The value of a draw in quasi-binary matches*

Casilda Lasso de la Vega †

University of the Basque Country

Oscar Volij ‡

Ben-Gurion University of the Negev

June 30, 2016

Abstract

A match is a recursive zero-sum game with three possible outcomes: player 1 wins, player 2 wins or there is a draw. Play proceeds by steps from state to state. In each state players play a "point game" and move to the next state according to transition probabilities jointly determined by their actions. We focus on quasi-binary matches which are those whose point games also have three possible outcomes: player 1 scores the point, player 2 scores the point, or the point is drawn (something that happens with probability less than 1) in which case the point game is repeated. 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 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. 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.[1] 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 are recursive game as defined by Everett [2]. Recursive games are a special case of stochastic games, which were earlier introduced by Shapley [8]. Matches have been the object of several empirical studies. For instance, both Walker and Wooders [10], using data on tennis, and Palacios-Huerta [6] using data on penalty kicks, show that players' behavior is broadly consistent with the minimax hypothesis. On the other hand, Apesteguia and Palacios-Huerta [1] observe a first-kicker advantage in penalty shootouts and Gonzalez-Díaz and Palacios-Huerta [4] find a similar anomaly in chess matches. This last paper also offers a brief theoretical analysis of a particular finite chess match.

Walker, Wooders and Amir [11] analyzed binary games and showed that under a

[1] In fact, the 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 so far also had 48 kicks.

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 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.

Strictly speaking however, Walker, Wooders and Amir's [11] result is proved for matches in which never-ending play is defined to be the worst outcome for both players, a feature that renders their matches non-zero sum games. In this paper we extend their analysis to a different class of games, which we call *quasi-binary* matches, and obtain an arguably stronger result. Quasi-binary matches differ from the ones considered in Walker, Wooders and Amir [11] in two aspects. First, they *are* zero-sum games: the match payoff function awards 1 to the winner, -1 to the loser, and in the event of an infinite play 0 to both players. And second, they are matches 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 it is repeated. Like in binary games, from any state play may move to one of at most two states. Unlike binary games, play may also stay in the current state for some time.

In this paper 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 in which players' strategies prescribe minimax play 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 thus equilibrium behavior at that state is independent of the structure of the match.

Before we describe these equilibria, notice 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.

As mentioned above, there are two ways of labeling the states to which the players can move from each state. 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, every stationary equilibrium of the match prescribes minimax play in these zero-sum games.

To illustrate the main result, consider 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.[2] 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 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

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

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, since this match satisfies a simple monotonicity condition, 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 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.

This paper generalizes the foregoing example for all quasi-binary matches. Specifically, denoting P^k the matrix of probabilities that govern the outcomes of the point played at state k, we can find a value of the draw e^k and build a matrix $P^k(e^k)$ which is obtained from P^k by first interpreting one of the outcomes as winning the point and the other as losing it, and by evaluating a draw as worth e^k of a point. Our main result says that for any quasi-binary game, choosing minimax mixtures of the point game $P^k(e^k)$ in state k constitutes a stationary equilibrium. Furthermore, when a simple monotonicity condition is satisfied, all the stationary equilibria of the match are of this type.

The paper is organized as follows. Section 2 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 main result.

2 Matches

2.1 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 $S = \{0, 1, \ldots, K+1\}$. 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 $1, \ldots, I_k$ and $1, \ldots, J_k$, 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, $\mathrm{val}(A)$ 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 $\mathrm{val}(A)$. 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 $\mathrm{val}(A)$. Recall that for $A = (a_{ij} | i \in \mathcal{I}, j \in \mathcal{J})$ and $B = (b_{ij} | i \in \mathcal{I}, j \in \mathcal{J})$, $|\mathrm{val}(A) - \mathrm{val}(B)| \leq \max_{ij} |a_{ij} - b_{ij}|$ 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 $\mathrm{val}(B) = \alpha\, \mathrm{val}(A) + \beta$.

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

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

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 $k = 1, \ldots, K$, 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.

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 $t = 0, 1, 2, \ldots$ is denoted by $H_t = S \times (\mathcal{I} \times \mathcal{J} \times S)^t$. A typical history of length t is $h_t = (s_0, (i_1, j_1, s_1), \ldots, (i_t, j_t, s_t)) \in H_t$. Here, the initial state is $s_0 \in S$ and at stage $\tau = 1, \ldots t$, 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 $\chi : H \to \Delta_{\mathcal{I}}$ that prescribes a mixed action $\chi(h_t) = (\chi_1(h_t), \ldots, \chi_I(h_t))$ to be used by player 1 after every finite history h_t. Similarly, a strategy for player 2 is a map $\psi : H \to \Delta_{\mathcal{J}}$ that prescribes a mixed action $\psi(h_t) = (\psi_1(h_t), \ldots, \psi_J(h_t))$ to be used by player 2 after every finite history h_t. *Stationary strategies* are strategies that depend only on the current state. Thus, a stationary strategy for player 1 can be represented by a vector $\vec{x} = (x^0, \ldots, x^{K+1})$, where for each $k \in S$, $x^k = (x_1^k, \ldots, x_I^k)$ is a mixed action for player 1. Similarly, a stationary strategy for player 2 is a vector $\vec{y} = (y^0, \ldots, y^{K+1})$ 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 of length 0, $h_0 \in H_0$,

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

And for histories of length $t = 1, 2, \ldots$ 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 $t = 1, 2, \ldots$, the current state is k' is given by

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

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).

As mentioned earlier, when state 0 is reached, player 1 wins and gets a payoff of 1 from player 2 and if state $K+1$ is reached, player 1 loses and pays 1 to player 2. It is not necessarily true, however, that any pair of strategies leads to one of these two states with probability 1. In the case that there is no winner we specify the players' payoffs to be 0.

We can now define the 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 $u^k : X \times Y \to [-1, 1]$ 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 $u^0(\chi, \psi) \equiv 1$ and $u^{K+1}(\chi, \psi) \equiv -1$. We denote by Γ the collection of matches $\{\Gamma^k : k = 1, \ldots, K\}$ and remark that Γ is fully determined by the set of states S and by the set of point matrices $(P^k)_{k=1}^K$.

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 $k = 1, \ldots, K$ we say that $(v^1, \ldots, v^K)$ is the value of Γ. If $\chi_\varepsilon \in X$ is such that $u^k(\chi_\varepsilon, \psi) \geq v^k - \varepsilon$ for $\varepsilon > 0$ and for all $\psi \in Y$, we say that χ_ε is ε-optimal for player 1 in Γ^k. Similarly, if $\psi_\varepsilon \in Y$ is such that $u^k(\chi, \psi_\varepsilon) \leq v^k + \varepsilon$ for $\varepsilon > 0$ and for all $\chi \in X$, we say that ψ_ε is ε-optimal for player 2 in Γ^k. 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 $k \in \{1, \ldots, K\}$.

As mentioned before, Γ is a recursive game. Everett [2] shows that recursive games have a value, and Mertens and Neyman [5] prove more generally that when streams of payoffs are undiscounted all stochastic games with finite state and action spaces have a value. Further results in recursive games can be found in Flesch, Thuijsman and Vrieze [3] and in Vieille [9].

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 $\alpha = (\alpha^1, \ldots, \alpha^K) \in \mathbb{R}^K$ we can define the matrix game $A^k(\alpha)$ as follows:

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

As a direct application of Theorems 2, 3 and 6 of Everett [2] we have the following observation which plays a fundamental role in our analysis.

Observation 1 For $k = 1, \ldots, K$, Γ^k has a value v^k and this value satisfies $v^k = \mathrm{val}(A^k(v^1, \ldots, v^k))$. Furthermore, for every $\varepsilon > 0$ there exist stationary strategies $\vec{x}_\varepsilon \in \vec{X}$ and $\vec{y}_\varepsilon \in \vec{Y}$ that are ε-optimal for players 1 and 2, respectively, in Γ^k, $k = 1, \ldots, K$.

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

2.2 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 $M(\vec{x}, \vec{y}) = (\mu^{ss'}(\vec{x}, \vec{y}) | s, s' \in S)$ 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

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

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 $P^k = (p^k_{ij} | i \in \mathcal{I}; j \in \mathcal{J})$, for $k = 1, \ldots, K$. For each state k, define the set of its immediate successors, or simply *successors*, to be

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

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*. The set of k's proper successors is denoted by $\hat{S}(k)$.

Definition 1 A match is quasi-binary if for each state $k = 1, \ldots, K$ the number of its proper successors is exactly two, and $p^{kk}_{ij} < 1$ 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 $p^{kk}_{ij} = 1$ 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 3 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 $k = 1, \ldots, K$ 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^k_{ij} in the ijth entry by the corresponding expected payoff $p^{kw(k)}_{ij} + p^{kk}_{ij}\varepsilon$. Formally, for each $\varepsilon \in [0, 1]$ we define the matrix game $P^k(\varepsilon)$ by letting its ijth entry be $p^{kw(k)}_{ij} + p^{kk}_{ij}\varepsilon$, 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.[3] 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 $f^k : [0, 1] \to [0, 1]$ be the function defined by $f^k(\varepsilon) = \mathrm{val}(P^k(\varepsilon))$. 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 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.

---

[3] 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 $\{p^{kw(k)}_{ij} | i \in \mathcal{I}, j \in \mathcal{J}\}$ and with this amended definition, the ensuing analysis would remain valid.

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. □

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 is needed. In particular, no prior knowledge of the value of Γ is required. The next proposition, however, shows that when $v^{w(k)} > v^{\ell(k)}$, 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 $(v^1, \ldots, v^K)$ 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 $\epsilon^k = (v^k - v^{\ell(k)})/(v^{w(k)} - v^{\ell(k)})$. By Proposition 1, the value of the draw in state k is the unique fixed point of the function $f^k : [0,1] \to [0,1]$ 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 $v^k = \text{val}(A^k(v^1, \ldots, v^k))$ 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 ijth 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,

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

and consequently,

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

□

The foregoing proposition justifies calling e^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 ε-optimal play) a payoff close to $v^{w(k)}$, or $v^{\ell(k)}$, respectively. Therefore, since $v^{w(k)} > v^{\ell(k)}$, player 1 has a guaranteed expected payoff close to $v^{\ell(k)}$ and hence what is really at stake in state k is close to $v^{w(k)} - v^{\ell(k)}$. When the point is drawn, the players remain in state k, in which case player 1 gets an expected payoff close 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 e^k, the unique fixed point identified in Proposition 1, is precisely this proportion – hence its interpretation as the value of a draw.

In the next definition we identify 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 $k = 1, \ldots, K$ 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 $k = 1, \ldots, K$, (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 $k = 1, \ldots, K$

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

Proof : Since $\vec{x} = (x^0, \ldots, x^{K+1})$ and $\vec{y} = (y^0, \ldots, y^{K+1})$ constitute a pair of minimax-stationary strategies, for $k = 1, \ldots, K$, (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. □

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.

We will prove the theorem in two stages. We first identify in Section 4.2 a natural way of labeling the successors of each state, and we later show, in Section 4.3, that this labeling is the one mentioned in the theorem. Before that we discuss the result.

4.1 Discussion

a) *Interpretation of the result.* The labeling $w(k), \ell(k)$ identified in the theorem, along with the associated value of a draw e^k, allows the following interpretation of the possible transitions from state k. Moving to $w(k)$ is interpreted 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 adopts this interpretation and aim at maximizing their respective expected shares of the point at stake.

b) *Computation of the value of the match.* Theorem 1 allows us to compute the value v of Γ in a relatively 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 $2^K - 1$ comparisons we identify the value of the match. In fact, 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 $s = 1, \ldots, K$.

c) *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 of Γ along the lines described in item a) and then use it 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.

d) *Necessity of the restriction to quasi-binary matches.* For the purposes of Theorem 1, the condition on quasi-binary matches 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. The payoffs 1 and -1 represent the transition to the 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.

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

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.

The idea of the labeling is as follows. 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 $(v^1, \ldots, v^K)$ 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 $S^- = \{k \in S : v^k < 0\}$. Define a binary relation $\to$ on S^+ as follows: for $k \in S^+$, $k \to k'$ if k' is a proper successor of k with $v^{k'} \geq v^k$ and if for all $j \in \mathcal{J}$ there exists $i \in \mathcal{I}$ such that $p_{ij}^{kk'} > 0$. In other words, $k \to k'$ if k' has a value at least as large as the value of k and player 2 cannot prevent a transition from k to k'.

Similarly, define a binary relation $\bar{\to}$ on S^- as follows: for any $k \in S^-$, $k \bar{\to} k'$ if k' is a proper successor of k with $v^{k'} \leq v^k$ and if for all $i \in \mathcal{I}$ there exists $j \in \mathcal{J}$ such that $p_{ij}^{kk'} > 0$.

We now iteratively classify the elements of S^+ into disjoint subsets. Let $S_0^+ = \{0\}$. Also let $S_1^+ = \{s \in S^+ \setminus S_0^+ : s \to 0\}$ be the set of states with positive value from which player 1 can guarantee a positive probability of winning the match in one step. In general, define for $n = 1, 2, \ldots$

$$S_{n+1}^+ = \{s \in S^+ \setminus \cup_{\nu=0}^n S_\nu^+ : \text{either there exists } s' \in S_n^+ \text{ with } s \to s' \text{ or } \hat{S}(s) \subseteq \cup_{\nu=0}^n S_\nu^+\}.$$

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

Similarly, we iteratively classify the states in S^- into disjoint subsets as follows:

$$S^-_{m+1} = \{s \in S^- \setminus \cup_{\nu=0}^m S^-_\nu : \text{either there exists } s' \in S^-_m \text{ with } s \xrightarrow{-} s' \text{ or } \hat{S}(s) \subseteq \cup_{\nu=0}^m S^-_\nu\}.$$

Since the number of states in S^+ is finite, there must be an N such that $S^+_N \neq \emptyset$ and $S^+_{N+\nu} = \emptyset$ for all $\nu = 1, 2, \ldots$. Similarly, there must be an M such that $S^-_M \neq \emptyset$ and $S^-_{M+\nu} = \emptyset$ for all $\nu = 1, 2, \ldots$. The following claim, whose proof can be found in the Appendix, states that the subsets defined above form a partition of S^+ and of S^-, respectively.

Claim 1 The collection $\{S^+_0, \ldots, S^+_N\}$ forms a partition of S^+, and $\{S^-_0, \ldots, S^-_M\}$ forms a partition of S^-.

We can now proceed to label the proper successors of the states in $\{1, \ldots, K\}$. Consider first a state $s \in S^+$. By the previous claim, $s \in S^+_{n+1}$ for some n. Let s_1, s_2 be its two proper successors and assume without loss of generality that $v^{s_1} \geq v^{s_2}$. Then we denote

$$w(s) = \begin{cases} s_1 & \text{if } v^{s_1} > v^{s_2} \\ s_1 & \text{if } v^{s_1} = v^{s_2} \text{ and } s_1 \in \cup_{\nu=0}^n S^+_\nu \\ s_2 & \text{if } v^{s_1} = v^{s_2} \text{ and } s_1 \notin \cup_{\nu=0}^n S^+_\nu \end{cases} \tag{3}$$

and denote by $\ell(s)$ the other successor.

Simlarly, let $s \in S^-$. By the previous claim, $s \in S^-_{m+1}$ for some m. Let s_1, s_2 be its two proper successors and assume without loss of generality that $v^{s_1} \geq v^{s_2}$. Then we denote

$$\ell(s) = \begin{cases} s_2 & \text{if } v^{s_1} > v^{s_2} \\ s_2 & \text{if } v^{s_1} = v^{s_2} \text{ and } s_2 \in \cup_{\nu=0}^m S^-_\nu \\ s_1 & \text{if } v^{s_1} = v^{s_2} \text{ and } s_2 \notin \cup_{\nu=0}^m S^-_\nu \end{cases}$$

and denote by $w(s)$ the other successor.

Finally, let $v^s = 0$ and denote by s_1, s_2 its two proper successors where $v^{s_1} \geq v^{s_2}$.

Then, we label $w(s) = s_1$ and $\ell(s) = s_2$.

We call any labeling built according to the above procedure *a natural labeling.*[4] Notice that this labeling satisfies $v^{w(s)} \geq v^{\ell(s)}$ for all $s \in 1, \ldots, K$. The following example illustrates the construction of a natural labeling.

Example 1 Consider the following match. The set of states is $S = \{s_0, s_1, s_2, s_3, s_4\}$. States s_0 and s_4 are absorbing. If the former is reached, player 1 wins the match and if the latter is reached player 2 wins the match. The payoffs for player 1 in these two absorbing states are 1 and -1, respectively. The match is characterized by the following point matrices where instead of s_0 and s_4 we write the respective payoffs 1 and -1.

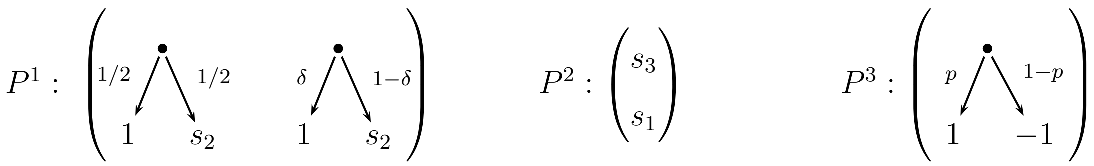

In state s_1 only player 2 has a non-trivial choice, in state s_2 only player 1 has a non-trivial choice, and in state s_3 there is a single action pair.

Case 1: $\delta = 0$ and $1/2 < p < 1$. In this case, player 1 can guarantee a payoff of at least $2p - 1 > 0$ in all Γ^k by choosing his first action in state 2. Similarly, player 2 can guarantee that player 1 gets no more than $2p - 1$ by choosing his second action in state 1. Therefore the value of Γ is given by $v^1 = v^2 = v^3 = 2p - 1 > 0$. Hence, $S^+ = \{s_0, s_1, s_2, s_3\}$ and $S^- = \{s_4\}$. Since the successors of s_1 are s_0 and s_2 and since $v^0 = 1 > 2p - 1 = v^2$, according to the natural labeling we have that $w(s_1) = s_0$ and $\ell(s_1) = s_2$. Similarly, $w(s_3) = s_0$ and $\ell(s_3) = s_4$. In order to complete the natural labeling, we build the partition of S^+ mentioned in Claim 1. By definition, $S_0^+ = \{s_0\}$. Although, s_0 is a proper successor of both s_1 and s_3, only $s_3 \in S_1^+$. Indeed, that $s_3 \in S_1^+$ is clear because $s_3 \to s_0$. That $s_1 \notin S_1^+$ follows from the fact that since $\delta = 0$, player 2 can prevent a transition to s_0 by choosing his second action in state s_1. Hence, $S_1^+ = \{s_3\}$. Since $s_2 \to s_3$, we have that $S_2^+ = \{s_2\}$. Finally, since $s_1 \to s_2$, we have

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

that $S_3^+ = \{s_1\}$. Therefore, by applying (3) we obtain that $w(s_2) = s_3$ and $\ell(s_2) = s_1$.
Case 2: $0 < \delta < 1/2$ and $1/2 < p < 1$. In this case player 1 can guarantee that in Γ^1 and in Γ^2 he wins the match by choosing his second action in state s_2. Consequently, the value of the match is given by $v^1 = v^2 = 1$ and $v^3 = 2p - 1 > 0$. As before, $w(s_3) = s_0$ and $\ell(s_3) = s_4$. Since the successors of s_2 are states s_1 and s_3 and since $v^1 > v^3$, we have that $w(s_2) = s_1$ and $\ell(s_2) = s_3$. Since $\delta > 0$, player 2 can no longer prevent a transition from state s_1 to state s_0, and consequently $S_0^+ = \{s_0\}$, $S_1^+ = \{s_1, s_3\}$, and $S_2^+ = \{s_2\}$. Applying (3) we obtain that the natural labeling is $w(s_1) = s_0$ and $\ell(s_1) = s_2$. Notice the labels of state s_2 are different depending on whether $\delta = 0$ or $\delta > 0$ which shows that a small change in the entries of the point matrix in one state can affect the natural labeling in other states.
Case 3: $\delta = 0$ and $0 < p < 1/2$. In this case the value of Γ is given by $v^1 = v^2 = 0$, and $v^3 = 2p - 1 < 0$. As can be checked, the resulting natural labeling is given by $w(s_1) = s_0$, $w(s_2) = s_1$, and $w(s_3) = s_0$, (and $\ell(s_1) = s_2$, $\ell(s_2) = s_3$, and $\ell(s_3) = s_4$).

We now illustrate Theorem 1 by computing the equilibrium minimax-stationary strategies of the matches described in the above example.

Example 1 (cont.)
Case 1: $\delta = 0$ and $1/2 < p < 1$. We have already shown that the natural labeling in this case is given by $w(s_1) = s_0$, $w(s_2) = s_3$, and $w(s_3) = s_0$, (and $\ell(s_1) = s_2$, $\ell(s_2) = s_1$, and $\ell(s_3) = s_4$). Since Γ is not just a quasi-binary match but also a binary game, the corresponding matrices $P^k(\varepsilon)$ with respect to this labeling are constant and are given by

$$P^1(\varepsilon) : \begin{pmatrix} 1/2 & 0 \end{pmatrix} \qquad P^2(\varepsilon) : \begin{pmatrix} 1 \\ 0 \end{pmatrix} \qquad P^3(\varepsilon) : \begin{pmatrix} p \end{pmatrix}$$

It can be checked that the associated minimax-stationary strategies dictate that player 1 chooses his first action in state s_2, and player 2 chooses his second action in state s_1. Consistent with Theorem 1 they constitute an equilibrium of Γ.

Case 2: $0 < \delta < 1/2$ and $1/2 < p < 1$. It can be checked that in this case the matrices $P^k(\varepsilon)$ become

$$P^1(\varepsilon) : \begin{pmatrix} 1/2 & \delta \end{pmatrix} \qquad P^2(\varepsilon) : \begin{pmatrix} 0 \\ 1 \end{pmatrix} \qquad P^3(\varepsilon) : \begin{pmatrix} p \end{pmatrix}$$

Therefore, the associated minimax-stationary strategies prescribe that player 1 chooses his second action in state s_2 and player 2 his first action in state s_1.

Case 3: $\delta = 0$ and $0 < p < 1/2$. In this case the associated minimax-stationary strategies prescribe that player 1 chooses his second action in state s_2 and player 2 chooses his second action is state s_1. Notice that these strategies lead to a never-ending cycle involving states s_1 and s_2, and consequently to a tie in the match.

Theorem 1 states that any pair of minimax-stationary strategies 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 natural labeling. Therefore, any modification in the structure of the match that involves neither a change in the point matrix P^k nor in the natural labeling, will leave the equilibrium behavior in state k unaffected. However, in some non-generic matches even a small change in the point matrix of a state different from k may drastically alter the equilibrium behavior in state k. To see this compare cases 1 and 2 in Example 1. We have seen that whether δ is positive or not affects the natural labeling, and as a consequence, affects the associated minimax-stationary strategies as well. Indeed, when $\delta = 0$ the minimax-stationary strategy of player 1 dictates that he chooses his first action while when $\delta > 0$ it prescribes his second action. Theorem 1 shows that this kind of interstate influence is possible only if the changes in the point games affect the natural labeling.

4.3 Proof of Theorem 1

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

Fix a natural labeling and let $(\vec{x}^*, \vec{y}^*)$ be a minimax-stationary strategy pair with respect to it. 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 [7], Theorem 7.1.9). Therefore, it is enough to show that

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

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

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

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

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

It follows that

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

To see this, let $k \in \{1, \ldots, K\}$. 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

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

Denoting $v = (v^0, v^1, \ldots, v^{K+1})'$, 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 $k = 1, \ldots, K$, 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 $\limsup_{t\to\infty} \sum_{s=1}^{K} \mu_t^{ks}(\vec{x}^*, \vec{y})\, v^s \leq 0$. And to prove this it is enough to show that for all states s with $v^s > 0$, except for $s = 0$, $\lim_{t\to\infty} \mu_t^{ks}(\vec{x}^*, \vec{y}) = 0$. The Markov matrix $M(\vec{x}^*, \vec{y})$ induces a partition of S into recurrent classes and possibly a transient set.[5] We will end the proof by showing that all states s with positive value, except for state 0, are transient states and thus $\lim_{t\to\infty} \mu_t^{ks}(\vec{x}^*, \vec{y}) = 0$.

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 $v^s > 0$. By Claim 1, there exists a unique $n(s)$ such that $s \in S^+_{n(s)}$. Without loss of generality assume that $v^s \geq v^{s'}$ for all $s' \in C$ and that $n(s) \leq n(s')$ for all s' such that $v^s = v^{s'}$. Consider now state $w(s)$. There are two cases.

Case 1: $w(s) \in C$. In this case, by the choice of s we have that $v^{w(s)} = v^s$. We will show that this case is impossible since it would imply that $n(w(s)) < n(s)$ contradicting our choice of s.

[5] A 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.

Case 1.1: $v^{w(s)} = v^s = v^{\ell(s)}$. Since $s \in S^+_{n(s)}$ and since $\hat{S}(s) \cap \cup_{\nu=0}^{n(s)-1} S^+_\nu \neq \emptyset$, by (3) we have that $w(s) \in \cup_{\nu=0}^{n(s)-1} S^+_\nu$, namely $n(w(s)) < n(s)$.

Case 1.2 : $v^{w(s)} = v^s > v^{\ell(s)}$. In this case $s \not\to \ell(s)$. Then since $s \in S^+_{n(s)}$, we must have either that $w(s) \in \cup_{\nu=0}^{n(s)-1} S^+_\nu$ and $s \to w(s)$, or that $w(s), \ell(s) \in \cup_{\nu=0}^{n(s)-1} S^+_\nu$. In either case, $w(s) \in \cup_{\nu=0}^{n(s)-1} S^+_\nu$, namely $n(w(s)) < n(s)$.

Case 2: $w(s) \notin C$. In this case $\ell(s) \in C$. We will show that this case is also impossible since it would imply that $v^{\ell(s)} = v^s$ and that $n(\ell(s)) < n(s)$, contradicting our choice of s. The fact that $w(s) \notin C$ 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 $s \not\to w(s)$. Then, since $s \in S^+_{n(s)}$, by a similar argument as the one used in Case 1.2, we have that $\ell(s) \in \cup_{\nu=0}^{n(s)-1} S^+_\nu$. To see that $v^{\ell(s)} = v^s$, note that by Observation 1, $v^{w(s)} \geq v^s \geq v^{\ell(s)}$ and if $v^{w(s)} > v^{\ell(s)}$, since $e^s = 0$, by Proposition 2, $v^{\ell(s)} = v^s$.

Since both cases are impossible, we conclude that all the recurrent states have non-positive value – hence the transience of the states with positive value. □

4.4 A partial converse

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, let us go back to the match in Example 1 with $\delta = 0$ and $1/2 < p < 1$, and recall that the value of this match is $v^1 = v^2 = v^3 = 2p - 1$. Consider the following pair of stationary strategies. Player 1 chooses his two actions with equal probabilities in state s_2 and player 2 chooses his second action in state s_1. It can be checked that, independent of the initial state, player 1's strategy guarantees that he gets a payoff of at least $2p - 1$ and that player 2's strategy guarantees that player 1 gets a payoff of at most $2p - 1$. Consequently, these strategies constitute an equilibrium of Γ. However, player 1's strategy is not a minimax-stationary strategy with respect to any labeling since no matter how the successors of s_2

are labeled, the corresponding minimax-stationary strategy will never prescribe mixing between his actions in s_2.

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 a natural labeling.

Let Γ be a quasi-binary match and let $(v^1, \ldots, v^K)$ be its value. Extend it so that $v^0 = 1$ and $v^{K+1} = -1$ and let $v = (v^0, \ldots, v^{K+1})$. 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 $(\vec{x}^*, \vec{y}^*)$ be a stationary equilibrium of Γ and let (w, ℓ) be a natural labeling. Let $k \in 1, \ldots, K$. Since $v^{w(k)} > v^{\ell(k)}$, by Proposition 2

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

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

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

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

Puterman [7], Chapter 7),

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

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

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

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

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

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 $P^k(e^k)$. □

Appendix

Proof of Claim 1: We prove only the first statement. The proof of the other one is analogous and is left to the reader. By definition of the sets in $\{S_0^+, \ldots, S_N^+\}$, it is clear that they are pairwise disjoint. In order to show that their union is S^+ it is enough to show that if $S^+ \setminus \cup_{\nu=0}^n S_\nu^+ \neq \emptyset$ then $S_{n+1}^+ \neq \emptyset$.

Assume by contradiction that $S_{n+1}^+ = \emptyset$ even though $S^+ \setminus \cup_{\nu=0}^n S_\nu^+ \neq \emptyset$. Then, for any $s \in S^+ \setminus \cup_{\nu=0}^n S_\nu^+$, since $\hat{S}(s) \not\subseteq \cup_{\nu=0}^n S_\nu^+$, at most one of its successors is in $\cup_{\nu=0}^n S_\nu^+$. And if s' is such a successor we have that $s \not\to s'$. That is, either $v^{s'} < v^s$ or player 2 can guarantee that the next state is not s', namely there exists $j \in \mathcal{J}$ s.t. $p_{ij}^{ss'} = 0$ for all $i \in \mathcal{I}$. Let $k \in S^+ \setminus \cup_{\nu=0}^n S_\nu^+$ such that $v^k \geq v^{k'}$ for all $k' \in S^+ \setminus \cup_{\nu=0}^n S_\nu^+$. Let $\vec{y}$ be a stationary strategy for player 2 that guarantees that from any $s \in S^+ \setminus \cup_{\nu=0}^n S_\nu^+$, the next state s' is not in $\cup_{\nu=0}^n S_\nu^+$ unless $v^{s'} < v^s$. By the forgoing discussion, such strategy exists. Let $\varepsilon > 0$ be such that $\varepsilon < v^k$ and $2\varepsilon < \min\{|v^s - v^{s'}| : v^s \neq v^{s'}, s, s' \in S\}$). Also, let $\vec{y}_\varepsilon$ be an ε-optimal strategy for player 2 and consider the following strategy for player 2 in Γ^k.

$$\psi(h_t) = \begin{cases} \vec{y}(h_t) & \text{if for all } \tau \leq t,\ s_\tau \in S^+ \setminus \cup_{\nu=0}^n S_\nu^+ \\ \vec{y}_\varepsilon(h_t) & \text{if for some } \tau \leq t,\ s_\tau \notin S^+ \setminus \cup_{\nu=0}^n S_\nu^+. \end{cases}$$

Strategy ψ makes sure that after any history $h_t = (s_0, (i_1, j_1, s_1), \ldots, (i_t, j_t, s_t))$, as long as all the states s_τ, $\tau \leq t$ have been in $S^+ \setminus \cup_{\nu=0}^n S_\nu^+$, the next state s_{t+1} will not be in $\cup_{\nu=0}^n S_\nu^+$, unless $v^{s_{t+1}} < v^{s_t}$ in which case it may be in $\cup_{\nu=0}^n S_\nu^+$. The only way to ever move to a state in $\cup_{\nu=0}^n S_\nu^+$, and in particular, to state 0, is to make a transition from some state $s \in S^+ \setminus \cup_{\nu=0}^n S_\nu^+$ to a state $s' \notin S^+ \setminus \cup_{\nu=0}^n S_\nu^+$ with $v^{s'} < v^s \leq v^k$. But as soon as the system moves from a state in $S^+ \setminus \cup_{\nu=0}^n S_\nu^+$ to a state not there, player 2 switches to the ε-optimal strategy $\vec{y}_\varepsilon$.

Let χ be any stationary strategy for player 1. Since the only way to ever reach state 0 is to go through a state $s \notin S^+ \setminus \cup_{\nu=0}^n S_\nu^+$ with $v^s < v^k$, we have that

$$\begin{aligned} u^k(\chi, \psi) &\leq \max\{0, v^s + \varepsilon : s \text{ with } v^s < v^k\} \\ &\leq \max\{0, v^k - 2\varepsilon + \varepsilon : s \text{ with } v^s < v^k\} \\ &= v^k - \varepsilon \end{aligned}$$

where the second and third inequalities follow from our choice of ε. This inequality, since it holds for every $\chi \in X$, contradicts the fact that v^k is the value of Γ^k.

References

[1] Apesteguia, J., and Palacios-Huerta, I. Psychological pressure in competitive environments: Evidence from a randomized natural experiment. *The American Economic Review* (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] Gonzalez-Díaz, J., and Palacios-Huerta, I. Cognitive performance in competitive environments: Evidence from a natural experiment. Tech. rep., Working Paper, 2016.

[5] Mertens, J.-F., and Neyman, A. Stochastic games. *International Journal of Game Theory 10*, 2 (1981), 53–66.

[6] Palacios-Huerta, I. Professionals play minimax. *The Review of Economic Studies 70*, 2 (2003), 395–415.

[7] Puterman, M. L. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 1994.

[8] Shapley, L. S. Stochastic games. *Proceedings of the National Academy of Sciences of the United States of America 39*, 10 (1953), 1095–1100.

[9] Vieille, N. Two-player stochastic games ii: The case of recursive games. *Israel Journal of Mathematics 119*, 1 (2000), 93–126.

[10] Walker, M., and Wooders, J. Minimax play at wimbledon. *American Economic Review* (2001), 1521–1538.

[11] Walker, M., Wooders, J., and Amir, R. Equilibrium play in matches: Binary markov games. *Games and Economic Behavior 71*, 2 (2011), 487–502.