Adversarial Leakage in Games

Noga Alon* Michal Feldman† Yuval Emek‡ Moshe Tennenholtz§

Abstract

While the minimax strategy has become the standard, and most agreed-upon solution for decision-making in adversarial settings, as discussed in game theory, computer science and other disciplines, its power arises from the use of mixed strategies, aka probabilistic algorithms. Nevertheless, in adversarial settings we face the risk of information leakage about the actual strategy instantiation. Hence, real robust algorithms should take information leakage into account. To address this fundamental issue, in this paper we introduce the study of *adversarial leakage* in games. We consider two models of leakage. In both of them the adversary is able to learn the value of b binary predicates about the strategy instantiation. In one of the models these predicates are selected after the decision-maker announces its probabilistic algorithm, and in the other one they are decided in advance. We give tight results about the effects of adversarial leakage in general zero-sum games with binary payoffs, as a function of the level of leakage captured by b in both models. We also compare the power of adversarial leakage in the two models, and the robustness of the original minmax strategies of games to adversarial leakage. Finally, we study the computation of optimal strategies for adversarial leakage models. Together, our study introduces a new framework for robust decision-making, and provides rigorous fundamental understanding of its properties.

1 Introduction

Decision-Making lies in the foundations of fields such as Economics, Operations Research, and Artificial Intelligence. The question of what should be the action to be taken by a decision-maker when facing an uncertain environment, potentially consisting of other decision makers, is a fundamental problem which led to a wide variety of models and solutions. The only type of situations for which this question got an agreed-upon answer is in the context of two-player zero-sum games. This setting can model any situation in which a decision-maker aims at maximizing his guaranteed payoff. When mixed strategies are allowed, such desired behavior, termed an agent's minimax (or safety level) strategy, leads to a well defined expected utility (known as the *value* of the game). Moreover, when presented explicitly in a matrix form the computation of a minimax strategy is polynomial (by solving a linear program). Various equilibrium concepts have been considered in the game-theoretic literature, but none of them provides a prescriptive advice to a decision-maker which will be as acceptable as the safety-level strategy solution in adversarial settings. Moreover, the safety-level strategy has been advocated for some non zero-sum settings as well.

*Tel-Aviv University, and Microsoft Israel R&D Center.

†School of Business Administration, The Hebrew University of Jerusalem, and Microsoft Israel R&D Center.

‡Microsoft Israel R&D Center.

§Microsoft Israel R&D Center, and Technion-Israel Institute of Technology.

Much of the power of a minimax strategy is associated with the use of mixed strategies, aka randomized algorithms. In such algorithms the randomization phase is assumed to be done in a private manner by the decision-maker, and no information about the instantiation selected in that phase is assumed to be revealed. However, in reality nothing is really private; for example, competitors will always strive to obtain the private actions of a business; hence, information leakage should be considered. As a result, it may be of interest to study the effects of adversarial leakage, where a limited amount of information on an agent's instantiation of its mixed strategy may leak in an adversarial manner. We believe that only by considering this situation, it will be possible to construct robust strategies when acting in an adversarial setting.

Our model of adversarial leakage is general. We consider a two-player zero-sum game in strategic form (aka matrix form), where the MAX player is our decision-maker and the MIN player is the adversary. Both MAX and MIN have a set of (pure) strategies they can choose from. MAX chooses a mixed strategy, that is, a probability distribution over its pure strategies. The MIN player may base its action on the value of b binary predicates defined on MAX' pure strategies; each such predicate is a boolean formula on the set of strategies, whose value is determined according to the actual instantiation of MAX' mixed strategy. The parameter b determines the level of information leakage; MAX would like to maximize his guaranteed expected payoff against any choice of such b predicates.

We consider two settings, distinguished by the information structure assumed in them. In the **Strong Model**, the MAX player chooses a mixed strategy, which is observable by the MIN player, who can then act upon it in determining the b predicates. In the **Weak Model**, on the other hand, the MIN player chooses the b predicates first, and MAX can observe it and act upon it in choosing his mixed strategy.

The information structure of the Weak Model gives the MAX player a potential strategic advantage in the game. Quantifying this advantage is one of our objectives in this paper. This question is particularly interesting in light of the minimax theorem, which essentially states that in a standard two-player zero-sum game, if mixed strategies are allowed, gaining information about the opponent's mixed strategy prior to playing does not give the agent possessing it any strategic advantage.

Other intriguing questions arise in this setting of adversarial leakage. What would be the best mixed strategy for the MAX player? How well will the original minimax strategy of the game perform? What is the computational complexity of finding the optimal strategy under information leakage? We address all these questions, focusing our attention on general two-person games, where the decision-maker has m strategies to choose from and the adversary has n strategies to choose form, and the decision maker payoffs are either 1 or 0. This is known to be a highly applicable model, as it captures games in which a goal is either achieved or not.

Our results : For the Strong Model, if the value of the game is $q = 1 - \epsilon$ (for small positive ϵ) and 2^b is much smaller than $1/\epsilon$, then MAX can ensure value close to 1 (at least $1 - 2^b\epsilon$), and this is tight. To do so, she simply uses the optimal mixed strategy for the usual game, with no communication. On the other hand, if 2^b is much bigger than $1/\epsilon$, for every mixed strategy of MAX the MIN player can ensure value close to zero (at most $e^{-2^b\epsilon}$.) Therefore, for EVERY such game with value $1 - \epsilon$, which is close to 1, a sharp transition occurs at b which is about $\log(1/\epsilon)$. If b is slightly smaller, the value stays close to 1, if it slightly larger, the value drops to nearly zero.

For games with value q bounded away from 1, even one bit enables MIN to square the value and

drop it to at most q^2, and every additional bit squares the value again. There are also examples showing that this is essentially tight. For any fixed value $q < 1$, $\log\log m + O_q(1)$ bits suffice to enable MIN to drop the value to precisely 0.
For the Weak Model, the situation is different. Clearly, here MAX is in a better shape, hence if the value of the game is $q = 1 - \epsilon$, MAX can still ensure a value close to 1 if the number of bits is much smaller than $\log(1/\epsilon)$. Here, however, there are examples in which she can do much better, and in fact can ensure no essential drop in the value as long as the number of communication bits is somewhat smaller than $\log\log m$. More precisely, for any fixed value q, $0 < q < 1$ and every large polynomially related n, m, there are examples of games represented by a binary m by n payoff matrix with value $q + o(1)$, so that even if the number of communication bits b is $\log\log m - O(1)$, MAX can ensure that the value will stay roughly q.

Somewhat surprisingly, once b is a bit bigger, that is, $b = \log\log m + O(1)$, the MIN player can already ensure value 0 in any game with a fixed value $q < 1$. Thus, in the examples above, nearly $\log\log m$ bits have essentially no effect on the value (in the Weak Model), while slightly more bits already suffice to drop the value to 0.

Note that, in contrast to leakage-free games, where no advantage is gained by observing the opponent's mixed strategy prior to playing (due to the minimax theorem), in settings of adversarial leakage, such information can contribute a great deal to the informed player. In particular, when the number of communication bits is somewhat smaller than $\log\log m$, there are examples of games with value $q + o(1)$, so that in the Strong Model, every additional bit squares the value of MAX, while in the Weak Model, MAX can ensure that the value will stay roughly the same.

With respect to computation complexity, computing the optimal strategy in the Strong Model (for the MAX player) against b bits is poly-time for any fixed b, while this problem becomes NP-hard to compute, or even to approximate within any factor, for a general b. In the Weak Model, the optimal strategy of MAX can be computed in polynomial time for every b. As for the MIN player, computing the optimal predicates is polynomial for a fixed number of bits, but is NP-hard in general.

2 Model

We consider two-player zero-sum games defined by m by n matrices with $\{0,1\}$-entries, where m, n are polynomially related (that is, $\log m = \Theta(\log n)$). The row player is the MAX player and the column player is the MIN player. $M_{i,j}$ is the payoff of the MAX player if MAX and MIN play row i and column j, respectively. The payoff of the MIN player is then $-M_{i,j}$. The matrix M is known for both players[1].

Given a matrix M and an integer $b \geq 0$, we describe a precise setting of adversarial leakage, as follows.

1. The MAX player chooses a distribution vector $\mathbf{p} = (p_1, \ldots, p_m)$ on $[m]$.
2. The MIN player chooses a b-bit leakage function, $f : [m] \to \{0,1\}^b$.
3. The MAX player realizes $i \in [m]$ according to p (i.e., chooses row i with probability p_i).

[1]While we focus on the natural binary case, some of our results hold for any matrix with entries in $[0,1]$ as well, while some become non-interesting or easily seen to be false.

4. The MIN player observes $f(i)$ (for i realized by MAX) and chooses a strategy $j \in [n]$.

5. The MAX and the MIN players receive payoffs $M_{i,j}$ and $-M_{i,j}$, respectively.

The two settings we consider, referred to as the Strong Model and the Weak Model, are distinguished according to the order in which the first two steps take place. In the Strong Model, step 1 precedes step 2; i.e., the MAX player first chooses a distribution $\mathbf{p}$, and the MIN player, knowing the distribution $\mathbf{p}$, chooses the function f. In the Weak Model, in contrast, step 2 above comes before step 1; i.e., the MIN player first chooses a leakage function f, and only then, the MAX player chooses a distribution $\mathbf{p}$, knowing the mapping f.

In both models, the MIN player chooses two functions, namely a function $f : [m] \to \{0,1\}^b$ (with or without observing p, depending on the model), and a function $g : \{0,1\}^b \to [n]$ once observing $f(i)$.

We denote by $v_p(M,b)$ the worst-case expected payoff (from the perspective of the MAX player) over all possible functions f and g chosen by the MIN player with b bits, under the strategy p. When not clear in the context, we shall denote the values in the two models by $v_p^{(A)}(M,b)$ and $v_p^{(A)}(M,b)$, respectively.

We also denote by $\mathbf{p}_b^*$ an optimal strategy of the MAX player against b bits; i.e., $\mathbf{p}_b^* \in argmax_{p\in\Delta(m)} v_p(M,b)$, where $\Delta(m) = \{(p_1,\ldots,p_m) \in \mathbb{R}^m | \sum_i p_i = 1 \text{ and } p_i \geq 0 \ \forall i\}$. In addition, $v(M,0)$ will be used to denote the value of the game M (with no leakage), and $v(M,b)$ will be used to denote the optimal payoff of the MAX player under adversarial leakage of b bits.
All logarithms are in base 2, unless otherwise specified.

3 Adversarial leakage in the Strong Model

We first show that for any m by n matrix with $\{0,1\}$-entries of value $q = 1-\epsilon$, the MAX player can guarantee herself at least a payoff of $1-2^b\epsilon$. This can be done, in particular, by playing the minimax strategy.

Proposition 3.1. *Let M be an m by n matrix with $\{0,1\}$ entries. Let $q = 1-\epsilon$ be the value of the game defined by M, that is, $v(M,0) = 1-\epsilon$. Then, for every $b \geq 0$, $v_{\mathbf{p}_0^*}(M,b) \geq 1-2^b\epsilon$.*

Proof. For every $w \in \{0,1\}^b$, let $S^w = \{i | f(i) = w\}$, and let $p^w = \sum_{i\in S^w} p_i$. Fix some column j. Since $1-\epsilon$ is the value of the game, it holds that for every w,

$$\sum_{i\in S^w} p_i M_{i,j} + \sum_{i\notin S^w} p_i M_{i,j} \geq 1-\epsilon.$$

Since $M_{i,j} \leq 1$ for every i,j, we get:

$$\sum_{i\in S^w} p_i M_{i,j} + \sum_{i\in[m]\setminus S^w} p_i \geq 1-\epsilon.$$

Substituting $\sum_{i\in S^w} p_i = 1-p^w$ and rearranging the last inequality yields:

$$\sum_{i\in S^w} p_i M_{i,j} \geq p^w - \epsilon. \tag{1}$$

Let $g(w)$ be the column chosen by MIN (in step 4) in response to the word w. The expected payoff of MAX is then given by the expression $\sum_{w\in\{0,1\}^b}\sum_{i:f(i)=w} p_i \cdot M_{i,g(w)}$ and the expected payoff of MAX conditioned on the event that some row $i \in S^w$ is played is given by the expression $\sum_{i:f(i)=w} \frac{p_i}{p^w} \cdot M_{i,g(w)}$, which is at least $\frac{1}{p^w}(p^w - \epsilon)$, by Equation 1. Since some row $i \in S^w$ is played with probability p^w, we get that the expected utility of MAX under f is at least:

$$\sum_{w\in\{0,1\}^b} p^w \frac{1}{p^w}(p^w - \epsilon) = 1 - 2^b\epsilon.$$

□

The above bound is tight, as established by the following proposition:

Proposition 3.2. *For every $\epsilon > 0$ and every n, there exists an n by n matrix with $\{0,1\}$-entries of value $1-\epsilon$ such that $v_{\mathbf{p}_0^*}(M,b) = v_{\mathbf{p}_b^*}(M,b) = 1 - 2^b\epsilon$.*

Proof. Let $n = 1/\epsilon$ and consider the matrix M in which $M_{i,i} = 0$ for every i, and $M_{i,j} = 1$ for every $i \neq j$. From symmetry considerations, both the minimax strategy and the optimal strategy against b bits of the MAX player is the uniform distribution over the rows, that is, $\mathbf{p}_0^*(i) = \mathbf{p}_b^*(i) = 1/n = \epsilon$ for every i. Let f be a function which imposes the following partition on the rows: each one of the first $2^b - 1$ rows constitutes its own subset, and the remaining rows constitute the last subset. That is, the partition imposed by f is: $\{1\}, \{2\}, \ldots, \{2^b - 1\}, [m] - \{1, \ldots, 2^b - 1\}$. In this case, if one of the first $2^b - 1$ rows is chosen (each with probability ϵ), the MAX player's payoff is 0 (since the MIN player can choose a column j so that $M_{i,j} = 0$). If one of the remaining rows is chosen (with a total probability of $1 - (2^b - 1)\epsilon$), it is easy to see that the payoff obtained by the MAX player is $\frac{\frac{1}{\epsilon} - 2^b}{\frac{1}{\epsilon} - (2^b - 1)}$. The expected payoff of the MAX player under f by playing $\mathbf{p}_0^*$ or $\mathbf{p}_b^*$ (the uniform distribution) is therefore given by:

$$v_{\mathbf{p}_0^*}(M,b) = v_{\mathbf{p}_b^*}(M,b) = (2^b - 1)\epsilon \cdot 0 + \left(1 - (2^b - 1)\epsilon\right) \cdot \frac{\frac{1}{\epsilon} - 2^b}{\frac{1}{\epsilon} - (2^b - 1)} = 1 - 2^b\epsilon.$$

□

The above two propositions essentially say that for games with value $q = 1 - \epsilon$ and b such that $2^b\epsilon = o(1)$, MAX can guarantee a payoff of about q^{2^b} by playing the safety level, and this is optimal. The case of general q and b, however, requires more work, and this is the focus of the following statement.

We next show that for every strategy of MAX the MIN player can always ensure a payoff of at most q^{2^b}.

Theorem 3.3. *Let M be an m by n matrix with $\{0,1\}$ entries. Let q be the value of the game defined by M, that is, $q = v(M,0)$. Then, for every $b \geq 0$ and every distribution vector p of the MAX player, $v_p(M,b) \leq q^{2^b}$.*

Proof. Put $p^{(1)} = p$, and let $j_1 \in [n]$ be a pure strategy of MIN (a column of M) ensuring a value of $q_1 \leq q$ against the mixed strategy $p^{(1)}$ of MAX. Such a pure strategy must exist since q is the value of the game. Define $S_1 = \{i \in [m] \mid M_{i,j_1} = 0\}$. It holds that $\sum_{i\in S_1} p_i^{(1)} M_{i,j_1} + \sum_{i\in[m]-S_1} p_i^{(1)} M_{i,j_1} = q_1$, hence $\sum_{i\in[m]-S_1} p_i^{(1)} = q_1$.

Let $p^{(2)}$ be the probability distribution defined by restricting $p^{(1)}$ to the rows in $[m]-S_1$, namely,

$$p_i^{(2)} = \begin{cases} p_i^{(1)}/q_1 & \text{if } i \in [m] - S_1; \\ 0 & \text{otherwise.} \end{cases}$$

Let j_2 be a pure strategy of MIN ensuring a value of $q_2 \leq q$ against the mixed strategy $p^{(2)}$ of MAX. Once again, such a pure strategy must exist since q is the value of the game. Define $S_2 = \{i \in [m] - S_1 \mid M_{i,j_2} = 0\}$. As before, it holds that $\sum_{i \in S_2} p_i^{(2)} M_{i,j_2} + \sum_{i \in [m]-S_1-S_2} p_i^{(2)} M_{i,j_2} = q_2$, hence $\sum_{i \in [m]-S_1-S_2} p_i^{(2)} = q_2$.

Continuing in this manner for 2^b steps, we obtain a partition of $[m]$ to 2^b subsets $S_1, \ldots, S_{2^b}$ with corresponding columns $j_1, \ldots, j_{2^b}$ such that $M_{i,j_k} = 0$ for every $1 \leq k \leq 2^b$ and $i \in S_k$. For convenience we index the words in $\{0,1\}^b$ by $w_1, \ldots, w_{2^b}$ and fix

$$f(i) = w_k \text{ for every } 1 \leq k \leq 2^b \text{ and } i \in S_k \ ;$$

$$g(w_k) = j_k \text{ for every } 1 \leq k \leq 2^b \ .$$

The above construction guarantees that when MAX plays according to p, and MIN follows f and g, the payoff is 0 with probability

$$\begin{aligned} & 1 - q_1 + q_1 \left(1 - q_2 + q_2 \left(1 - q_3 + q_3 \left(\cdots \left(1 - q_{2^b}\right) \cdots\right)\right)\right) \\ = \ & 1 - q_1 + q_1 - q_1 q_2 + q_1 q_2 - q_1 q_2 q_3 + q_1 q_2 q_3 - \cdots - q_1 \cdots q_{2^b} = 1 - q_1 \cdots q_{2^b} \ . \end{aligned}$$

It follows that $v_p(M,b) = q_1 \cdots q_{2^b} \leq q^{2^b}$ as required. □

As a corollary of Theorem 3.3, we get the following.

Corollary 3.4. *Let M be an m by n matrix with $\{0,1\}$ entries, and let q be the value of the game defined by M, as above. If $q^{2^b} < 1/m$, then for every distribution vector p of the MAX player, $v_p(M,b) = 0$. Therefore, for every fixed q satisfying $0 < q < 1$, $b = \log\log m + O_q(1)$ suffices to ensure value 0 for the MIN player.*

Proof. Let p be the uniform distribution on $[m]$. Theorem 3.3 guarantees the existence of $f : [m] \to \{0,1\}^b$ and $g : \{0,1\}^b \to [n]$ such that if MAX plays according to p and MIN follows f and g, then the expected payoff is at most $q^{2^b} < 1/m$. We argue that if MIN follows f and g, then in fact, the expected payoff is 0. Indeed, since $p_i = 1/m$ for every $i \in [m]$, a positive expected payoff is possible only if it is at least $1/m$, which derives a contradiction. It follows that for every $w \in \{0,1\}^b$ and for every $i \in [m]$ such that $f(i) = w$, we must have $M_{i,g(w)} = 0$. But this means that the same f and g guarantees that the expected payoff is 0 regardless[2] of the mixed strategy of MAX. □

Remark: The corollary is essentially the known simple fact (proved in [4], [5]) that the ratio between the fractional cover and the integer cover of a hypergraph with m edges is at most $\ln m$.

The following theorem shows that both Theorem 3.3 and Corollary 3.4 are essentially tight.

[2] Note that as here the same f and g work against any mixed strategy of MAX, this works in Model B as well, providing a proof of Theorem 4.5.

Theorem 3.5. *For every real $0 < q < 1$, for every integer $b \geq 0$ and for every large polynomially related m and n satisfying $q^{2^b} m > 2^b \log n$, there exists an m by n $\{0,1\}$-matrix M that satisfies*
(i) $v(M,0) = q \pm o(1)$, where the $o(1)$-term tends to 0 as m and n grow; and
(ii) if $p = (p_1, p_2, \ldots, p_m)$ is the uniform distribution on the rows, then $v_p(M,b) \geq (1 - o(1))q^{2^b}$ (and thus $v_p(M,b) = (1 \pm o(1))q^{2^b}$, by Theorem 3.3).
In particular, for, say, $m = n^2$ and $b \leq \log\log m - \Theta(1)$, $v_p(M,b) > 0$.

Proof. Let M be a random m by n matrix with $\{0,1\}$-entries obtained by choosing each entry $M_{i,j}$, randomly and independently, to be 1 with probability q and 0 with probability $1 - q$. We show that M satisfies the assertion of the theorem with positive probability.

Since m, n are large and are polynomially related, almost surely (that is, with probability that tends to 1 as m, n tend to infinity) every row of M has $(1 \pm o(1))qn$ 1-entries, and every column of M has $(1 \pm o(1))qm$ 1 entries. This follows easily by the standard known estimates for Binomial distributions, see, for example, [3]. This implies that the value of the game is $(1 \pm o(1))q$: indeed, if MAX (respectively, MIN) plays according to the uniform distribution on the rows (resp., columns), then it guarantees an expected payoff of at least (resp., at most) $(1 \pm o(1))q$. Thus (i) holds almost surely.

We establish the assertion by showing that (ii) holds with a positive probability (bounded away from 0). For that purpose, we argue that for every choice of a set $J \subset [n]$ of size $|J| = 2^b$, the number of indices $i \in [m]$ so that $M_{i,j} = 1$ for all $j \in J$ is, almost surely, $(1 \pm o(1))q^{2^b} m$. Indeed, for a fixed choice of a set J, the random variable X that counts the number of such indices i is a Binomial random variable with parameters m and q^{2^b}. Therefore the probability that X is not $(1 \pm o(1))q^{2^b} m$ decreases exponentially with $q^{2^b} m > 2^b \log n = \log\left(n^{2^b}\right)$. The assertion is established by union bound over all $\binom{n}{2^b} < n^{2^b}$ choices of the set J. □

Remark: The proof of existence of M as in the last theorem is probabilistic. In Appendix A we give similar explicit examples using either finite geometries, or character sum estimates (Weil's Theorem [7]), or any example of small sample spaces supporting nearly 2^b-wise independent random variables. See [6] or [1] for some such examples.

4 Adversarial leakage in the Weak Model

The following result deals with the Weak Model. It shows that in sharp contrast to the situation with Model A, here there are examples in which $\log\log m - O(1)$ bits of information do not enable the MIN player to gain any significant advantage.

Theorem 4.1. *For every real q, $0 < q < 1$, for every positive δ and for all large polynomially related n, m satisfying*

$$[(\frac{10}{\delta})]^{n2^b} < \delta^{m/\sqrt{n}} \quad \textit{and} \quad [\frac{q(1-q)}{10}]^{2^b} \geq \frac{1}{\sqrt{n}},$$

there is an m by n matrix M with $\{0,1\}$-entries so that the value of the game it determines $v(M,0)$ is $q + o(1)$, and $v^{(B)}(M,b) \geq q - \delta$.

In particular, if $m = n^2$ and $b = \log\log m - \Theta(1)$, $v^{(B)}(M,b)$ is essentially equal to $v(M,0)$, where $v^{(B)}$ denotes the value under the Weak Model.

The proof of the above theorem is more complicated than the ones in the previous section, and requires several preparations. We need the following known result.

Lemma 4.2 ([2], Lemma 3.2)**.** *Let Y be a random variable with expectation $E[Y] = 0$, variance $E[Y^2]$ and fourth moment $E[Y^4] \leq k(E[Y^2])^2$, where k is a positive real. Then Prob$[Y \geq 0] \geq \frac{1}{2^{4/3}k}$.*

Using the above lemma, we prove the following.

Lemma 4.3. *Let q be a real, $0 < q < 1$, and let $p = (p_1, p_2, \ldots, p_n)$ be a distribution vector on $[n]$, that is, $p_j \geq 0$ for all j and $\sum_j p_j = 1$. Let $X_1, X_2, \ldots, X_n$ be independent, identically distributed indicator random variables, where each X_j is 1 with probability q (and 0 with probability $1 - q$.) Define $X = \sum_{j=1}^{n} X_j p_j$. Then the probability that X is at least its expectation (which is q) is bigger than $\frac{q(1-q)}{10}$.*

Proof. Define $Y_j = X_j p_j - E[X_j p_j] = X_j p_j - q p_j$, and $Y = \sum_j Y_j$. By linearity of expectation $Y = X - E[X]$, and $E[Y] = 0$. In order to apply the previous lemma, we compute the variance of Y, and estimate its forth moment.

$$Var[Y] = \sum_j Var[Y_j] = \sum_j [q(1-q)^2 p_j^2 + (1-q) q^2 p_j^2]$$

$$= q(1-q) \sum_j p_j^2.$$

Similarly

$$E[Y^4] = \sum_j E[Y_j^4] + 6 \sum_{i<j} E[Y_i^2] E[Y_j^2]$$

$$= \sum_j [q(1-q)^4 p_j^4 + (1-q) q^4 p_j^4] + 6 \sum_{i<j} q^2 (1-q)^2 p_i^2 p_j^2$$

$$\leq q(1-q) \sum_j p_j^4 + 6 \sum_{i<j} q^2 (1-q)^2 p_i^2 p_j^2 \leq \frac{1}{q(1-q)} [q^2 (1-q)^2 \sum_j p_j^4 + 6 \sum_{i<j} q^2 (1-q)^2 p_i^2 p_j^2]$$

$$\leq \frac{3}{q(1-q)} (Var[Y])^2.$$

The desired result now follows from Lemma 4.2, (using the fact that $2^{4/3} \cdot 3 < 10$). □

Remark: For $q \leq 1/2$ the estimate in the lemma is tight, up to a constant factor. Indeed, for $p = (1, 0, 0, \ldots, 0)$ the probability that X is at least q is precisely the probability that $X_1 = 1$, which is q. For $q = 1/k$ with k being an integer there is a simpler argument showing that in this case the probability that X is at least its expectation is at least q (which is precisely tight). The idea is to choose the random vector $(X_1, X_2, \ldots, X_n)$ by first choosing, for each $1 \leq j \leq n$, a random uniform number n_j in $\{1, 2, , \ldots, k\}$ with all choices being independent, and then by selecting a uniform random $Z \in \{1, 2, \ldots, k\}$, defining X_j to be 1 iff $n_j = Z$. Since the sum $\sum_{Z \in [k]} (\sum_{j: n_j = Z} p_j) = 1$, it follows that for each choice of the values n_j, there is at least one Z so that $\sum_{j: n_j = Z} p_j \geq 1/k$, and therefore the probability that the obtained random sum is at least $q = 1/k$ is at least $1/k$, as claimed. Note that for some values of q the probability that X is at least q is strictly smaller than q. Indeed, for example, if $q = 0.501$ and the vector p is $(0.5, 0.5, 0, 0, \ldots, 0)$, then the probability that X is at least q is the probability that $X_1 = X_2 = 1$, which is q^2, that is, roughly $q/2$.

Corollary 4.4. *Let v be a random vector of length n with $\{0,1\}$ entries obtained by selecting each entry, randomly and independently, to be 1 with probability q, and 0 with probability $1-q$. Let $\mathbf{P}$ be any fixed set of distribution vectors of length n. Then, the probability that the inner product of the vector v with each of the vectors $p \in \mathbf{P}$ is at least q, is at least $(\frac{q(1-q)}{10})^{|\mathbf{P}|}$.*

Proof. By the the FKG Inequality (c.f., e.g., [3], Chapter 6.), the probability that the inner product of v with each of the vectors $p \in \mathbf{P}$ is at least q is greater or equal to the product of these probabilities. But according to Lemma 4.3, for every $p \in \mathbf{P}$, the probability that the inner product of v with p is at least q is greater than $\frac{q(1-q)}{10}$. The desired result follows. □

We are now ready to state the proof of Theorem 4.1:

Proof. (sketch) Take a δ-net $\mathcal{N}$ of distributions of length n with respect to the ℓ_1-norm. Apply the last lemma to every set $\mathbf{P}$ of 2^b of them to conclude, using the union bound, that almost surely for every such set there is a pure strategy of the MAX player that ensures her value at least q with respect to each of these mixed strategies. Here we use the fact that $|\mathcal{N}| \leq [(\frac{10}{\delta})]^n$, and hence there are at most $[(\frac{10}{\delta})]^{n2^b}$ ways to choose a set of 2^b members of $\mathcal{N}$. The desired result follows. The missing details are deferred to the full version. □

Somewhat surprisingly, even in the Weak Model, although there are examples in which the MIN player cannot decrease the value by much using at most $\log\log m - O(1)$ bits of information, if he is allowed to use $\log\log m + O(1)$ bits, he can always decrease the value to 0. This is described in the next (simple) result, which, together with the previous theorem, exhibits an unexpected sharp phase transition at $b = \log\log m$.

Theorem 4.5. *Let M be an m by n matrix with $\{0,1\}$ entries, and let q be the value of the game defined by M. If $q^{2^b} < 1/m$, then $v^{(B)}(M,b) = 0$. Therefore, for every fixed q satisfying $0 < q < 1$, $b = \log\log m + O_q(1)$ suffice to ensure value 0 for the MIN player, even in the Weak Model.*

The proof of Theorem 4.5 follows from the proof of Corollary **??**.

5 Optimal strategy computation

We begin with a simple example of a {0,1} matrix M with a minimax strategy $\mathbf{p}_0^*$ that satisfies (i) $v_{\mathbf{p}_0^*}(M,0) = 1/2 \pm o(1)$; (ii) $v_{\mathbf{p}_0^*}(M,1) = 0$; and (iii) $v_{\mathbf{p}_1^*}(M,1) = \Omega(1)$. This shows that playing the minimax strategy may be a naive behavior for $b > 0$, and hence motivates the computation of better strategies. The matrix M showing the above is of dimension 9×14 and it is depicted in Table **??**. The minimax strategy in M is playing the first two rows with probability $1/2$ each, yielding a value of $1/2$. One can easily verify that this is a unique optimal strategy for $b = 0$. For $b = 1$, however, the value of this strategy is clearly 0. Yet, by playing the uniform distribution on the bottom 7 rows, MAX ensures an expected payoff of $\Omega(1)$, even for $b = 1$, as detailed in the proof of Theorem ??.

We first consider the computational complexity of finding the optimal strategy for the MAX player in the Strong Model. The following theorem shows that computing the optimal strategy against b bits is poly-time for any fixed b.

Theorem 5.1. *Given an m by n $\{0,1\}$-matrix and a fixed b, computing the optimal strategy against b bits ($\mathbf{p}_b^*$) in the Strong Model is poly-time.*

Proof. For every $w \in \{0,1\}^b$, let A_w denote the event that MAX plays $i \in [m]$ s.t. $f(i) = w$. The LP that computes the optimal strategy against b bits is given by:

$$\begin{aligned} &\text{maximize} \quad z \\ &\text{subject to} \quad \forall f : [m] \to \{0,1\}^b, \forall g : \{0,1\}^b \to [n], \sum_{w \in \{0,1\}^b} Pr_p(A_w) \cdot E[\text{MAX's payoff} | A_w \wedge g(w)] \geq z \\ &\qquad \sum_{i \in [m]} p_i = 1 \\ &\qquad \forall i \in [m], p_i \geq 0 \end{aligned} \tag{2}$$

Substituting $E[\text{MAX's payoff} | A_w \wedge g(w)] = \frac{1}{Pr_p(A_w)} \sum_{i:f(i)=w} p_i M_{i,g(w)}$ yields the following equivalent LP:

$$\begin{aligned} &\text{maximize} \quad z \\ &\text{subject to} \quad \forall f : [m] \to \{0,1\}^b, \forall g : \{0,1\}^b \to [n], \sum_{w \in \{0,1\}^b} \sum_{i:f(i)=w} p_i M_{i,g(w)} \geq z \\ &\qquad \sum_{i \in [m]} p_i = 1 \\ &\qquad \forall i \in [m], p_i \geq 0 \end{aligned} \tag{3}$$

For any fixed b, the above LP has a polynomial number of variables $(m+1)$ and an exponential number of constraints (since there are 2^{mb} possible f functions). However, a closer analysis shows that the above LP is poly-time. Every g function marks 2^b columns with vectors $w_1, \ldots, w_{2^b} \in \{0,1\}^b$. If we can find for every given g function the f function that minimizes MAX's payoff, we can use that f function in the constraint instead of going over all possible f functions. But this is simple: for every function g, we define a function $f_g : [m] \to [n]$ to be the function that assigns each row $i \in [m]$ to the column marked by w_j that minimizes $M_{i,g(w_j)}$ among all marked columns. (Note that the design of f_g does not depend on p). Based on the above observation, the following LP is equivalent to LP 3:

$$\begin{aligned} &\text{maximize} \quad z \\ &\text{subject to} \quad \forall g : \{0,1\}^b \to [n], \sum_{w \in \{0,1\}^b} \sum_{i:f_g(i)=w} p_i M_{i,g(w)} \geq z \\ &\qquad \sum_{i \in [m]} p_i = 1 \\ &\qquad \forall i \in [m], p_i \geq 0 \end{aligned} \tag{4}$$

Since there is a polynomial number of g functions for any fixed b, the above LP computes the optimal strategy against b bits in polynomial time. □

We next show that for general b, computing the optimal strategy against b bits is NP-hard. Moreover, we show that it is NP-hard to approximate the optimal value by any factor.

Theorem 5.2. *Given an m by n $\{0,1\}$-matrix, It is NP-hard to approximate $v_{\mathbf{p}_b^*}(M,b)$ by any factor in the Strong Model.*

Proof. In order to prove the inapproximability of the optimal value, we show that it is NP-hard to distinguish between the case in which the value is positive and the case in which the value is zero. That is, we show that the following problem is NP-hard:

Instance: A game $M \in \{0,1\}^{m\times n}$, and an integer b.
Question: Is $v_{\mathbf{p}_b^*}(M,b) > 0$?

In order to show NP-hardness, we present a polynomial-time reduction from the NP-hard Set Cover (SC) problem to our problem.

An instance of the SC problem is composed of a finite set of elements $U = \{1,\ldots,m\}$, a collection C of element subsets of U, $C = \{C_1,\ldots,C_r\}$ s.t. for every $1 \le i \le r$, $C_i \subseteq U$, and an integer k. The question is whether there is a subset $C' \subseteq C$, $|C'| \le k$, such that every element in U belongs to at least one member of C'.

We next give the details of the reduction from SC to our problem. Given an instance of SC $\langle U, C, k\rangle$, we construct the following instance of our problem. We define a matrix $M \in \{0,1\}^{m\times n}$, with $m = |U|$ and $n = r$, s.t. $M_{i,j} = 0 \Leftrightarrow i \in C_j$, and set $b = \log k$. We next show that there is a set cover of size smaller or equal to k if and only if $v_{\mathbf{p}_b^*}(M,b) = 0$.

Sufficiency: Suppose the size of the set cover is greater than k. Then, we show that taking the uniform distribution over the whole action set (i.e. setting $p_i = \frac{1}{m}\ \forall i \in [m]$) yields $v_{\mathbf{p}_b^*}(M,b) > 0$.

Consider Inequality 4 and let p be the uniform distribution as described above. For every choice of g, the left-hand side of the inequality is composed of a finite set of summands. In order to show that the obtained payoff is greater than zero, it is sufficient to show that at least one summand is greater than zero. Indeed, since the set cover is greater than $k = 2^b$, for every set of 2^b columns, S, there exists $i \in [m]$, call it i', s.t. $M_{i',j} = 1$ for every $j \in S$, and also $p(i') > 0$ (since p has a full support). Therefore, $p(i')M_{i',g(f_g(i'))} > 0$ for every g. Consequently, $v_{\mathbf{p}_b^*}(M,b) > 0$.

Necessity: Suppose there exists a set cover of size at most $k = 2^b$. Then, there is a set of columns S, $|S| \le 2^b$, s.t. for every $i \in [m]$ there exists $j \in |S|$ for which $M_{i,j} = 0$. Let g be a function that maps every $w \in \{0,1\}^b$ to a different column in S (arbitrarily). By the choice of S, it must hold that for every $i \in [m]$, $M_{i,g(f_g(i))} = 0$. Therefore, for every distribution $\mathbf{p}$, every summand in Inequality 4 equals zero. Consequently, $v_{\mathbf{p}_b^*}(M,b) = 0$

□

In contrast to the last theorem, computing the optimal strategy $\mathbf{p}_b^*$ after observing the function f (as in the Weak Model) is poly-time, using the following algorithm. Given the function $f : [m] \to \{0,1\}^b$, for every $w \in \{0,1\}^b$, let $S^w = \{i : f(i) = w\}$, and let M^w denote the sub-matrix $M^w \in \mathbb{R}^{|S^w|\times n}$, induced by S^w. For every M^w, compute the minimax strategy of M^w, denoted p^w, through the corresponding LP, and let v^w denote the value of the game M^w. MAX will play p^w for the w that yields the highest value of v^w.

Finally, we consider the computational complexity of finding the optimal f function for the MIN player in the Weak Model. For a general b, the exact same reduction from Set Cover, presented in the proof of Theorem 5.2, shows that computing the optimal f function is NP-hard and that it is NP-hard to find an f function that approximates the optimal value (for the MIN player) within any factor.

For a fixed b, however, the MIN player can compute the optimal f function as follows. Consider all $\binom{n}{2^b}$ possible g functions, and for every g function (defined by a set of 2^b columns), let $f_g : [m] \to [n]$ be the best (from the perspective of MIN) f function given g; i.e., f_g maps every row to the column out of the 2^b columns with minimal payoff. For every g, MAX will play p^w (i.e., minimax strategy, computed by a linear program) for the $w \in \{0,1\}^b$ that maximizes his payoff among all the sub-games M^w (as specified above). Knowing that, MIN will choose the g function that minimizes this maximal value, and its corresponding f_g function.

References

[1] N. Alon, O. Goldreich, J. Håstad and R. Peralta, *Simple constructions of almost k-wise independent random variables*, Proc. 31^{st} IEEE FOCS, St. Louis, Missouri, IEEE (1990), pp. 544-553.

[2] N. Alon, G. Gutin and M. Krivelevich, Algorithms with large domination ratio, J. Algorithms 50 (2004), 118-131.

[3] N. Alon and J. H. Spencer, *The Probabilistic Method, Third Edition*, Wiley, 2008, xv+352 pp.

[4] D. S. Johnson, Approximation algorithms for combinatorial problems, J. Comput. System Sci. 9 (1974), 256-278.

[5] L. Lovász, On the ratio of optimal and fractional covers, Discrete Mathematics 13 (1975), 383-390.

[6] J. Naor and M. Naor, *Small-bias probability spaces: efficient constructions and applications*, Proc. 22^{nd} annual ACM STOC, ACM Press (1990), pp. 213-223.

[7] A. Weil, Sur les courbes algebriques et les varietes qui s'en deduisent, Actualites Sci. et Ind. No. 1041 (1948), Herman, Paris.

A Explicit Constructions

In this appendix we describe several explicit constructions of n by n $\{0,1\}$-matrices representing games with value q, such that if the MIN player has b bits and b is smaller than $\log\log n + O(1)$, then the MAX player can guarantee a payoff of at least roughly q^{2^b}. This shows (by explicit examples) that the statements of Theorem 3.3 and Corollary 3.4 are essentially tight.

Example 1:

Let p be a prime power and let r be a positive integer. Fix $n = p^r - 1$. Let $M = (M_{u,v})$ be the following n by n binary matrix whose rows and columns are indexed by the set N of all nonzero vectors of length r over $GF(p)$. For each such u, v, $M_{u,v} = 1$ if and only if the two vectors u and v are orthogonal over $GF(p)$ (namely, their inner product over $GF(p)$ is zero). Note that M is a symmetric matrix, and every row and every column of it contains exactly $p^{r-1} - 1$ 1-entries. Indeed, this is the number of non-zero solutions of a single linear equation in r variables over $GF(p)$. It is easy to check that the minimax strategy of the game determined by M is the uniform distribution over N, yielding a value of $q = \frac{p^{r-1}-1}{p^r-1}$. Note that for large $n = p^r - 1$ this is very close to $1/p$.

We claim that for every set $J \subseteq N$ of at most $\log_p n$ columns, there are at least $p^{r-|J|} - 1$ rows u so that $M_{u,v} = 1$ for every $v \in J$. Note that if $p^{r-|J|}$ is large, then this number is very close to $q^{|J|}n$, implying that by playing the uniform distribution on the rows of M, MAX can ensure a value close to $q^{|J|}$. Note also that if $b \leq \log r - O(1) = \log\log n - O_p(1)$, then 2^b is much smaller than r, and hence $p^{r-|J|}$ is large provided $|J| \leq 2^b$. Fix a subset $J \subseteq N$ of cardinality at most $\log_p n$. By definition, row u satisfies $M_{u,v} = 1$ for every $v \in J$ if and only if the inner product of u and v over $GF(p)$ is zero for every $v \in J$. This is a homogeneous system of $|J|$ linear equations in the r variables representing the coordinates of u. This system clearly admits at least $p^{r-|J|} - 1$ non-trivial solutions; each such non-trivial solution corresponds to a row with the desired properties, proving the claim.

This completes the description of the first set of examples. Note that it works for every value q which is about $1/p$, where p is a prime power.

Example 2 (sketch):

Let p be a prime and let M be a $p \times p$ binary matrix, where $M_{i,j} = 1$ if and only if $i - j$ is a quadratic residue modulo p (where here zero is considered a quadratic residue). The value of the game represented by M is $(p+1)/(2p)$ which, for large p, is roughly $1/2$. Using Weil's Theorem, it is not difficult to show that for every subset S of Z_p of size at most $(0.5 - \delta)\log p$, the number of rows i such that $M_{i,j} = 1$ for all $j \in S$, is $(1 + o(1))\frac{p}{2^{|S|}}$. A similar example holds for characters of other orders instead of the quadratic character, providing examples with values close to $1/d$ for any desired positive integer $d > 1$ (where here we have to choose a prime p so that d divides $p - 1$- by Dirichlet's Theorem on primes in arithmetic progressions it is known that there are infinitely many such primes for any such d).

More generally, one can use any construction of small sample spaces supporting nearly 2^b-wise independent binary random variables to supply additional examples. We omit the details, which will appear in the full version of the paper.