The duality gap for two-team zero-sum games*

Leonard J. Schulman[1] and Umesh V. Vazirani[2]

1 **Caltech, Engineering & Applied Science MC305-16, Pasadena CA 91125, USA**
`schulman@caltech.edu`

2 **UC Berkeley, Computer Science, Berkeley CA 94720, USA**
`vazirani@eecs.berkeley.edu`

Abstract

We consider multiplayer games in which the players fall in two teams of size k, with payoffs equal within, and of opposite sign across, the two teams. In the classical case of $k = 1$, such zero-sum games possess a unique value, independent of order of play, due to the von Neumann minimax theorem. However, this fails for all $k > 1$; we can measure this failure by a duality gap, which quantifies the benefit of being the team to commit last to its strategy. In our main result we show that the gap equals $2(1 - 2^{1-k})$ for $m = 2$ and $2(1 - m^{-(1-o(1))k})$ for $m > 2$, with m being the size of the action space of each player. At a finer level, the cost to a team of individual players acting independently while the opposition employs joint randomness is $1 - 2^{1-k}$ for $k = 2$, and $1 - m^{-(1-o(1))k}$ for $m > 2$.

This class of multiplayer games, apart from being a natural bridge between two-player zero-sum games and general multiplayer games, is motivated from Biology (the weak selection model of evolution) and Economics (players with shared utility but poor coordination).

1998 ACM Subject Classification G.1.6 Optimization

Keywords and phrases multi-player games, duality gap, zero-sum games, evolution

Digital Object Identifier 10.4230/LIPIcs.CVIT.2016.23

1 Introduction

Games between *teams* of players are ubiquitous; in the economy this occurs most prominently in competition between firms. Another case of significance is that in which a team is a biological species and the players on the team are the genes of the species. What makes a set of players a team, in our idealization, is that in any outcome the players in the set receive identical payoffs.

Competition among firms or species need not be zero-sum; however, the zero-sum case will be the focus of this paper, being the most basic form of competition, and often an approximation to reality. Specifically, a two-team zero-sum game is a multiplayer game in which the players are partitioned into two sets A and B, and a real-valued payoff tensor (of dimension equal to the number of players) specifies the payoff conditional on player actions; this payoff accrues positively to each player of Team B and negatively to each player of Team A.

If perfect coordination within each team can be achieved, then a zero-sum interaction between two teams is nothing but a zero-sum interaction between two "virtual" players. In the biological setting, the opposite extreme is relevant: an important model in evolutionary

* LJS was supported in part by NSF grants 1319745/1618795 and BSF grant 2012333; UVV was supported in part by NSF Grant CCF-1410022.

42nd Conference on Very Important Topics (CVIT 2016).
Editors: John Q. Open and Joan R. Acces; Article No. 23; pp. 23:1–23:8
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

theory is the *weak selection* model (see [8, 1, 10, 9, 13, 14, 3, 7]), in which a species is a team, the genes are the players, the alleles of the gene are the possible actions of a player, and the allele frequencies are independent across genes. Likewise, the difficulty of coordination has long been treated in the economic literature as one of the forces limiting the size of firms [4].

This raises a natural question: does von Neumann's minimax theorem for zero-sum games continue to hold for a zero-sum team game where the individual members of the team play independently, and if not can the deviation be bounded? Formally, this deviation is expressed through the *duality gap* between the values of the game (always expressed as the payoff to Team B) under two scenarios: when all members of Team A must first commit to their randomized strategies, and then Team B gets to respond; and when all members of Team B must commit, and then Team A gets to respond. Surprisingly, this natural question does not seem to have been asked before.

In Section 3 we determine the range of this duality gap: for action spaces of size 2 we give the exact answer for all k (the size of the teams); for action spaces of any size $m > 2$ we determine the asymptotics of the answer as a function of k.

The key lemma in the lower bound on the duality gap for $m > 2$ may be of independent interest: fix a random set S of $g(m)$ m-ary strings of any length k. Then with high probability, any product distribution may place probability more than $m^{-k(1-o(1))}$ (i.e. much more than the uniform distribution would) on at most $O(\log g(m))$ strings in S.

Another interesting quantity is the *defensive gap*: the defensive gap of Team B is the difference between the payoff to Team B if Team A must play a product strategy, and the payoff to Team B if Team A can use a joint (but hidden from B) source of randomness. The defensive gap quantifies the reduced effectiveness of a team when its members are constrained to choose their actions from a product distribution. Note that the duality gap is the sum of the two teams' defensive gaps.

The defensive gap may be compared to two notions in the existing literature. Assume Team A goes first, and think just of the multiplayer game being played by the k players of Team A. (Since Team B can respond optimally, even deterministically, once the strategies of Team A have been fixed, we may ignore the players of Team B and consider their response merely part of the definition of the game being played by Team A.) Then, since the payoffs to all players in Team A are identical, the defensive gap is the difference between the value of the best correlated equilibrium [2] and the best Nash equilibrium [12]. It quantifies, if you will, the penalty for not coordinating. Again, since the payoffs are identical within the team, social welfare agrees with individual welfare, and so the defensive gap is, conceptually, a Price of Stability [11] (although that "price" is normally defined as a ratio and not a difference).

Finally, in Section 4 we go on to investigate how the value of a game can be affected by more incremental changes, specifically, by exchanging the order of play of two players of opposing teams who were playing (committing to their randomized strategy) in immediate succession. There are games with duality gap bounded away from 0, in which all these value changes tend to 0 in k; whereas there are other games, including games symmetric within each team, in which the largest such value change is bounded away from 0 as a function of k.

2 Preliminaries

We consider two-team zero-sum games in which Team A has k players each with m choices in its action space; likewise for Team B. The payoff (to Team B) is specified by a tensor T in $(\mathbb{R}^m)^{\otimes 2k}$. (In the biological setting, each of A and B is a species with a genome of k

genes, each taking on one of m possible alleles.) More formally to each player A_i, $i = 1, \dots k$ of Team A corresponds a vector space $U_i \cong \mathbb{R}^m$, and to each player B_j, $j = 1, \dots k$ of Team B, a vector space $V_j \cong \mathbb{R}^m$. Space U_i is spanned by a basis which we denote $u_{i0}, \dots, u_{i,m-1}$. Similarly for V_j. In this setting $T \in U_1^* \otimes \dots \otimes U_k^* \otimes V_1^* \otimes \dots \otimes V_k^*$. (With $*$ denoting dualization.) Letting $\boldsymbol{I} = (i_1, \dots, i_k) \in \{0, \dots, m-1\}^k$ represent an action of the players of Team A, and $\boldsymbol{J} = (j_1, \dots, j_k) \in \{0, \dots, m-1\}^k$ an action of the players of Team B, $T_{\boldsymbol{I}}^{\boldsymbol{J}}$ is the payoff to players of Team B (and minus the payoff to players of Team A).

In the case $m = 2$, the strategy of player A_i is specified by a parameter $0 \le p_i \le 1$ which is the probability with which he plays choice 0, i.e., vector u_{i0}. Likewise player B_j has a parameter $0 \le q_j \le 1$ which is the probability with which he plays choice 0, i.e., vector v_{j0}. For $m > 2$, the strategy of player A_i is specified by a probability distribution $(p_{i0}, \dots, p_{i,m-1})$ and the strategy of player B_j is specified by a probability distribution $(q_{j0}, \dots, q_{j,m-1})$. (Thus for $m = 2$, $(p_i, 1 - p_i)$ is shorthand for (p_{i0}, p_{i1}).)

We let T_p^q denote the expected payoff to Team B when the players of Team A use distributions p_i and those of Team B use distributions q_j. This notation generalizes the notation $T_{\boldsymbol{I}}^{\boldsymbol{J}}$, if one interprets $\boldsymbol{I}$ as the probability distribution on $\{0, \dots, m-1\}^k$ supported solely on $\boldsymbol{I}$ (and similarly for $\boldsymbol{J}$). Equivalently, T_p^q equals the scalar given by contracting T with the tensor product of the vectors $(p_{i0}, \dots, p_{i,m-1})$ (ranging over i) and $(q_{j0}, \dots, q_{j,m-1})$ (ranging over j).

By a standard argument, $\min_p \max_q T_p^q \ge \max_q \min_p T_p^q$. (We write everywhere min or max rather than inf or sup since the spaces are compact.) However, apart from the linear ($k = 1$) case, the gap $\min_p \max_q T_p^q - \max_q \min_p T_p^q$ can be positive.

Our purpose is to quantify this gap relative to the uniform norm $\|T\|_\infty = \max_{\boldsymbol{I},\boldsymbol{J}} |T_{\boldsymbol{I}}^{\boldsymbol{J}}|$. We define the duality gap of tensor T:

$$\operatorname{gap}(T) = \frac{\min_p \max_q T_p^q - \max_q \min_p T_p^q}{\|T\|_\infty} = \frac{\min_p \max_{\boldsymbol{J}} T_p^{\boldsymbol{J}} - \max_q \min_{\boldsymbol{I}} T_{\boldsymbol{I}}^q}{\|T\|_\infty} \tag{2.1}$$

where as above, $\boldsymbol{I}$ or $\boldsymbol{J}$ represent the pure strategy choosing that action.

The principal quantity of interest is

$$\operatorname{gap}_{m,k} = \max_T \operatorname{gap}(T) \tag{2.2}$$

ranging over games T for teams of size k and action spaces of size m. It is trivial that $\operatorname{gap}_{m,k} \le 2$. Moreover $\operatorname{gap}_{m,k}$ is nondecreasing in k (because one may ignore the actions of players after the k'th player on each team), and in m (because one may map all actions $\ge m - 1$ to action $m - 1$).

Here and throughout the paper, upper-case P and Q denote mixed strategies of virtual players; that is to say, each is a general probability tensor (a tensor with nonnegative entries summing to 1), $P \in U_1 \otimes \dots \otimes U_k$ and $Q \in V_1 \otimes \dots \otimes V_k$. Lower-case p and q denote product distributions, i.e., rank-one probability tensors.

Extending the existing notation, T_P^Q is the expected payoff to B when A (as a virtual player) uses distribution P and B uses distribution Q. It is also useful to employ the standard convention that a repeated index indicates tensor contraction over that index, so $P^{\boldsymbol{I}} T_{\boldsymbol{I}}^{\boldsymbol{J}} = T_P^{\boldsymbol{J}} \in V_1^* \otimes \dots \otimes V_k^*$ and $Q_{\boldsymbol{J}} T_{\boldsymbol{I}}^{\boldsymbol{J}} = T_{\boldsymbol{I}}^Q \in U_1^* \otimes \dots \otimes U_k^*$.

By strong LP duality we can define the value of the virtual player game by

$$\operatorname{Val} T = \min_P \max_{\boldsymbol{J}} \{P^{\boldsymbol{I}} T_{\boldsymbol{I}}^{\boldsymbol{J}}\} = \max_Q \min_{\boldsymbol{I}} \{Q_{\boldsymbol{J}} T_{\boldsymbol{I}}^{\boldsymbol{J}}\}. \tag{2.3}$$

Let P and Q be strategies achieving equality in (2.3). We can usefully refine the study of gap(T) by defining the *defensive gap* of Team A in tensor T as

$$\operatorname{gap}_A(T) = \frac{\min_p \max_{\boldsymbol{J}}\{p^{\boldsymbol{I}} T_{\boldsymbol{I}}^{\boldsymbol{J}}\} - \max_{\boldsymbol{J}}\{P^{\boldsymbol{I}} T_{\boldsymbol{I}}^{\boldsymbol{J}}\}}{\|T\|_\infty} = \frac{\min_p \max_{\boldsymbol{J}}\{p^{\boldsymbol{I}} T_{\boldsymbol{I}}^{\boldsymbol{J}}\} - \operatorname{Val} T}{\|T\|_\infty}$$

where, of course, p ranges over product distributions. Likewise the defensive gap of Team B is

$$\operatorname{gap}_B(T) = \frac{\min_{\boldsymbol{I}}\{Q_{\boldsymbol{J}} T_{\boldsymbol{I}}^{\boldsymbol{J}}\} - \max_q \min_{\boldsymbol{I}}\{q_{\boldsymbol{J}} T_{\boldsymbol{I}}^{\boldsymbol{J}}\}}{\|T\|_\infty} = \frac{\operatorname{Val} T - \max_q \min_{\boldsymbol{I}}\{q_{\boldsymbol{J}} T_{\boldsymbol{I}}^{\boldsymbol{J}}\}}{\|T\|_\infty}$$

The defensive gap quantifies the reduced effectiveness of a team of players (when forced to commit to a mixed strategy to which the other team has a chance to respond), as compared with a virtual player (in the same situation).

3 The Defensive Gaps and the Duality Gap

▶ **Theorem 1.** $\operatorname{gap}_{2,k} = 2(1 - 2^{1-k})$*, and for every* $m > 2$, $2(1 - m^{-(1-o(1))k}) \le \operatorname{gap}_{m,k} \le 2(1 - m^{1-k})$*. (The "*$o(1)$*" being w.r.t. k.) More specifically:*

1. *Upper bound on the defensive gaps and duality gap:*
For any $m \ge 2$*, and for any* T *having action spaces of size* m*,* $\operatorname{gap}_A(T) \le (1 - \operatorname{Val} T/\|T\|_\infty)(1 - m^{1-k})$ *and* $\operatorname{gap}_B(T) \le (1 + \operatorname{Val} T/\|T\|_\infty)(1 - m^{1-k})$*.*

2. *Lower bound on the duality gap:*
  a. $\operatorname{gap}_{2,k} \ge 2(1 - 2^{1-k})$*.*
  b. *For every* $m > 2$ *there is a function* $\varepsilon(k)$ *tending to 0 as* $k \to \infty$*, such that* $\operatorname{gap}_{m,k} \ge 2(1 - m^{-(1-\varepsilon(k))k})$*.*

Henceforth scale any $T \ne 0$ so that $\|T\|_\infty = 1$.

Proof.

Part (1): Upper Bounds on $\operatorname{gap}_A(T)$**,** $\operatorname{gap}_B(T)$ **and** $\operatorname{gap}_{m,k}$**.**

It suffices to show the claim for gap_A. The claim for gap_B follows by negating all entries of T, reversing the roles of the teams and applying the claim for gap_A. The claim for the duality gap follows because $\operatorname{gap}(T) = \operatorname{gap}_A(T) + \operatorname{gap}_B(T)$.

We start with the case $m = 2$. Given an arbitrary coordinated mixed strategy P for team A, we wish to "round" it to a rank-one probability tensor (a product distribution) that does no worse than the claimed defensive gap. The natural candidate for rounding would be the approximation by independent random variables having the same marginals as P. That is, set

$$p_1 = \sum_{i_2,\ldots,i_k} P^{0,i_2,\ldots,i_k}, \quad p_2 = \sum_{i_1,i_3,\ldots,i_k} P^{i_1,0,i_3,\ldots,i_k}, \quad \text{etc.} \tag{3.1}$$

and, letting

$$p = (p_1, 1-p_1) \otimes \ldots \otimes (p_k, 1-p_k), \tag{3.2}$$

use p as the rank-one strategy replacing P, and use it to bound the defensive gap. It turns out that this approach cannot be used to prove any bound on the defensive gap.

Surprisingly, there is a less obvious rank-one strategy which can be obtained from P and which provides a tight bound on the defensive gap.

For $0 \le x \le 1/2$ set

$$\beta(x) = \frac{(2x)^{1/k}}{2} \tag{3.3}$$

and for $1/2 < x \le 1$ set $\beta(x) = 1 - \beta(1-x)$. (Observe that β is increasing.)

Then use the rank-one strategy:

$$\tilde{p} = (\beta(p_1), \beta(1-p_1)) \otimes \ldots \otimes (\beta(p_k), \beta(1-p_k)) \tag{3.4}$$

We now claim that for every $\boldsymbol{I}$, $\tilde{p}^{\boldsymbol{I}} \ge 2^{1-k} P^{\boldsymbol{I}}$. We show this for $\boldsymbol{I} = \mathbf{0}$; all other $\boldsymbol{I}$ follow by the same argument. First note that for every $1 \le i \le k$, $p_i \ge P^{\mathbf{0}}$, and therefore $\beta(p_i) \ge \beta(P^{\mathbf{0}})$. So $\tilde{p}^{\mathbf{0}} = \beta(p_1) \cdots \beta(p_k) \ge \beta(P^{\mathbf{0}})^k$. Now there are two cases. If $P^{\mathbf{0}} \le 1/2$ then $\beta(P^{\mathbf{0}})^k = 2^{1-k} P^{\mathbf{0}}$ as desired. If $P^{\mathbf{0}} > 1/2$ then $\beta(P^{\mathbf{0}})^k = (1 - \frac{1}{2}(2 - 2P^{\mathbf{0}})^{1/k})^k$. Showing this is $\ge 2^{1-k} P^{\mathbf{0}}$ is equivalent to showing that

$$(2 - (2P^{\mathbf{0}})^{1/k})^k \ge 2 - 2P^{\mathbf{0}}. \tag{3.5}$$

Let $\varepsilon = (2P^{\mathbf{0}})^{1/k} - 1$ and note $0 \le \varepsilon \le 1$. Then (3.5) is equivalent to $(1-\varepsilon)^k + (1+\varepsilon)^k \ge 2$ which holds by the power-mean inequality (for k vs. 1).

From $\tilde{p}^{\boldsymbol{I}} \ge 2^{1-k} P^{\boldsymbol{I}}$ for every $\boldsymbol{I}$ and from $\|T\|_\infty \le 1$, we have that for every $\boldsymbol{J}$:

$$\tilde{p}^{\boldsymbol{I}} T_{\boldsymbol{I}}^{\boldsymbol{J}} \le 1 - 2^{1-k} + 2^{1-k} P^{\boldsymbol{I}} T_{\boldsymbol{I}}^{\boldsymbol{J}} \tag{3.6}$$

which (upper bounding $P^{\boldsymbol{I}} T_{\boldsymbol{I}}^{\boldsymbol{J}}$ by $\operatorname{Val} T$, and subtracting $\operatorname{Val} T$ from each side) completes the proof for $m = 2$.

The same proof technique works for $m > 2$, only there is more flexibility in the function β. Given that in an optimal virtual-player distribution P for Team A, the marginal distribution of player i is $p_i(0), \ldots, p_i(m-1)$, we require a new distribution $\beta_i(0), \ldots, \beta_i(m-1)$ for player i such that for every ℓ, $\beta_i(\ell)^k \ge m^{1-k} p_i(\ell)$, which is to say, $\beta_i(\ell) \ge (m p_i(\ell)^{1/k})/m$. Such a distribution exists due to the inequality $\sum_\ell (m p_i(\ell))^{1/k}/m \le 1$, which holds because by the power-mean inequality $(\sum_\ell (m p_i(\ell))^{1/k}/m)^k \le \sum_\ell (m p_i(\ell))/m = 1$. The rest of the details are as for $m = 2$.

Part (2a): Lower Bound on $\mathrm{gap}_{2,k}$.

Write $\mathbf{0} = (0, \ldots 0)$ and $\mathbf{1} = (1, \ldots 1)$.

Consider the following tensor.

▶ **Example 2.**

$$\begin{cases} G_{\boldsymbol{I}}^{\mathbf{0}} &= \begin{cases} -1 & \text{if } \boldsymbol{I} = \mathbf{0} \\ 1 & \text{otherwise} \end{cases} \\ G_{\boldsymbol{I}}^{\mathbf{1}} &= \begin{cases} -1 & \text{if } \boldsymbol{I} = \mathbf{1} \\ 1 & \text{otherwise} \end{cases} \\ G_{\mathbf{0}}^{\boldsymbol{J}} &= \begin{cases} 1 & \text{if } \boldsymbol{J} = \mathbf{1} \\ -1 & \text{otherwise} \end{cases} \\ G_{\mathbf{1}}^{\boldsymbol{J}} &= \begin{cases} 1 & \text{if } \boldsymbol{J} = \mathbf{0} \\ -1 & \text{otherwise} \end{cases} \\ G_{\boldsymbol{I}}^{\boldsymbol{J}} &= 0 \text{ for all other } \boldsymbol{I}, \boldsymbol{J} \end{cases} \tag{3.7}$$

An informal description of this game is that if both Team A and Team B choose actions in $\{\mathbf{0}, \mathbf{1}\}$, then the outcome is as it would be in the "matching pennies" game. If just one of the teams chooses an action in $\{\mathbf{0}, \mathbf{1}\}$, then that team wins. If neither team chooses an action in $\{\mathbf{0}, \mathbf{1}\}$, then the game is a tie.

(We incidentally note that the proof of Part (2a) does not depend on setting entries to 0 in the last line of 3.7; the argument will go through with each entry taking any value in $[-1, 1]$.)

Now consider any strategy $p = (p_1, \ldots, p_k)$ for Team A (recall these are the probabilities of action 0). The expected payoff for action $\boldsymbol{J} = \mathbf{0}$ of Team B is

$$G_p^{\mathbf{0}} = 1 - 2p_1 \cdots p_k$$

The expected payoff for $\boldsymbol{J} = \mathbf{1}$ is

$$G_p^{\mathbf{1}} = 1 - 2(1 - p_1) \cdots (1 - p_k)$$

By the arithmetic-geometric mean inequality, $G_p^{\mathbf{0}} \geq 1 - 2(\frac{1}{k} \sum p_i)^k$ and $G_p^{\mathbf{1}} \geq 1 - 2(\frac{1}{k} \sum (1 - p_i))^k$. So

$$\frac{1 - \max\{G_p^{\mathbf{0}}, G_p^{\mathbf{1}}\}}{2} \leq \min\{(\frac{1}{k} \sum p_i)^k, (\frac{1}{k} \sum (1 - p_i))^k\}$$

equivalently

$$\left(\frac{1 - \max\{G_p^{\mathbf{0}}, G_p^{\mathbf{1}}\}}{2}\right)^{1/k} \leq \min\{\frac{1}{k} \sum p_i, \frac{1}{k} \sum (1 - p_i)\}$$

The RHS is at most 1/2. So $\min_p \max\{G_p^{\mathbf{0}}, G_p^{\mathbf{1}}\} \geq 1 - 2^{1-k}$.

A similar argument applied to the strategy q of Team B establishes that $\max_q \min\{G_{\mathbf{0}}^q, G_{\mathbf{1}}^q\} \leq -1 + 2^{1-k}$. Adding the two contributions, Part (2a) of the theorem follows.

Part (2b): Lower Bound on $\mathrm{gap}_{m,k}$**,** $m > 2$**.**

Proof: We non-constructively exhibit a game establishing the lower bound. We start by selecting, for a function $g(m)$ to be specified, $g(m)$ strings $S = (s^1, \ldots, s^{g(m)})$, each s^j chosen independently and uniformly in $\{0, \ldots, m-1\}^k$ (Think of m as arbitrary but fixed while $k \to \infty$). The first idea of the proof lies in the interesting fact that with high probability, this set (whose size is independent of k) has the property that any product distribution may place probability more than $m^{-k(1-o(1))}$ on at most $\sim \log g(m)$ strings in S.

To describe the second idea, we start by associating the strings of S with players in a tournament. (A tournament is a digraph in which there is one directed edge between every pair of distinct vertices.) The game is then as follows. Associate the strings of S with the vertices of the tournament in an arbitrary way. If neither team chooses a string in S, the game is a tie. If one team chooses a string in S and the other does not, the first team wins. If both teams choose strings in S, the winning vertex is that which points toward the other (with a tie for identical strings). In all cases a win means a payoff of 1 and a tie a payoff of 0.

Key to the second idea is that by a result of Erdös [5] (later constructively in [6]) it is possible to choose the tournament so that it has no dominating set of size $\sim \log g(m)$. (A set of vertices is dominating if every vertex outside the set is pointed to by at least one vertex in the set.)

Now we claim that whichever team goes second can achieve payoff $1 - m^{-(1-o(1))k}$. The argument is the same for both teams, so say Team A goes first and let p be its product strategy. Let R be the $\sim \log g(m)$ strings in S accorded highest probability in p; all other strings of S have probability less than $m^{-(1-o(1))k}$. Team B responds to p with an $s \in S$ which dominates all of R. Team B wins unless Team A selects an $s' \in S - R$ (and sometimes even then). The payoff to Team B is therefore at least $1 - g(m)m^{-(1-o(1))k}$ which is at least $1 - m^{-(1-o(1))k}$. ◀

4 Order Refinements

It is natural to consider a more general scenario in which players of the two teams commit to their strategies in some (not necessarily strict) alternation. That is to say, let π be any bijection from $\{1,\ldots,2k\}$ to $\{A_1,\ldots,A_k,B_1,\ldots,B_k\}$. If $\pi(\ell)=A_i$ for some i then let $M(\ell)$ be the quantifier $\min_{p_i}$; if $\pi(\ell)=B_j$ for some j then let $M(\ell)=\max_{q_j}$. Then the value of game T with respect to order π is defined to be

$$V(T,\pi)=M(1)\ldots M(2k)T_p^q.$$

In particular, let π^{AB} be an order in which all the members of Team A go first, that is, $\pi^{AB}(\ell)=A_\ell$ for $\ell\le k$, and $\pi^{AB}(\ell)=B_{\ell-k}$ for $\ell>k$. (Note that V is invariant under exchange of same-team players with adjacent quantifiers.) Likewise let π^{BA} be an order in which Team B goes first. Then the duality gap of game T is

$$\mathrm{gap}(T)=V(T,\pi^{AB})-V(T,\pi^{BA}).$$

We now ask how much V may change when we change π by a single adjacent transposition.

A natural place to study this question is in *symmetric games*, by which we mean games invariant under permutation of the actions taken by members of a team. Even in this restrictive setting, a wide range of behaviors can occur.

For one extreme, we return to the game G of Example 2. We show that for any k, the order of the first three players can affect the outcome decisively.

▶ **Lemma 3.** *If $\ell\ge 3$ is the first time that Team B plays in π, then $V(G,\pi)\ge 1/2$. Likewise if $\ell\ge 3$ is the first time that Team A plays in π, then $V(G,\pi)\le -1/2$.*

For the other extreme, we have:

▶ **Theorem 4.** *There exist symmetric team games with any $k\ge 1$ players per team, with duality gap bounded away from 0 (as a function of k), but in which any adjacent transposition in the order of play changes the value of the game only by $O(1/k)$.*

Proofs of the lemma and theorem will be provided in the journal version of the paper.

5 Discussion

We have characterized the possible range of the duality gap. The examples which achieved large gap were highly structured. It would be interesting to find natural conditions on a game (particularly a symmetric game) that ensure small duality gap.

It would be interesting to extend our inquiry to teams (possibly more than two) competing in non-zero-sum games.

References

1 H Akashi. Inferring weak selection from patterns of polymorphism and divergence at "silent" sites in Drosophila DNA. *Genetics*, 139(2).

2 R Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.

3 E Chastain, A Livnat, C Papadimitriou, and U Vazirani. Algorithms, games, and evolution. *Proc. National Academy of Sciences*, 2014.

4 R H Coase. The nature of the firm. *Economica, New Series*, 4(16):386–405, Nov. 1937.

5 P Erdös. On a problem in graph theory. *Math. Gaz.*, 47:220–223, 1963.
6 R L Graham and J H Spencer. A constructive solution to a tournament problem. *Canad. Math. Bull.*, 14(1):45–48, 1971.
7 R Mehta, I Panageas, and G Piliouras. Natural selection as an inhibitor of genetic diversity: Multiplicative weights updates algorithm and a conjecture of haploid genetics. In *Proc. ITCS*, page 73. ACM, 2015.
8 T Nagylaki. The evolution of multilocus systems under weak selection. *Genetics*, 134(2):627–647, 1993.
9 MA Nowak, A Sasaki, C Taylor, and D Fudenberg. Emergence of cooperation and evolutionary stability in finite populations. *Nature*, 428:646–650, 2004. `doi:10.1038/nature02414`.
10 T Ohta. Near-neutrality in evolution of genes and gene regulation. *Proc. National Academy of Sciences*, 99(25):16134–16137, 2002. `doi:10.1073/pnas.252626899`.
11 T Roughgarden and E Tardos. Ch. 17: Introduction to the inefficiency of equilibria. In N Nisan, T Roughgarden, E Tardos, and V V Vazirani, editors, *Algorithmic Game Theory*. Cambridge U Press, 2007.
12 J Von Neumann. Zur theorie der gesellschaftsspiele. *Math. Ann.*, 100:295–320, 1928.
13 B Wu, PM Altrock, L Wang, and A Traulsen. Universality of weak selection. *Physical Review E*, 82:046106, 2010.
14 B Wu, J García, C Hauert, and A Traulsen. Extrapolating weak selection in evolutionary games. *PLoS Comput Biol*, 9(12), 2013. `doi:10.1371/journal.pcbi.1003381`.