of the match that involves neither a change in the point matrix P^k nor in the equilibrium labeling, will leave the equilibrium behavior in state k unaffected. This is not the case for the discounted match Γ_λ . Indeed, even a small perturbation of the point matrix in a state will generically affect equilibrium behavior in all other states.⁶

c) *Computation of the value of the match.* Theorem 1 allows us to compute the value v of Γ in a conceptually easy manner. To see this, for each of the 2^K possible labelings l , let (\vec{x}_l, \vec{y}_l) be a minimax-stationary strategy pair with respect to it, and let $(u^s(\vec{x}_l, \vec{y}_l))_{s \in S}$ be the corresponding payoffs. (Recall that minimax-stationary strategies can be computed without knowing v .) In order to identify v it is enough to compare these payoffs as follows. Take any two distinct payoff vectors $(u^s(\vec{x}_l, \vec{y}_l))_{s \in S}$ and $(u^s(\vec{x}_m, \vec{y}_m))_{s \in S}$ corresponding to labels l and m , and assume that for some state k , $u^k(\vec{x}_l, \vec{y}_l) > u^k(\vec{x}_m, \vec{y}_m)$. Next calculate the payoff in Γ^k when player 1 uses \vec{x}_l and player 2 uses \vec{y}_m . If $u^k(\vec{x}_l, \vec{y}_m) > u^k(\vec{x}_m, \vec{y}_m)$ then we conclude that \vec{y}_m does not guarantee that player 1 gets a payoff less or equal $u^k(\vec{x}_m, \vec{y}_m)$, which means that $u^k(\vec{x}_m, \vec{y}_m)$ is not the value of Γ^k . If $u^k(\vec{x}_l, \vec{y}_m) < u^k(\vec{x}_l, \vec{y}_l)$ then we conclude that \vec{x}_l does not guarantee that player 1 gets a payoff of at least $u^k(\vec{x}_l, \vec{y}_l)$, which means that $u^k(\vec{x}_l, \vec{y}_l)$ is not the value of Γ^k . Since at least one of the above inequalities must hold, we conclude that at least one of the above vectors of payoffs is not the value of Γ . Since Theorem 1 guarantees that there is one labeling l^* such that $(u^s(\vec{x}_{l^*}, \vec{y}_{l^*}))_{s \in S}$ is the value of Γ , after at most $2^K - 1$

⁶We should note, though, that for binary matches the above-mentioned invariance does hold. That is, if Γ'_λ is a binary discounted match that is obtained from the binary discounted match Γ_λ by slightly perturbing the non-zero entries of its point matrices such that for state k , $v_\lambda^{w(k)} > v_\lambda^{\ell(k)} \Leftrightarrow v_{\lambda'}^{w(k)} > v_{\lambda'}^{\ell(k)}$, then equilibrium behavior in state k is the same in both Γ_λ and Γ'_λ .

comparisons we identify the value of the match. In fact, as will be seen later, one needs only to consider payoffs $(u^s(\vec{x}_l, \vec{y}_l))_{s \in S}$ that are consistent with their labelings, namely $u^{w(s)}(\vec{x}_l, \vec{y}_l) \geq u^{\ell(s)}(\vec{x}_l, \vec{y}_l)$ for $s = 1, \dots, K$.

d) *Computation of the equilibrium minimax-stationary strategies.* Theorem 1 says not only that every quasi-binary match Γ has an equilibrium but also that it has an equilibrium which is relatively easy to compute. To do this, compute the value v of Γ along the lines described in item c) and then, if the equilibrium labeling has not yet been identified, use v to build the labeling mentioned in Theorem 1 which, as will be seen, can be done once v is known. The equilibrium strategies are the minimax-stationary strategies associated with this labeling.

e) *Alternative existence proof.* If we were interested in just showing the existence of equilibrium, we could do so by applying standard techniques based on the martingale property of the value and ignoring the concept of the value of a draw altogether. Specifically, it can be shown that a strategy pair (\vec{x}^*, \vec{y}^*) where for each state $k = 1, \dots, K$, x^{*k} and y^{*k} are optimal in $A^k(v)$ for players 1 and 2, respectively, constitutes an equilibrium as long as for each state k such that both its successors have the same value, x^{*k} and y^{*k} are completely mixed. These strategies, however, lack the natural interpretation we are seeking for, i.e., players do not identify one successor that they should try to move to and another that should be avoided. (The reason a player chooses a completely mixed action in a state whose successors have equal values is not necessarily that he is indifferent between moving to each one of them. Instead, he allows himself to act as if he did not know which state he should move to since by mixing he is making sure that were the current state to recur, he will eventually make the right choice.) Furthermore, in order to compute the above-mentioned equilibrium one needs to know the value of the game for which the alternative existence proof is of no help whatsoever. Neither does it highlight the fact that equilibrium play in each state is independent of the point matrices in the other states. As we have tried to argue in items a) – d) our result does not have these limitations.

f) *Necessity of the restriction to quasi-binary matches.* As will be clear from the proof, Theorem 1 can be immediately generalized to quasi-binary *recursive games*. In other words the restriction to games with only two absorbing states is unnecessary. However, the condition that $p_{ij}^{kk} < 1$ for all $i \in \mathcal{I}$ and $j \in \mathcal{J}$ cannot be dispensed with. Example 1 in Everett [2], summarized in the following matrix, illustrates this point.

$$P^1 : \begin{pmatrix} s_1 & 1 \\ 1 & -1 \end{pmatrix}$$

In this match, there is only one non-absorbing state, denoted by s_1 , and if players choose the first row and the first column, they remain in s_1 with probability 1. Following the usual practice, the payoffs 1 and -1 represent the transition to the corresponding absorbing states. As Everett shows, the value of Γ is 1 but player 1 cannot guarantee this payoff. Specifically, while player 1 can obtain a payoff as close to 1 as he wishes by choosing the mixed action $(1 - \varepsilon, \varepsilon)$ at every stage, he cannot guarantee a payoff of 1 since, for every one of his strategies, player 2 has a reply that yields a payoff less than 1.

Neither can the restriction to no more than two proper successors per state be relaxed, as the following two-state version of Everett's example demonstrates.

$$P^1 : \begin{pmatrix} s_2 & 1 \\ 1 & -1 \end{pmatrix} \quad P^2 : \begin{pmatrix} s_1 & 1 \\ 1 & -1 \end{pmatrix}$$

This match is obtained from the previous one by cloning the only non-absorbing state and amending the point matrices so that when players choose the first row and the first column, there is a transition from one state to its clone. Therefore, this match does not have an equilibrium. Note, however, that although the probability of remaining in the current state is 0 for all action pairs, both states have three proper successors.

4.2 The natural labeling

Before we prove the theorem we will construct an algorithm that labels the proper successors of the states. We will later show that any pair of minimax-stationary strategies with respect to this labeling is an equilibrium of Γ . If we were interested in showing only the existence of equilibrium, we could dispense with the construction of the natural labeling altogether and resort to standard arguments in stochastic games. However, the identification of this labeling allows us to build an equilibrium with a natural interpretation in terms of the value of a draw.

The idea of the labeling is as follows.⁷ Consider a state s and let s_1 and s_2 be its proper successors. If these successors have different values, then the one with the highest value will be labeled $w(s)$ and the one with the lowest value will be labeled $\ell(s)$. However, when they have the same value the choice of labels is not obvious and must be made carefully. There are three cases to consider. If $v(s_1) = v(s_2) > 0$, the state denoted by $w(s)$ will be a proper successor from which player 1 can guarantee a positive probability of winning the match by following a path of states, not including s , with non-decreasing values. If $v(s_1) = v(s_2) < 0$, the state denoted by $\ell(s)$ will be a proper successor from which player 2 can guarantee a positive probability of winning the match by following a path of states, not including s , with non-increasing values. Finally, if $v(s_1) = v(s_2) = 0$, any labeling of s 's successors will do. We next define a partition of the set of states that will allow us to identify the above-described $w(s)$ and $\ell(s)$.

Let (v^1, \dots, v^K) be the value of Γ and extend it so that $v^0 = 1$ and $v^{K+1} = -1$. Let $S^+ = \{k \in S : v^k > 0\}$ and $S^- = \{k \in S : v^k < 0\}$. For any $k \in S$, and $S' \subseteq S$, we write $k \rightarrow S'$ if for all $j \in \mathcal{J}$ there exists $i \in \mathcal{I}$ such that $\sum_{k' \in S'} p_{ij}^{kk'} > 0$. In other words, $k \rightarrow S'$ if player 2 cannot prevent a transition from k to some state in S' . Similarly, we write $k \bar{\rightarrow} S'$ if for all $i \in \mathcal{I}$ there exists $j \in \mathcal{J}$ such that $\sum_{k' \in S'} p_{ij}^{kk'} > 0$. In other

⁷There is an alternative way of building an equilibrium labeling, according to which $w(s)$ is chosen to be a successor of s with the highest discounted value for a sufficiently high discount factor. Figuring out, however, the appropriate discount factor and the associated discounted value is not an easy task.

words, $k \bar{\rightarrow} S'$ if player 1 cannot prevent a transition from k to some state in S' .

We now iteratively classify the elements of S^+ into disjoint subsets. Let $S_0^+ = \{0\}$ and for $n = 0, 1, 2, \dots$ let

$$S_{n+1}^+ = \{s \in S^+ \setminus \cup_{\nu=0}^n S_\nu^+ : v^s \geq v^{s'} \text{ for all } s' \in S^+ \setminus \cup_{\nu=0}^n S_\nu^+ \text{ and } s \rightarrow \cup_{\nu=0}^n S_\nu^+\}.$$

The set S_{n+1}^+ contains the states with maximum value among those not yet classified from which player 1 can guarantee a positive probability of a transition to a state that has already been classified.

Note that as long as $S^+ \setminus \cup_{\nu=0}^n S_\nu^+$ is not empty, S_{n+1}^+ is not empty either. Indeed, if there was no state in $\arg \max\{v^s : s \in S^+ \setminus \cup_{\nu=0}^n S_\nu^+\}$ with $s \rightarrow \cup_{\nu=0}^n S_\nu^+$, there would be a strategy for player 2 that guarantees that if play were to leave the set $\arg \max\{v^s : s \in S^+ \setminus \cup_{\nu=0}^n S_\nu^+\}$, it would do so through a state s' with $v^{s'} < v^s$. This would contradict the fact that player 1 can guarantee a payoff as close to v^s as he wishes.

Similarly, we iteratively classify the states in S^- into disjoint subsets as follows: we set $S_0^- = \{K+1\}$ and for $m = 0, 1, 2, \dots$ we let

$$S_{m+1}^- = \{s \in S^- \setminus \cup_{\nu=0}^m S_\nu^- : v^s \leq v^{s'} \text{ for all } s' \in S^- \setminus \cup_{\nu=0}^m S_\nu^- \text{ and } s \bar{\rightarrow} \cup_{\nu=0}^m S_\nu^-\}.$$

For $s \in S^+$, we denote by $n(s)$ the stage at which it was classified, namely, the index n such that $s \in S_n^+$. Similarly, for $s \in S^-$, we denote by $m(s)$ the index m such that $s \in S_m^-$. We can now proceed to label the proper successors of the states in $\{1, \dots, K\}$. Consider a state $s \in \{1, \dots, K\}$. Let s_1, s_2 be its two proper successors. Then we denote

$$w(s) = \begin{cases} s_1 & \text{if } v^{s_1} > v^{s_2} \\ s_1 & \text{if } v^{s_1} = v^{s_2} > 0 \text{ and } n(s_1) \leq n(s_2) \\ s_1 & \text{if } v^{s_1} = v^{s_2} < 0 \text{ and } m(s_1) > m(s_2) \\ s_2 & \text{otherwise} \end{cases} \quad (3)$$

and denote by $\ell(s)$ the other successor. In words, if s 's two successors have different

value, $w(s)$ is the one with the highest value. If they have the same value and it is positive, $w(s)$ is chosen to be the successor that has been classified earlier. If they have the same value and it is negative, $w(s)$ is chosen to be the successor that has been classified later. If both successors have a value of 0, any of them can be chosen to be $w(s)$. We call any labeling built according to the above procedure *a natural labeling*.⁸

Note that if $v^s > 0$, then $n(w(s)) < n(s)$. Indeed, one of s 's successors must have been classified before s (otherwise s could not have been classified), and by definition, $w(s)$ is the proper successor of s 's that has been classified the earliest.

For any state s , the natural labeling designates as $w(s)$ the proper successor with the highest value and in the case where both successors have equal value, it chooses one according to a specific tie-breaking rule. The following example shows that the outcome of this tie-breaking rule cannot be ignored.

Consider a simple one-person decision problem with three states, $S = \{s_0, s_1, s_2\}$. State s_0 is the only absorbing state and if reached, the player wins. The match is characterized by the following matrices in which again, in accordance with the usual practice, a transition to the absorbing state is denoted by the corresponding payoff:

$$P^1 : \begin{pmatrix} s_2 \\ 1 \end{pmatrix} \quad P^2 : \begin{pmatrix} s_1 \\ 1 \end{pmatrix}$$

It is clear that the value of all states is 1; the player can guarantee a win by choosing his second action in s_1 and s_2 . There are two possible ways to break the tie of the successors of both s_1 and s_2 . The natural labeling sets $w(s_1) = w(s_2) = s_0$, and indeed the minimax-stationary strategies with respect to this labeling constitutes an equilibrium. If we ignore the tie-breaking rule and set $\ell(s_1) = \ell(s_2) = s_0$, we obtain minimax-stationary strategies that lead to an infinite cycle and thus to a payoff of 0.

⁸There may be more than one natural labeling. For our analysis, any of them will do.

4.3 Proof of Theorem 1

We now show that any minimax-stationary strategy pair with respect to the natural labeling constitutes an equilibrium of Γ .

Let (\vec{x}^*, \vec{y}^*) be a minimax-stationary strategy pair with respect to the natural labeling. In order to show that it is an equilibrium of Γ^k we will show that \vec{x}^* guarantees a payoff of at least v^k for player 1 in Γ^k . The fact that \vec{y}^* guarantees that player 1 gets a payoff of at most v^k in Γ^k is analogous and is left to the reader. Finding a strategy $\psi^* \in Y$ that minimizes $u^k(\vec{x}^*, \cdot)$ is a Markov decision problem with the expected total reward criterion. Consequently, it has a stationary solution (see Puterman [8], Theorem 7.1.9). Therefore, it is enough to show that

$$u^k(\vec{x}^*, \vec{y}) \geq v^k \quad k = 1, \dots, K$$

for all stationary strategies \vec{y} of player 2. Let $\vec{y} = (y^0, \dots, y^{K+1})$ be a stationary strategy for player 2. The fact that \vec{x}^* guarantees e^k in the point game $P^k(e^k)$ for $k = 1, \dots, K$ implies that

$$\sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} x_i^{*k} y_j^k (p_{ij}^{kw(k)} + p_{ij}^{kk} e^k) \geq e^k \quad k = 1, \dots, K.$$

Let $M(\vec{x}^*, \vec{y}) = (\mu^{kk'}(\vec{x}^*, \vec{y}) | k, k' \in S)$ be the Markov transition matrix induced by the strategy pair (\vec{x}^*, \vec{y}) . Using equation (2), the above inequality can be written as

$$\mu^{kw(k)}(\vec{x}^*, \vec{y}) + \mu^{kk}(\vec{x}^*, \vec{y}) e^k \geq e^k \quad k = 1, \dots, K. \quad (4)$$

It follows that

$$\mu^{kw(k)}(\vec{x}^*, \vec{y}) v^{w(k)} + \mu^{kk}(\vec{x}^*, \vec{y}) v^k + \mu^{k\ell(k)}(\vec{x}^*, \vec{y}) v^{\ell(k)} \geq v^k \quad k = 1, \dots, K. \quad (5)$$

To see this, let $k \in \{1, \dots, K\}$. The natural labeling ensures that $v^{w(k)} \geq v^{\ell(k)}$. If $v^{w(k)} = v^{\ell(k)}$, inequality (5) is trivially satisfied since in this case, by Observation 1,

$v^{w(k)} = v^k = v^{\ell(k)}$. And if $v^{w(k)} > v^{\ell(k)}$, inequality (5) is obtained by multiplying (4) by $v^{w(k)} - v^{\ell(k)}$, adding $v^{\ell(k)}$ to both sides and applying Proposition 2. Taking into account that k has no successors except for $w(k)$, k and $\ell(k)$, we can rewrite inequality (5) as

$$\mu^{k0}(\vec{x}^*, \vec{y}) + \sum_{s=1}^K \mu^{ks}(\vec{x}^*, \vec{y}) v^s - \mu^{kK+1}(\vec{x}^*, \vec{y}) \geq v^k \quad k = 1, \dots, K.$$

Denoting $v = (v^0, v^1, \dots, v^{K+1})'$, we can rewrite the above inequality in matrix notation as

$$M(\vec{x}^*, \vec{y}) \cdot v \geq v.$$

Iterating, we obtain that $M^t(\vec{x}^*, \vec{y}) \cdot v \geq v$ for all t . In other words, for each $k = 1, \dots, K$, we have that

$$\mu_t^{k0}(\vec{x}^*, \vec{y}) + \sum_{s=1}^K \mu_t^{ks}(\vec{x}^*, \vec{y}) v^s - \mu_t^{kK+1}(\vec{x}^*, \vec{y}) \geq v^k \quad \text{for all } t.$$

Since $u^k(\vec{x}^*, \vec{y}) = \mu_\infty^{k0}(\vec{x}^*, \vec{y}) - \mu_\infty^{kK+1}(\vec{x}^*, \vec{y})$, in order to show that $u^k(\vec{x}^*, \vec{y}) \geq v^k$ it is enough to show that $\limsup_{t \rightarrow \infty} \sum_{s=1}^K \mu_t^{ks}(\vec{x}^*, \vec{y}) v^s \leq 0$. And to prove this it is enough to show that for all states s with $v^s > 0$, except for $s = 0$, $\lim_{t \rightarrow \infty} \mu_t^{ks}(\vec{x}^*, \vec{y}) = 0$. The Markov matrix $M(\vec{x}^*, \vec{y})$ induces a partition of S into recurrent classes and possibly a transient set.⁹ We will end the proof by showing that all states s with positive value, except for state 0, are transient states and thus $\lim_{t \rightarrow \infty} \mu_t^{ks}(\vec{x}^*, \vec{y}) = 0$.

Let C be a recurrent class different from $\{0\}$. We will show that all states in C have non-positive value. Let $s \in C$ and assume by contradiction that $v^s > 0$. Without loss of generality assume that $n(s) \leq n(s')$ for all $s' \in C$ such that $v^{s'} > 0$. Since $v^{w(s)} \geq v^s > 0$ and since $n(w(s)) < n(s)$, we have that $w(s) \notin C$. This means that $\mu^{sw(s)}(\vec{x}^*, \vec{y}) = 0$. By equation (4), and since $\mu^{ss}(\vec{x}^*, \vec{y}) < 1$, we obtain that $e^s = 0$. Namely, player 2 can prevent a transition from s to $w(s)$. That is, we must have that $s \not\rightarrow \{w(s)\}$. But since

⁹A set C is a *recurrent* class if $\sum_{k' \in C} \mu^{kk'}(\vec{x}^*, \vec{y}) = 1$ for all $k \in C$ and no proper subset of C has this property. A state is *transient* if there is a positive probability of leaving and never returning.

$s \rightarrow \cup_{\nu=0}^{n(s)-1} S_{\nu}^{+}$ we must have that $\ell(s) \in \cup_{\nu=0}^{n(s)-1} S_{\nu}^{+}$. This implies that $n(\ell(s)) < n(s)$, which in turn implies that $v^{\ell(s)} \geq v^s$. By our choice of s , this means that $\ell(s) \notin C$. We have obtained that neither $\ell(s)$ nor $w(s)$ is in C , which contradicts the fact that $\mu^{ss}(\vec{x}^*, \vec{y}) < 1$. \square

4.4 Monotonicity and a partial converse

As mentioned before, it is not necessarily so that every stationary equilibrium of a quasi-binary match is a minimax-stationary strategy pair with respect to some labeling. To see this, consider again the example that appears at the end of section 4.2. Consider the stationary strategy according to which player 1 chooses his two actions with equal probabilities in both states. It can be checked that, independent of the initial state, this strategy guarantees a payoff of 1; hence it is an equilibrium strategy. However, it is not a minimax-stationary strategy with respect to any labeling since no matter how the successors of s_1 are labeled, the corresponding minimax-stationary strategy will never prescribe that player 1 should mix between his actions in s_1 .

We now present a partial converse of Theorem 1. It says that when for every state both of its proper successors have different values, any stationary strategy equilibrium of Γ is a minimax-stationary strategy pair with respect to the natural labeling.

Let Γ be a quasi-binary match and let (v^1, \dots, v^K) be its value. Extend it so that $v^0 = 1$ and $v^{K+1} = -1$ and let $v = (v^0, \dots, v^{K+1})$. Note that if the proper successors of a given state k have different values, then $v^{w(k)} > v^{\ell(k)}$ for any natural labeling (w, ℓ) . We say that Γ satisfies *monotonicity* if for every state both its proper successors have different values. Notice that if Γ satisfies monotonicity, there is a unique natural labeling.

Theorem 2 Let Γ be a quasi-binary match that satisfies monotonicity. A pair of stationary strategies is an equilibrium of Γ only if it is a pair of minimax-stationary strategies with respect to the natural labeling.

Proof : Let (\bar{x}^*, \bar{y}^*) be a stationary equilibrium of Γ and let (w, ℓ) be the natural labeling. Let $k \in 1, \dots, K$. Since $v^{w(k)} > v^{\ell(k)}$, by Proposition 2

$$e^k = \frac{v^k - v^{\ell(k)}}{v^{w(k)} - v^{\ell(k)}}. \quad (6)$$

We need to show that x^{*k} guarantees that player 1 gets a payoff of at least e^k in $P^k(e^k)$ and that y^{*k} guarantees that player 1 gets a payoff of at most e^k in $P^k(e^k)$.

Since (\bar{x}^*, \bar{y}^*) is an equilibrium of Γ^k ,

$$v^k = u^k(\bar{x}^*, \bar{y}^*) \geq u^k(\chi, \bar{y}^*) \quad \text{for all } \chi \in X. \quad (7)$$

Since \bar{y}^* is a stationary strategy, the problem of finding a strategy for player 1 that maximizes $u^k(\cdot, \bar{y}^*)$ is a Markov decision problem (with the expected total-reward criterion). Equation (7) says that \bar{x}^* is one of its solutions and that it attains v^k . Therefore (see Puterman [8], Chapter 7),

$$v = \max_{\bar{x} \in \bar{X}} M(\bar{x}, \bar{y}^*)v \quad (8)$$

where $M(\bar{x}, \bar{y}^*)$ is the Markov matrix induced by the stationary strategy pair (\bar{x}, \bar{y}^*) . This means that, using equation (2), for every $k = 1, \dots, K$,

$$\begin{aligned} v^k &= \max_{\bar{x} \in \bar{X}} \sum_{k'=0}^{K+1} \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} x_i^k y_j^{*k} p_{ij}^{kk'} v^{k'} \\ &= \max_{\bar{x} \in \bar{X}} \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} x_i^k y_j^{*k} \sum_{k'=0}^{K+1} p_{ij}^{kk'} v^{k'} \\ &= \max_{\bar{x} \in \bar{X}} \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} x_i^k y_j^{*k} (p_{ij}^{kw(k)} v^{w(k)} + p_{ij}^{kk} v^k + p_{ij}^{k\ell(k)} v^{\ell(k)}). \end{aligned}$$

Subtracting $v^{\ell(k)}$ from both sides and then dividing the result by $v^{w(k)} - v^{\ell(k)}$ (which can

be done since this difference is positive) using equation (6) we find that

$$e^k = \max_{\vec{x} \in \bar{X}} \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} x_i^k y_j^{*k} (p_{ij}^{kw(k)} + p_{ij}^{kk} e^k).$$

This shows that y^{*k} guarantees that player 1 gets at most e^k in $P^k(e^k)$.

A similar argument shows that x^{*k} guarantees that player 1 gets at least e^k in $P^k(e^k)$. \square

We end this section by discussing the difference between quasi-binary matches and Walker, Wooders and Amir's [15] binary Markov games. As mentioned in the introduction, one small difference is that while in the former, states may have three successors, in the latter they have only two. The main difference, however, is that whereas in the former a player's payoff consists of the difference between the probability of winning and the probability of losing the match, in the latter it is simply the probability of winning the match. Formally, in a binary Markov game player 1's payoff from a strategy pair (χ, ψ) is $W_1^k(\chi, \psi) = \mu^{k0}(\chi, \psi)$, and player 2's payoff is $W_2^k(\chi, \psi) = \mu^{kK+1}(\chi, \psi)$.

To illustrate this difference, consider the following game form. There are four non-absorbing states, denoted $s_1, s_2, s_3,$ and s_4 . States s_1 and s_4 are trivial states that if reached, a lottery is obtained. Specifically, in s_1 , player 1 wins the match with probability 1/4 and loses the match with probability 3/4, and similarly in s_4 , player 1 wins the match with probability 3/4 and loses the match with probability 1/4. In state s_2 both players can guarantee a transition to state s_3 , and in state s_3 both players can force a transition to s_2 .¹⁰ The point matrices are as follows:

¹⁰This game resembles a chess position in which players can either transition to an endgame with unfavorable odds or keep playing safe. It is usually the case that players prefer the second option and thus the game ends in a draw.

$$P^1 : \left(\begin{array}{c} \bullet \\ 1/4 \swarrow \quad \searrow 3/4 \\ 1 \quad -1 \end{array} \right) \quad P^2 : \begin{pmatrix} s_1 & s_3 \\ s_3 & s_3 \end{pmatrix} \quad P^3 : \begin{pmatrix} s_4 & s_2 \\ s_2 & s_2 \end{pmatrix} \quad P^4 : \left(\begin{array}{c} \bullet \\ 3/4 \swarrow \quad \searrow 1/4 \\ 1 \quad -1 \end{array} \right)$$

If we endow players 1 and 2 with the payoff functions $u^k(\phi, \psi)$ and $-u^k(\phi, \psi)$, respectively, we obtain a (quasi) binary match whose unique equilibrium outcome results in an infinite play that orbits around states s_2 and s_3 . Indeed, since at state s_2 each player can force a transition to s_3 , no matter how we label its successors all minimax stationary strategies will result in a transition to state s_3 . Similarly, since at state s_3 each player can force a transition to s_2 , no matter how we label its successors all minimax stationary strategies will result in such a transition. The values of the states are $v_1 = -1/2$, $v_2 = v_3 = 0$ and $v_4 = 1/2$, and therefore the natural labeling is $w(s_2) = s_3$, $\ell(s_2) = s_1$, $w(s_3) = s_4$, $\ell(s_3) = s_2$. According to Theorem 1, all pairs of minimax stationary strategies with respect to this labeling are an equilibrium. And according to Theorem 2 they are the only stationary equilibria of the game.

Alternatively, if we endow the players with the payoff functions W_1^k and W_2^k we obtain a binary Markov game as defined by Walker, Wooders and Amir [15]. Thus defined, this game is not constant-sum and as a consequence it has three stationary equilibrium outcomes with three different payoffs.

One equilibrium consists of players 1 and 2 choosing the second row and column, respectively, both in states s_2 and s_3 . This equilibrium leads to a never-ending cycle orbiting around s_2 and s_3 , yielding a winning probability of 0 for both players. Notice that in this equilibrium one of the players uses a weakly dominated action.

Besides this equilibrium, there are two additional kinds of equilibria, both of which lead to finite play. In one of them, player 1 chooses row 2 in state s_2 , and in state s_3 players 1 and 2 choose the first row and first column, respectively, with positive probability. In these equilibria, if the initial state is s_2 or s_3 player 1 wins the game with probability $3/4$ and player 2 wins with probability $1/4$. In the other kind of equilibria,

player 2 chooses column 2 in state s_3 , and in state s_2 players 1 and 2 choose the first row and first column, respectively, with positive probability. In these equilibria, if the initial state is s_2 or s_3 player 1 wins the game with probability $1/4$ and player 2 wins with probability $3/4$.

As we mentioned above, no matter how we label the states' successors, all minimax stationary strategy pairs will result in a never ending cycle orbiting around s_2 and s_3 , and in a 0 payoff for both players. It turns out that there are infinitely many such pairs and except for one, none of them constitutes an equilibrium. Specifically, as long as none of the players chooses a weakly dominated action with probability one, one of them will choose his first action with positive probability in state s_2 and thus the other player can profitably deviate by choosing his own first action and causing a transition to state s_1 , with the resulting a positive payoff.

This is a robust example where the non-zero sum nature of the binary Markov games essentially transforms the match into a coordination game. As a result, no matter the labeling, almost none of the pairs of minimax-stationary strategies is an equilibrium of the binary Markov game. Walker, Wooders and Amir's [15] equilibrium theorem states that if a certain labeling satisfies a monotonicity condition, then any pair of minimax-stationary strategies with respect to it constitutes an equilibrium. Their monotonicity condition's role is to restrict the class of binary Markov games to exclude examples like this one. In contrast, by modeling matches as standard recursive games, our result requires no such restriction.

References

- [1] APESTEGUIA, J., AND PALACIOS-HUERTA, I. Psychological pressure in competitive environments: Evidence from a randomized natural experiment. *The American Economic Review* 100, 5 (2010), 2548–2564.

- [2] EVERETT, H. Recursive games. *Contributions to the Theory of Games* 3, 39 (1957), 47–78.
- [3] FLESCH, J., THUIJSMAN, F., AND VRIEZE, O. J. Recursive repeated games with absorbing states. *Mathematics of Operations Research* 21, 4 (1996), 1016–1022.
- [4] GAURIOT, R., PAGE, L., AND WOODERS, J. Nash at Wimbledon: Evidence from half a million serves. Tech. rep., QuBE, 2016.
- [5] GONZALEZ-DÍAZ, J., AND PALACIOS-HUERTA, I. Cognitive performance in competitive environments: Evidence from a natural experiment. *Journal of Public Economics* 139 (2016), 40–52.
- [6] MERTENS, J.-F., AND NEYMAN, A. Stochastic games. *International Journal of Game Theory* 10, 2 (1981), 53–66.
- [7] PALACIOS-HUERTA, I. Professionals play minimax. *The Review of Economic Studies* 70, 2 (2003), 395–415.
- [8] PUTERMAN, M. L. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 1994.
- [9] SCHWALBE, U., AND WALKER, P. Zermelo and the early history of game theory. *Games and economic behavior* 34, 1 (2001), 123–137.
- [10] SHAPLEY, L. S. Stochastic games. *Proceedings of the National Academy of Sciences of the United States of America* 39, 10 (1953), 1095–1100.
- [11] THUIJSMAN, F. Optimality and equilibria in stochastic games. *CWI Tracts* 82 (1992), 1–107.
- [12] THUIJSMAN, F., AND VRIEZE, K. Note on recursive games. In *Game Theory and Economic Applications, Lecture Notes in Economics and Mathematical Systems* 389 (1992), Springer, pp. 133–145.

- [13] VIEILLE, N. Two-player stochastic games ii: The case of recursive games. *Israel Journal of Mathematics* 119, 1 (2000), 93–126.
- [14] WALKER, M., AND WOODERS, J. Minimax play at Wimbledon. *American Economic Review* (2001), 1521–1538.
- [15] WALKER, M., WOODERS, J., AND AMIR, R. Equilibrium play in matches: Binary markov games. *Games and Economic Behavior* 71, 2 (2011), 487–502.
- [16] ZERMELO, E. Über eine anwendung der mengenlehre auf die theorie des schachspiels. In *Proceedings of the fifth international congress of mathematicians* (1913), vol. 2, II, Cambridge UP, Cambridge, pp. 501–504.