Selection in the Presence of Noise: The Design of Playoff Systems

Micah Adler * Peter Gemmell # Mor Harchol ** Richard M. Karp ##
Claire Kenyon ***

November 7, 1995

1 Introduction

In every sport, playoffs and tournaments are used to select the best among a set of competing players or teams. In this paper we consider the optimal design of such systems. We seek designs that are optimally efficient, in the sense that they minimize the number of rounds or the number of games needed to select the best player with a stated probability. Our models reflect the fact that the better player in games between two players or teams does not always win. As a consequence, the problems we consider are not equivalent to choosing the smallest of n elements in the standard comparison model.

We assume that there are n players. There is an initially unknown one-to-one correspondence between the set of players and the index set $\{1, 2, \cdots, n\}$. The player corresponding to index j is called Player j. Player 1 is the best player in the following sense: in any game between Player 1 and Player j, Player 1 wins with probability p_{1j}, where $p_{1j} > 1/2$; draws are not allowed. The goal is to determine the identity of Player 1 with probability at least $1 - \Delta$, where Δ is a given constant. The games are played in *rounds* where, in each round, each player participates in at most one game. We seek to minimize the number of rounds (or, in some cases, the expected number of rounds) needed to achieve the goal. A secondary objective is to minimize the number of games.

*Computer Science Division, UC Berkeley, Berkeley, CA 94720.

#Sandia National Laboratories, Albuquerque, NM 87185. Part of work completed while at Computer Science Division, UC Berkeley, Berkeley, CA 94720., supported by NSF grant number CCR-9201092.

**Computer Science Division, UC Berkeley, Berkeley, CA 94720., supported by the National Physical Science Consortium Fellowship and NSF grant number CCR-9201092.

##Computer Science Division, UC Berkeley, Berkeley, CA 94720 and International Computer Science Institute, Berkeley, CA. Supported by NSF grant number CCR-9005448.

***Research supported by International Computer Science Institute, Berkeley, CA. Mailing address: LIP, ENS-Lyon, 46 allée d'Italie, 69364 Lyon Cedex 07, France.

We consider three different models, which differ in their assumptions about the outcomes of games that do not involve Player 1.

The Adversary Model This model assumes that the outcomes of all games that do not involve Player 1 are under the control of an adversary; i.e., they are completely unpredictable.

The Strong Transitivity Model This model assumes that there is a fixed ranking of the players such that a higher-ranked player always has at least a 50 percent chance of beating a lower-ranked player and, for any fixed player, the stronger the opponent, the lower the probability of winning. These assumptions have been widely adopted in connection with certain problems of statistical inference using paired comparisons [D]. We formalize this as follows:

- There is a matrix (p_{ij}) such that, whenever Player i faces Player j, Player i wins with probability p_{ij} , where $p_{ij} \geq 0$ and $p_{ij} + p_{ji} = 1$.
- $p_{ij} \geq 1/2$ whenever $i < j$.
- If $i < j < k$ then $p_{ik} \geq \max(p_{ij}, p_{jk})$; this property is called *strong transitivity*.

The Discriminating Model This is a special case of the strong transitivity model in which the matrix P is *discriminating*: i.e., if $i < j$ then, for $k = 1, 2, \cdots, n-1$, $\frac{p_{ki}}{p_{kj}} \geq \frac{p_{k+1,i}}{p_{k+1,j}}$. This condition can be interpreted as saying that, the weaker the common opponent k, the greater the significance of a loss to Player k in distinguishing between Player i and Player j. The condition holds for a number of natural concrete assumptions about the matrix P, including the following: Player i has a *strength* μ_i, and his performance in any given game is a normal random variable with mean μ_i and standard deviation 1. When two players meet the one with the higher performance wins.

Observe that the adversary model is the most general. The strong transitivity model is a special case of the adversary model, and the discriminating model is a special case of the strong transitivity model.

In each of the three models we may assume either **known win probabilities** - i.e., that the row (p_{1j}) is known to the algorithm - or **unknown win probabilities** - i.e., that the row (p_{1j}) is unknown to the algorithm. Note that, even in the case of known win probabilities, the algorithm is provided with no information about the win probabilities for games that do not involve Player 1, even though these probabilities are well-defined in the strong transitivity and discriminating models.

Thus we have six cases to consider, corresponding to three possible models for the outcomes of games, and, within each of these, two possible assumptions about the knowledge available to the algorithms. We designate a case by an ordered pair in which the first component (ADV, TRANS or DISC) indicates the model for the outcomes of games and the second component (K or U) indicates whether the win probabilities are known or unknown. Thus (ADV, K) denotes the adversary model with known win probabilities. In each of the six cases we allow our algorithms to be randomized, and restrict attention to algorithms that, for all choices of Δ, n and (p_{1j}), select Player 1 with probability at least $1-\Delta$.

In the case of known win probabilities it is possible to define a nonuniform algorithm that is "pointwise optimal"; i.e., for every fixed choice of the (known) win probabilities (p_{1j}), it minimizes the worst-case expected number of rounds. Let $T_{ADV}(\Delta, n, (p_{1j}))$, $T_{TRANS}(\Delta, n, (p_{1j}))$ and $T_{DISC}(\Delta, n, (p_{1j}))$ denote the worst-case expected running time of this pointwise optimal algorithm in the cases (ADV,K), (TRANS,K) and (DISC,K) respectively.

In the three cases involving unknown win probabilities no such pointwise optimal algorithm exists, and thus a more elaborate definitional framework is required in order to describe the complexity of the problem in all cases. The function $F(\Delta, n, (p_{1j}))$ is called an *upper bound on complexity* if, for every choice of Δ, n and (p_{1j}), there is an algorithm that runs within an expected number of rounds bounded above by $F(\Delta, n, (p_{1j}))$, a *lower bound on complexity* if, for every choice of Δ, n and (p_{1j}), every algorithm requires an expected number of rounds greater than or equal to F, and an *existential lower bound on complexity* if for every algorithm there exists a choice of Δ, n and (p_{1j}) for which the expected number of rounds is at least $F(\Delta, n, (p_{1j}))$.

1.1 Main Results

Let $a_j = p_{1j} - \frac{1}{2}$ and let $\epsilon_i = \frac{\sum_{j=2}^{2^i} a_j}{2^i - 1}$. Thus a_j is Player 1's advantage against Player j, and ϵ_i is Player 1's average advantage against Players $2, 3, \cdots, 2^i$. We assume throughout that $\Delta \leq \frac{1}{3}$.

Observe that $\lg n$ is a lower bound on complexity for all models. This is because the declared champion of any tournament algorithm, which is correct more than half the time, must be indirectly compared with at least $\frac{n}{2}+1$ inputs.

For all six cases we demonstrate a lower bound on complexity of $LB = \frac{1}{\lg(1+\frac{\lg(n)}{\lg(\frac{1}{\Delta})})} \sum_{i=1}^{\lg n} \frac{1}{\epsilon_i^2}$. We also derive an existential lower bound of $\frac{n}{16}\frac{1}{\epsilon_1^2}$ for the case (ADV,U).

Our upper bounds on complexity are as follows:

Known Win Probabilities

	Upper Bound
ADV	$O\left(\lg(\frac{1}{\Delta})\sum_{i=1}^{\lg n}\frac{1}{\epsilon_i^2} + \lg(\frac{1}{\Delta})\lg\lg(\frac{1}{\Delta})\frac{1}{\epsilon_1^2}\right)$
TRANS	$O\left(\lg(\frac{1}{\Delta})\sum_{i=1}^{\lg n}\frac{1}{\epsilon_i^2} + \lg(\frac{1}{\Delta})\lg\lg(\frac{1}{\Delta})\frac{1}{\epsilon_1^2}\right)$
DISC	$\min($ $O\left(\lg(\frac{1}{\Delta})\sum_{i=1}^{\lg n}\frac{1}{\epsilon_i^2} + \lg(\frac{1}{\Delta})\lg\lg(\frac{1}{\Delta})\frac{1}{\epsilon_1^2}\right)$, $O\left(LB + \ln(\frac{1}{\Delta})\lg^2 n\right))$

Unknown Win Probabilities

	Upper Bound
ADV	?
TRANS	$O(\sum_{i=1}^{\lg n}(\frac{1}{\epsilon_i^2}(\lg\lg\frac{1}{\epsilon_i} + \lg\frac{1}{\Delta} + \lg n)))$
DISC	$O(\sum_{i=1}^{\lg n}(\frac{1}{\epsilon_i^2}(\lg\lg\frac{1}{\epsilon_i} + \lg\frac{1}{\Delta} + \lg n)))$

In Section 3 we prove the lower bound on complexity LB for the case (DISC,K). Since this is the most favorable case from the point of view of the algorithm, this lower bound applies to all six cases.

In Section 4 we prove an upper bound on complexity of $O\left(\lg(\frac{1}{\Delta})\sum_{i=1}^{\lg n}\frac{1}{\epsilon_i^2} + \lg(\frac{1}{\Delta})\lg\lg(\frac{1}{\Delta})\frac{1}{\epsilon_1^2}\right)$ for the case (ADV,K). Since the adversary model is the least favorable for the algorithm designer, this upper bound also applies to the cases (TRANS,K) and (DISC,K). This upper bound is particularly interesting, as it is achieved using a variant of a common method of pairing chess players called the Swiss System, in which players with equal scores are matched whenever possible. The bound holds within the adversary model, which makes no assumptions about the outcomes of matches not involving Player 1, and yet comes within roughly a $\lg\lg n$ factor of a lower bound that applies even under the rather specific assumptions of the discriminating model. Thus the result shows that our variant of the Swiss system is both efficient and robust. We note that, even in the symmetric case where forall j greater than or equal to 2, $p_{1j} = \frac{1}{2} + \epsilon$, our algorithm beats any obvious variant of a knockout tournament by a factor of $\lg\lg(n)$. In the general case, where the p_{1j} values will vary, our algorithm beats any obvious knockout variant by a factor of as much as $\lg(n)\lg\lg(n)$.

In Section 5 we improve the upper bound on complexity for (DISC,K) in the case when the best player's advantage over the second best player is small.

The remaining sections assume the first row is unknown. In Section 6 we prove an existential lower bound on complexity for (ADV,U). This lower bound implies that uniformly efficient selection procedures are not possible in the adversary model, when the win probabilities are unknown.

In Section 7 we prove an upper bound on complexity for (TRANS,U) in the case where there are just two players. This is used in Section 8 to prove an upper bound on complexity for (TRANS,U) in the general case of n players. This upper bound also applies to the case (DISC,U). Observe that in contrast to the adversary model, efficient selection procedures are possible under the strong transitivity and discriminating models with unknown win probabilities.

In deriving our upper bounds, we have concentrated on asymptotic results and our constants are sometimes very large. Nevertheless, we believe that our results do provide insights into the design of real-world playoff systems.

2 Previous Results

The problem of selecting the best of n players using unreliable comparisons was addressed in [RGL], where Ravikumar, Ganesan and Lakshmanan assume that the total number of erroneous outcomes is less than some absolute upper bound e. They show that $(e+1)n-1$ comparisons are necessary and sufficient to find the best player.

In [FPRU], Feige, Peleg, Raghavan and Upfal choose a probabilistic model, assuming that each comparison has a fixed probability p of being erroneous, and that successive comparisons are independent. The goal is then to select the best player with probability at least $1-\Delta$, for some fixed confidence level Δ. They give a parallel algorithm operating within $O(\lg n)$ rounds and $O(n)$ comparisons. Although they do not point this out, their proof does not depend on any assumptions about the outcomes of games that do not involve Player 1, and thus implies an upper bound on complexity of order $\lg n$ for (ADV,K), (TRANS,K) and (DISC, K) in the case where p_{1j} is equal to a constant greater than $1/2$ for all j. These bounds are optimal up to constant factors.

Previous work on the strong transitivity model has focused on classical "knockout tournaments", where a player is eliminated as soon as he or she loses a game, and a total of $(n-1)$ games are played ([ChuH], [Hw], [I], [CheH]). The tournament can be represented as a tree, each leaf containing a player and each internal node containing the winner of a game between its two children. Since knockout tournaments are not very efficient in selecting the best player with unreliable games, several authors have considered generalizations of the knockout tournament in which each node of the tournament tree represents a match between two players extending over a series of games, rather than a single game. Such schemes are studied in [HM].

The present paper seems to be the first to give a lower bound on the number of rounds required to select the best player with probability $1-\Delta$ in the strong transitivity model, and the first to consider the adversary and discriminating models at all (although special cases of the discriminating model are studied in [Br] and [Th]).

Coping with unreliable information has also been studied in other contexts. In particular, searching with erroneous comparisons was initiated by Rivest, Meyer, Kleitman, Winklmann and Spencer [RMKWS], assuming that the number of errors is less than e. Pelc [P] studied that problem in the probabilistic model (with fixed error probability p). Pelc [P], and Aslam and Dhagat [AD] worked on the model of "linearly bounded errors", where they assume that there is a constant r such that each initial sequence of i comparison questions receives at most ri erroneous answers.

3 A Lower Bound for $T_{DISC,K}$

Let $LB = \frac{\sum_{i=1}^{\lg(n)} \frac{1}{\epsilon_i^2}}{\lg(1+\frac{\lg(n)}{\lg(\frac{1}{\Delta})})}$.

Theorem 1 *If T is the expected number of rounds used by a tournament algorithm to find the best of n players with confidence $1-\Delta$, where the algorithm is given only the first row of the discriminating probability matrix, (p_{ij}), then $T = \Omega(LB)$.*

Proof. (Sketch)

We give a construction which extends any vector (p_{1j}) to an $n \times n$ matrix (p_{ij}) for which the lower bound holds (even if the algorithm is given the entire matrix (p_{ij})). For any real μ let $X(\mu)$ be a random variable that has the normal distribution with mean μ and standard deviation 1. Given p_{1j}, choose $\mu_1, \mu_2, \cdots, \mu_n$ such that, if X_{μ_1} and X_{μ_j} are independent, then $Pr[x_{\mu_1} < X_{\mu_j}] = p_{1j}$. Now define the rest of the matrix (p_{ij}) by the rule $p_{ij} = Pr[x_{\mu_i} < X_{\mu_j}]$ where X_{μ_i} and X_{μ_j} are independent. This matrix of win probabilities has the following interpretation: whenever Player i participates in a game he draws a value from the normal distribution with mean μ_i and standard deviation 1; in each game, the player with the smaller value wins.

Now we can prove our lower bound using the "little birdie" principle. Suppose that, in each round, instead of being told the winner of each game, the algorithm is told the actual values that the players draw from their normal distributions. Then, given the extra information, the algorithm is faced with the following inference problem: The constants $\mu_1 \leq \mu_2, \cdots \leq \mu_n$ are given. We have n independent random variables such that, for each i, one of these random variables is normal with mean μ_i and standard deviation 1, but we have no knowledge as to which random variable has which mean. We want to determine, with probability $1 - \Delta$, which random variable has mean μ_1. We proceed in rounds where, in each round, we draw a sample from each of the distributions. It is easy to show that, for any stopping rule, the expected number of rounds is $\Omega(LB)$. The theorem follows from this fact.

In the full paper, we show that, even in the case where the values for the samples are given, rather than merely the outcomes of the games, any stopping rule requires $T = \Omega(\frac{\sum_{i=1}^{\lg(n)} \frac{1}{\epsilon_i^2}}{\lg(1+\frac{\lg(n)}{\lg(\frac{1}{\Delta})})})$ samples from each distribution, on the average, to identify the distribution of minimum mean with confidence $1 - \Delta$.

Corollary 2 *The lower bound of Theorem 1 holds in the strongly transitive and adversary models.*

4 An Upper Bound on $T_{ADV,K}$

In this section we assume known win probabilities. We will determine the best player with error $\leq \Delta$. The adversary can answer in any way he likes, provided that when 1 plays j, 1 wins with probability p_{1j}. Our algorithm is motivated by the Swiss System, a widely used method of pairing chess players.

Let k be a number such that $2(480 \lg(k) \lg(\frac{1}{\Delta}))^2 + 240 \lg k \lg \frac{1}{\Delta} = \sqrt{k}$.

```
Program Best-Player-Against-
      Adversary(n, Δ)
Do until ≤ k players left:
      Let m = number of players   .
   Assign all players a win
      score of 0.
   Let phase = 1
   Do while phase
         ≤ t = 480(lg m) lg(1/Δ)
      Pair up randomly all
         players with an equal
         win score, have each
         pair play 3 1/ε²_{(lg m)/4}
         games, and declare as
         winner the member of
         each pair who has won
         the majority of the
         games.
      The winners of the phase
         increase their win
         score by 1.
      The losers in the phase
         keep the same win score.
      Any odd (unpaired) player
         is assumed to have won
         the phase.
      phase = phase + 1
   Throw out all players whose
      win score < t/2 + (√3/√2)√(lg m)√t,
      except for those players
      who were ever in a win
      score category containing
      fewer than 2 ⁴√m players
      during this stage.
Run a standard knockout tournament
   among the remaining k players.
```

In the comments below, we'll use the term *stage* to describe one iteration of the outermost loop. We denote the number of players remaining at the beginning of stage i by m_i, and we denote the number of phases in stage i by $t_i = 480(\lg m_i) \lg(\frac{1}{\Delta})$.

Theorem 3 *The above algorithm runs in* $11520 \lg \frac{1}{\Delta} \sum_{i=1}^{\lg n} \frac{1}{\epsilon_i^2} + 864 \frac{1}{\epsilon_1^2} \lg \frac{1}{\Delta} \lg \lg \frac{1}{\Delta}$ *rounds and the probability it fails to determine the best player in the presence of a malicious adversary is at most* Δ.

Proof. (Sketch)
We establish the following lemmas:

- In a given stage i, the probability that the best player fails to survive for the next stage is at most $\Delta^{\lg(m_i)}$.
- For all w, in stage i, after t_i phases, the number of players with exactly w wins is at most

$$\frac{1}{2^{t_i}}\left(\begin{array}{c} t_i \\ w \end{array}\right) m_i + 1$$

- In any stage i, the number of players who survive stage i is at most $\sqrt{m_i}$. So the total number of stages is at most $\lg\lg(n)$.

When we're down to k players, we run a knockout tournament. A knockout tournament involves pairing up the k players and then playing T games between each pair. The member of each pair with the smaller number of wins is then thrown out, and the process is repeated with the $\frac{k}{2}$ remaining players. We will use $T = \frac{3}{2}\lg(\frac{2\lg k}{\Delta})\frac{1}{\epsilon_1^2}$ in our knockout tournament.

We show:

- Pr[Best player is killed off during knockout tournament] $\leq \frac{\Delta}{2}$
- The number of rounds for the knockout tournament $= T\lg k \leq 864\frac{1}{\epsilon_1^2}\lg\frac{1}{\Delta}\lg\lg\frac{1}{\Delta}$.

5 Better Upper Bound on $T_{DISC,K}$

Let us say that the matrix (p_{ij}) is *discriminating* if the following inequality holds whenever $j > i$: $\frac{p_{k+1,j}}{p_{k+1,i}} \geq \frac{p_{kj}}{p_{ki}}$. This inequality may be interpreted as follows: if Player i is stronger than Player j, then the ratio of i's loss probability to j's loss probability against a common opponent is an increasing function of the common opponent's strength. That is, the weader the common opponent k, the greater the significance of a loss to Player k in distinguishing between Player i and Player j.

We mention two commonly used models that lead to discriminating matrices.

The Bradley-Terry Model[Br] Player i has a *strength* π_i, where $\pi_1 \geq \pi_2 \geq \cdots \geq \pi_n$, and $p_{ij} = \frac{\pi_i}{\pi_i+\pi_j}$. This model applies when each player is a Geiger counter and, in any game, the first counter to click wins.

The Thurstone-Mosteller Model[Th] Player i has a *strength* s_i, and his *performance* in any given match is a random variable drawn from the normal distribution with mean s_i and standard deviation 1. When two players are matched, the one with the higher performance is the winner. It follows that p_{ij} is just the probability that a normal random variable with variance 2 is less than or equal to $s_i - s_j$. Thus $p_{ij} = \frac{1}{2\sqrt{2\pi}}\int_{-\infty}^{s_i-s_j} e^{-\frac{x^2}{4}}dx$, and a brief calculation shows that (p_{ij}) is discriminating.

Throughout this section we assume that (p_{ij}) is discriminating.

We consider the following algorithm:

```
Program Best-Player-Discriminating-
   Matrix-Algorithm-A(players, Δ)
   Initially, S = {1, 2, ···, n}.
   Do until |S| = 1
      Pair up the players in S
         randomly and play a game
         between each pair.
      Delete from S all players
         with at least T losses.
   The sole remaining player is
      declared the champion.
```

Theorem 4 *The number of rounds required by Algorithm A to determine the best player with probability* $1-\Delta$ *is* $O(LB\lg n)$*, where LB is the lower bound from Section 3.*

Proof. (Sketch)

- We use Azuma's Martingale Tail Inequality to show that the probability that Player 1 gets eliminated before Player j does is at most $e^{-\frac{a_{1j}^2 T}{2}}$.
- We note that the number of days cannot exceed $T\lg n$ and the number of games cannot exceed Tn.
- We show that in order to get $\sum_{j=2}^{n} e^{-\frac{a_{1j}^2 T}{2}} < \Delta$, it suffices that $T = c \cdot LB$.

5.1 The Case of Closely Matched Players

We continue to assume that the matrix (p_{ij}) is discriminating. We consider the following very simple algorithm.

```
Program Best-Player-Discriminating-
   Matrix-Algorithm-B(players, Δ)
   Do until n − 1 players have
         accrued at least T losses:
      On each day, pair the players
         randomly and play a game
         between each pair.
   Declare as winner any player who
      has the minimum number of
      losses.
```

The analysis of Algorithm A applies to this algorithm as well and shows that we can choose $T = O(T^*)$. We shall show that the number of days required for Algorithm B will be $O(\max(\lg n, T))$ with high probability, rather than $O(T \lg n)$, provided that the players are fairly evenly matched, in the sense that each player has a probability of at least q of losing against a random player, where q is a fixed positive constant. Thus we suppose that $\frac{1}{n-1}\sum_{i=2}^{n} p_{1i} \geq q$.

Theorem 5 *If Player 1's probability of losing against a random player is bounded below by a positive constant q, then Algorithm B requires $O(LB)$ rounds to determine the best player with probability $1-\Delta$, where LB is the lower bound from Section 3 .*

Proof. (Sketch) We show by Chernoff bounds that there exists a constant a such that if competitions were allowed to run for $an(\max(\lg n, LB)$ days, then with probability $1-\Delta$, all players would accrue at least T losses.

Theorem 6 *The condition that Player 1's probability of losing against a random player is bounded below by a positive constant q can be achieved in $O(\lg^2 n \ln(\frac{1}{\Delta}))$ rounds with probability $1-\Delta$.*

Proof. (Sketch) As long as Player 1's average probability of losing against the remaining players is less than q, then we can eliminate a constant fraction of the remaining players in a subtournament of $\lg n$ rounds.

6 A Lower Bound on $T_{ADV,U}$

We will show that, when Players $2 \ldots n$ are allowed to choose adaptively their winning probabilities amongst each other and when we do not know the best player's winning probabilities, we can not efficiently determine the best player. This corresponds to a situation where many dishonest, possibly mediocre, players are trying to prevent us from learning the identity of the honest best player.

Theorem 7 *There exists no algorithm which, in the presence of a malicious adversary and without knowing the best player's winning probabilities, can identify the best player, with confidence $1-\Delta$ for $\Delta < \frac{1}{3}$, within $\frac{n}{16}\frac{1}{\epsilon_1^2}$ rounds for all values of n and ϵ_1.*

Proof. (Sketch)
Suppose there exists an algorithm A which will output the best player, with probability at least $1-\Delta$, within $\frac{n}{16}\frac{1}{\epsilon_1^2}$ rounds.
We consider two scenarios:

1. First, we consider a situation where Player 1's winning probabilities against players $2 \ldots n$ are all equal to 1 and players $2 \ldots n$ having winning probability $\frac{1}{2}$ amongst each other. In this scenario, algorithm A will output Player 1 within $\frac{n}{4}$ rounds with probability at least $1-\Delta$.

2. We now consider a situation where Player 1's advantages over players $2 \ldots n$ are the same and equal to $\gamma = \epsilon_{\lg(n)} < \frac{1}{4n}$.

   The adversary's strategy is to choose a second best player, Player 2, and give that player probability 1 of beating players $3 \ldots n$. For $3 \leq i < j \leq n$, the adversary lets Player i beat Player j with probability $\frac{1}{2}$. So long as the best player can not be discriminated from players $3 \ldots n$, the algorithm will announce (incorrectly) that Player 2 is the best player within $\frac{n}{4}$ rounds.

7 Upper Bound on $T_{TRANS,U}$ for only Two Players

Throughout this section we assume the strong transitivity model.

7.1 Playing the Two Players Against Each Other

In this subsection we consider the following problem: *Given two unequal players with unknown relative strengths, determine, quickly and with high probability, which player is the better player.* In this subsection, we will show that one can determine the better player correctly within $\frac{12}{\epsilon^2}(2\lg\lg(\frac{1}{\epsilon}) + \lg(\frac{4}{\Delta}))$ games between the two players, with probability at least $1-\Delta$, where ϵ is the advantage of one player over the other. Note that ϵ is *not* part of the given information.

We start by observing that the above problem is the same as the problem of determining whether a biased coin is biased up or down, where we are not given the bias, ϵ, of the coin. Our goal for the coin is to determine the correct answer as quickly as possible with probability $1-\Delta$.

This problem was first addressed in [Fa] where Farrell proved that

$$limsup_{\epsilon\to 0}\epsilon^2 \cdot \frac{1}{\lg\lg\epsilon} \cdot (\text{Expected Number of flips}) \geq c,$$

where c is a constant.

Our algorithm works by guessing different values for the bias of the coin. In stage i, the guess is that the bias is $\epsilon_i = \frac{1}{2^i}$. Given this guess, we flip the coin T_i times. If the coin behaves as though it has bias $+\epsilon_i$ during these T_i flips, we output that the bias is UP. If the coin behaves as though it has bias $-\epsilon_i$ during these T_i flips, we output that the bias is DOWN. Otherwise, we proceed with stage $i+1$, in which the guess of the bias is halved.

```
Program
Determine Coin Bias Direction(Δ)
   Initially i = 1.
   Let ε_i = 1/2^i.
   Let T_i = 3/2 · 1/ε_i^2 ln(4i^2/Δ).
   Flip coin T_i times.
   If number of heads ≥ T_i(1/2 + ε_i),
         output:  UP.
   If number of heads ≤ T_i(1/2 − ε_i),
         output:  DOWN.
   Else:  i = i + 1.  Return to step 1.
```

Theorem 8 *With probability at least $1-\Delta$, the above algorithm outputs the correct direction of bias for a coin of bias $\pm\epsilon$ using at most $\frac{12}{\epsilon^2}(2\lg\lg(\frac{1}{\epsilon}) + \lg(\frac{4}{\Delta}))$ coin flips.*

Proof. (Sketch)

There are two sources of error in the above algorithm.

1. The algorithm returns an incorrect output.
2. The algorithm requires more than $\lg(\frac{1}{\epsilon})$ stages.

To bound error of the first type, we use Chernoff bounds to show that the error in stage i (i.e. the probability that the algorithm answers UP to a coin which is biased DOWN, or vice-versa, during stage i) is $\leq \frac{\Delta}{2i^2}$. Consequently, the total error over all stages of the algorithm is $\leq \frac{\Delta}{4}\sum_{i=1}^{\infty}\frac{1}{i^2} \leq \frac{\Delta}{2}$.

To bound error of the second type we use Chernoff bounds to show that the probability that the number of stages exceeds $\lg(\frac{1}{\epsilon})+1$ is less than $\frac{\Delta}{2}$.

7.2 Playing the Two Players Against Other Players

In this subsection we are given n players: two distinguished players, A and B, and $n-2$ other players, *other_players*. We are not given any information about the relative strengths of any of the players. Our goal is to determine which of A and B is the better player. We present an algorithm that determines which of A and B is the better player, with probability $1-\Delta$, by having A and B each play $\frac{12}{\epsilon^2}(2\lg\lg(\frac{1}{\epsilon}) + \lg(\frac{4}{\Delta}))$ rounds against random opponents chosen from among *other_players* and where ϵ is the difference in the average winning probabilities of A and B. Once again, we note that ϵ is *not* part of the given information.

```
Program Best-of-2-Players
            (A, B, other_players, Δ)
   Initially i = 1.
   Let ε_i = 1/2^i.
   Let T_i = 3/2 · 1/ε_i^2 ln(2i^2/Δ).
   A and B each play T_i rounds,
      where in each round,
      A and B play different
      randomly chosen opponents.
   If (A's wins - B's wins) ≥ T_i ε_i,
         output:  A:  better player
   If (B's wins - A's wins) ≥ T_i ε_i,
         output:  B:  better player
   Else:  i = i + 1.  Return to step 1.
```

Theorem 9 *With probability at least $1-\Delta$, the above algorithm outputs the identity of the better player within at most $\frac{12}{\epsilon^2}(2\lg\lg(\frac{1}{\epsilon}) + \lg(\frac{4}{\Delta}))$ rounds, where the average winning probability of the better player is ϵ + the average winning probability of the worse player.*

Proof.

The proof follows the same structure as that of the previous subsection, except that a Chernoff bound for sums of identically distributed 3-valued random variables, rather than for Bernoulli trials, is required because A and B are distinguished based on their performance against other players.

8 An Upper Bound on $T_{TRANS,U}$

Throughout this section we assume the strong transitivity model.

We assume that we are given n players with no information about their relative strengths, and we must determine, as quickly as possible, which player is best with confidence at least $1-\Delta$.

The algorithm which we describe runs in $\lg(n)$ stages. In stage i, we effectively run $\binom{2^{\lg(n)-i+1}}{2}$ **Best-of-2-Players** algorithms in parallel until we can eliminate half the remaining players.

```
Program Best-Player-Unknown-Strengths
                        (players, Δ)
   Do for i = 1 to lg(n):  (i designates
               the stage number).
      Initialize each remaining
         player's record at 0 wins
         and 0 losses.
      Play randomly-paired rounds
         among all remaining players.
      After each round, Player y is
         said to be dominated if
         there exists some player
         z_y, such that the program
         Best-of-2-Players
              (y, z_y, players − {y, z_y}, Δ/4n)
         would at this point declare z_y
         to be a better player than y.
      If at least half of all players
         are dominated, then remove
         all dominated players.
   Output the one remaining player.
```

Theorem 10 *With probability at least $1-\Delta$, **Best-Player-Unknown-Strengths**(players, Δ) outputs the best player within $\sum_{i=1}^{\lg(n)} \frac{12}{\epsilon_i^2}(2\lg\lg(\frac{1}{\epsilon_i}) + \lg(\frac{16n}{\Delta}))$ rounds.*

Proof. (Sketch) There are two sources of error

- In some stage, more than $R_i = \frac{12}{\epsilon_{\lg n-i-1}^2}(2\lg\lg(\frac{1}{2\epsilon_{\lg n-i-1}}) + \ln(\frac{16n}{\Delta}))$ rounds are played
- Player 1 gets eliminated

We show:

1. In stage i, $Pr[<\frac{1}{2}$ of all remaining players are eliminated within R_i rounds$]\leq \frac{\Delta}{2^i}\cdot\frac{1}{4}$. This is shown by observing that the best player's average winning probability over the bottom half of all remaining players is high.
2. $Pr[$ Player 1 is eliminated in stage $i] \leq \frac{\Delta}{2^i}\cdot\frac{1}{2}$.

9 Future Work

The functions T_{ADV}, T_{TRANS} and T_{DISC} are known only up to a factor of $\Theta(\lg\lg n)$. It would be of interest to determine their precise growth rates or, at least, to determine whether they all grow at the same rate.

We suspect that the winner of a playoff or a tournament is often not the best player, but is seldom among the weaker players. Thus, it would be of interest to study within, say, the strong transitivity model, the complexity of playoff systems which, with high probability, select as champion a player who is not necessarily the best, but is among the k best.

Throughout this paper we have assumed that a player cannot be replicated; i.e., that he or she can participate in only one game per round. The case of replicatable players is also of interest. Consider, for example, the problem of selecting the best of n chess programs where, on any day, m games can be played concurrently, but there is no restriction on the number of games in which a given program may participate. For this problem it would be of interest to analyze "survival of the fittest" strategies, in which each program has a "weight" which grows when it wins and shrinks when it loses, and the number of games in which a program participates on a given day is proportional to its weight.

We would like to thank Sridhar Rajagopalan, Troy Shahoumian, David Freedman and Manuel Blum for helpful conversations and useful suggestions.

References

[AD] J. A. Aslam and A. Dhagat. "Searching in the presence of linearly bounded errors", *Symp. on Theory of Computing* 1991, 486-493.

[Br] R. A. Bradley. "Some Statistical Methods In Taste Testing and Quality Evaluation". *Biometrics*, 1953, 9, pp. 22- 38.

[CheH] R. Chen and F. K. Hwang. "Stronger players win more balanced knockout tournaments". *Graphs and Combinatorics* 4, 95-99 (1988).

[ChuH] F. K. Chung and F. K. Hwang. "Do Stronger Players Win More Knockout Tournaments?" *J. of the Amer. Statistical Assoc.*, 1978, (73) 363, Theory and Methods Section.

[D] H. A. David. "The Method of Paired Comparisons". *Griffin and Co., Oxford University Press* (1963).

[Fa] R. H. Farrell. "Asymptotic Behavior of Expected Sample Size in Certain One Sided Tests". *Annals of Mathematical Statistics*, (35) 1964, p.36-72.

[FPRU] U. Feige, D. Peleg, P. Raghavan and E. Upfal. "Computing with Unreliable Information". In *Symposium on Theory of Computing*, 1990, 128-137.

[HM] B. R. Handa and V. Maitri. "On a knockout selection procedure". *Sankhya: the Indian J. of Statistics*, 1984, 46, A, Pt 2, 267-276.

[Hw] F. K. Hwang, Z. Z. Lin, Y. C. Yao. "Knockout tournaments with diluted Bradley-Terry preference schemes". *J. Statist. Planning and Inference*, 1991, 28, 99-106.

[I] R. B. Israel. "Stronger players need not win more knockout tournaments". *J. of the Amer. Statist. Association*, 1981, 76, 376, Theory and Models Section.

[P] A. Pelc. "Searching with known error probability". *Theoretical Computer Science* 63: 185-202, 1989.

[RGL] B. Ravikumar, K. Ganesan and K. B. Lakshmanan. " On selecting the largest element in spite of erroneous information ". *Lecture Notes in Computer Science, ICALP* 1987, 88-99.

[RMKWS] R. L. Rivest, A. R. Meyer, D. J. Kleitman, K. Winklmann and J. Spencer. "Coping with errors in binary search procedures", *Journal of Computer and System Sciences*, 20: 396-404, 1980.

[Th] L. L. Thurstone. "The Method of Paired Comparisons for Social Values," *J. Abnorm. Soc. Psychol.*, 1927, 21, 384- 400.