# An Information Complexity Approach to Extended Formulations 

Mark Braverman＊Ankur Moitra ${ }^{\dagger}$

February 13， 2014


#### Abstract

We prove an unconditional lower bound that any extended formulation that achieves an $O\left(n^{1-\epsilon}\right)$ approximation for clique has size $2^{\Omega\left(n^{\epsilon}\right)}$ ．There has been considerable recent interest in proving lower bounds for extended formulations．Fiorini et al［14］proved that there is no polynomial sized extended formulation for traveling salesman．Braun et al［7］proved that there is no polynomial sized $O\left(n^{1 / 2-\epsilon}\right)$－approximate extended formulation for clique．Here we prove an optimal and unconditional lower bound against extended formulations for clique that matches Håstad＇s［16］celebrated hardness result．Interestingly，the techniques used to prove such lower bounds have closely followed the progression of techniques used in communication complexity． Here we develop an information theoretic framework to approach these questions，and we use it to prove our main result．

Also we resolve a related question：How many bits of communication are needed to get $\epsilon$－ advantage over random guessing for disjointness？Kalyanasundaram and Schnitger［19］proved that a protocol that gets constant advantage requires $\Omega(n)$ bits of communication．This result in conjunction with amplification implies that any protocol that gets $\epsilon$－advantage requires $\Omega\left(\epsilon^{2} n\right)$ bits of communication．Here we improve this bound to $\Omega(\epsilon n)$ ，which is optimal for any $\epsilon>0$ ．


[^0]
## 1 Introduction

### 1.1 Background

A typical goal in combinatorial optimization is to minimize a linear objective function over a discrete set of solutions. A powerful change in perspective is obtained by representing these discrete solutions as vectors and taking their convex hull to get instead a continuous optimization problem. Thus to many discrete sets of combinatorial interest, we associate a polytope and through this paradigm algorithms for linear programming and geometric insights about polytopes and facets can be brought to bear to design new algorithms and approximation algorithms. Unfortunately the natural encoding of a discrete optimization problem as a continuous one often has an exponential number of facets. Nevertheless often it is possible to express a given polytope $P$ as the linear projection of a higher dimensional polytope $Q$ and in so doing to drastically reduce the number of facets needed. In this case, we call $Q$ an extension of the polytope $P$ (sometimes we will refer to $Q$ as an extended formulation instead). In general $P$ and $Q$ are related by

$$
P=\{x \mid(x, y) \in Q\}
$$

and we can draw an analogy with the more familiar case of Boolean formulas: every Boolean formula can be expressed without using quantifiers, but often a Boolean formula can be expressed much more succinctly using quantifiers. Indeed, in many well-known examples in combinatorial optimization, particular polytopes admit a much more compact description as the linear projection of a higher dimensional polytope. We will refer to the minimum number of facets of the higher dimensional polytope as the extension complexity of $P$.

For example, the permutahedron is the convex hull of permutations of the vector $\{1,2,3, \ldots, n\}$. Rado gave an exact characterization of this polytope using an exponential number of inequalities but omitting even a single one results in a polytope that strictly contains the permutahedron. However, if we represent each permutation instead as a permutation matrix then the permutahedron can be recovered from a linear projection and the Birkhoff-von Neumann theorem gives a set of $2 n^{2}$ linear constraints (that the matrix be doubly-stochastic) that exactly characterize the convex hull of permutation matrices. Another well-known example is the spanning tree polytope: Edmonds gave an exact characterization of this polytope using an exponential number of facets, and Martin [23] gave a compact extended formulation using $O\left(n^{3}\right)$ constraints. In fact, there are even general principles by which one can obtain a compact extended formulation: If the separation oracle for a linear programming problem itself relies on linear programming, then one can often give a polynomial-sized extended formulation. Also, dynamic programming algorithms typically lead to an extended formulation for the associated polytope whose size is roughly the space used by the algorithm. Lastly, Balas [4] proved that polytopes that admit a compact extended formulation are closed under disjunction.

There has been remarkable success in characterizing the facets of certain polytopes that arise in combinatorial optimization. In some cases, this can lead to faster algorithms (that avoid using the Ellipsoid algorithm). And yet another motivation is that having a convenient representation of a given polytope can make it easier to prove a statement about purely discrete objects. (For example, the $T$-join polytope and the spanning tree polytope play an important role in recent approximation algorithms for graphic TSP [25] and asymmetric TSP [3] respectively).

Yet there are many polytopes that we do not expect to have a simple description. For example, if we could find a polynomial-sized extended formulation expressing the convex hull of all traveling salesman tours, we could then use any one of a number of efficient algorithms known for linear programming to solve the traveling salesman problem in polynomial time. So certainly, believers in
$N P \subsetneq P \backslash$ poly $^{1}$ do not believe such an extended formulation exists (and conversely, if you believe that $P=N P$, it is quite reasonable to believe that an algorithm for 3 -SAT would involve linear programming). Remarkably, there are tools for understanding the extension complexity of a given polytope (i.e. for proving that there is no compact extended formulation). To describe these tools, we must introduce some notation:

In a ground breaking work, Yannakakis defined the notion of a slack matrix. Given a polytope $P$ with $v$ vertices and $f$ facets, the slack matrix is a $v \times f$ entry-wise nonnegative matrix whose $(i, j)$ th entry is $b_{j}-\left\langle a_{j}, v_{i}\right\rangle$ where $v_{i}$ is the $i$ th vertex of $P$ and $\left\langle a_{j}, x\right\rangle \leq b_{j}$ is the $j$ th constraint. This matrix provides a tight connection between two combinatorial concepts - the first is the extension complexity as defined earlier, and the second is the nonnegative rank: The nonnegative rank of a matrix $M$ is the smallest $r$ so that $M$ can be written as $M=\sum_{i=1}^{r} A_{i}$ where each $A_{i}$ is a nonnegative, rank one matrix. We will use $\operatorname{rank}^{+}(M)$ to denote the nonnegative rank of $M$. This concept plays an important role in various methods in machine learning (see [2] and [24]). And Yannakakis proved that the nonnegative rank of the slack matrix is precisely the extension complexity (of a given polytope $P$ ). This result is often referred to as Yannakakis's factorization theorem and exactly characterizes one important combinatorial parameter - the extension complexity - in terms of another - the nonnegative rank.

### 1.2 Extended Formulations and Communication Complexity

So if there are polytopes that we do not expect to have a simple description, how might we go about proving that the nonnegative rank of the slack matrix is indeed superpolynomial? The increasingly sophisticated techniques used to prove strong lower bounds on the nonnegative rank have followed in parallel with the techniques used to prove lower bounds on communication complexity! In order to compare the various works that prove lower bounds on nonnegative rank (and the technique that each one uses) it is helpful to consider each work as providing an answer to the question:

Question 1.1. If a matrix $M$ has small nonnegative rank, what does that imply?
Then each work uses a particular consequence of small nonnegative rank to derive a contradiction (with some communication complexity result).

- [14]: If $M$ has small nonnegative rank, then a nonnegative matrix factorization $M=\sum_{i=1}^{r} A_{i}$ gives a covering of the non-zero entries of $M$ by combinatorial rectangles (the support of each $A_{i}$ ). This in turn contradicts the lower bounds for the nondeterministic communication complexity of unique disjointness (when $M$ is chosen to be the correlation polytope).

Fiorini et al [14] used this connection to prove that there is no polynomial-sized extended formulation for the traveling salesman polytope. An important point is that this result (and others in this line) would follow immediately from $P \neq N P$, but in fact this result holds without any assumption!

- [7]: If $M$ has small nonnegative rank, then there is a large rectangle in $M$ with discrepancy bounded away from zero and this contradicts Razborov's rectangle corruption lemma.

Braun et al [7] use this to prove that any extended formulation for clique that achieves an $O\left(n^{1 / 2-\epsilon}\right)$ approximation has size at least $2^{n^{\epsilon}}$. What Braun et al [7] need is a technique to prove

[^1]lower bounds on the nonnegative rank of matrices that have no zero entries, and accomplish this through a generalization of Razborov's rectangle corruption lemma (for disjointness).

Interestingly, the result of [7] falls short of the known $N P$-hardness results for clique. If $P \neq$ $N P$, then the celebrated result of Håstad [16] and later Khot [21] and Zuckerman [33] directly imply that any extended formulation for clique that achieves an $O\left(n^{1-\epsilon}\right)$ approximation has superpolynomial size. Can we prove an analogue of this result without the assumption that $P \neq N P$ ?

In this paper, we give a new method to lower bound nonnegative rank.

- [this work]: If $M$ has small nonnegative rank, then we can use the factorization $M=\sum_{i} A_{i}$ to generate a uniform sample from $S$ - the set of pairs of disjoint strings - using too few bits of entropy.

This new method allows us to draw on some of the most powerful tools in communication complexity - information complexity. Information complexity has been recently shown [20] to contain as a sub-case many other lower bound techniques for two-player communication complexity. This is the first time that tools from information complexity have been used in the context of nonnegative rank, and we believe that this new method for proving lower bounds will be useful in other applications.

Our main theorem is (see Section 3 for background information on the polytope associated with clique):

Theorem 1.2. Any extended formulation for clique that achieves an $s+1$-approximation has size at least $2^{\Omega\left(\frac{n}{s+1}\right)}$.

Thus any extended formulation that achieves an $O\left(n^{1-\epsilon}\right)$-approximation must have size at least $2^{\Omega\left(n^{\epsilon}\right)}$ and our lower bound for extended formulations matches the celebrated hardness results of Håstad [16]. In fact, the best known approximation algorithm for clique is due to Feige [13] and achieves an approximation ratio of $O\left(n(\log \log n)^{2} /(\log n)^{3}\right)$ and our results even imply that any extended formulation that achieves this bound has size at least $\Omega\left(n^{\tilde{\Omega}\left((\log n)^{2}\right)}\right)$. This is the first example that we know of in which there is a polynomial time approximation algorithm and yet provably there is no polynomial-sized extended formulation that achieves the same approximation ratio! Although we remark that it is widely believed that there is no polynomial-sized extended formulation ${ }^{2}$ for matching even though there is an efficient algorithm ${ }^{3}$.

### 1.3 Disjointness

In order to explain our main technical contribution (even at a heuristic level) it helps to consider the communication complexity of disjointness. The techniques of Braun et al [7] break down at $n^{1 / 2}$ precisely because there is a gap in our understanding of the communication complexity of disjointness!

In this problem, Alice and Bob are given $a, b \in\{0,1\}^{n}$ and their goal is to determine whether or not these two strings are disjoint, that is, whether there is an index $i \in[n]$ such that $a_{i}=$

[^2]$b_{i}=1$. Kalyanasundaram and Schnitger [19] were the first to prove an $\Omega(n)$ lower bound on the communication complexity of any protocol that gets a constant advantage. In fact any linear lower bound in conjunction with amplification implies that any protocol that gets $\epsilon$-advantage requires $\Omega\left(\epsilon^{2} n\right)$ bits of communication ${ }^{4}$. Since then a number of simpler proofs have been given. Razborov [27] proved an elegant rectangle corruption lemma which yields a linear lower bound. Perhaps the most intuitive lower bound was given in the ground breaking work of Bar-Yossef et al [6] which used information theoretic arguments to prove that an protocol that gets advantage $\epsilon$ over random guessing must communicate at least $\Omega\left(\epsilon^{2} n\right)$ bits. However, this does not quite settle the asymptotic complexity of disjointness.

Question 1.3. Is there a protocol that gets advantage $1 / \sqrt{n}$ for disjointness and uses a subpolynomial amount of communication?

There is a simple to state technical issue that arises in these proofs, and makes $\Omega\left(\epsilon^{2} n\right)$ a barrier for these approaches: For example, the work of Bar-Yossef et al [6] works by reducing an $n$-bit communication problem - disjointness - to a one-bit problem - an AND. Indeed, any lower bound (say, $\gamma$ ) on the information revealed to the participants (about the inputs to Alice and Bob) based on the transcript implies a $\gamma n$ lower bound on the communication complexity of the $n$-bit problem. Yet, there is a protocol for AND that gets advantage $\epsilon$ and reveals only $O\left(\epsilon^{2}\right)$ bits of information (see Lemma 2.1)!

So in a sense, the $n$-bit problem cannot be reduced to $n$ instances of a one-bit problem! Yet here we are able to circumvent this issue. Instead of considering the advantage of a protocol, we consider a more nuanced quantity - a matrix that describes the probability of outputting a one for each pair of input bits for Alice and Bob. Suppose this matrix is:

$$
N=\left[\begin{array}{l}
N_{00}, N_{01} \\
N_{10}, N_{11}
\end{array}\right]
$$

Then the advantage of the protocol is at least $\gamma$ if and only if $N_{00}, N_{10}, N_{01} \geq 1 / 2+\gamma$ and $N_{11} \leq 1 / 2-\gamma$. Indeed, there is a protocol for AND that gets $N_{00}=1 / 2+5 \epsilon+\Theta\left(\epsilon^{2}\right), N_{10}=$ $N_{01}=1 / 2+\epsilon+\Theta\left(\epsilon^{2}\right), N_{11}=1 / 2-3 \epsilon+\Theta\left(\epsilon^{2}\right)$, and reveals $O\left(\epsilon^{2}\right)$ bits of information. Another way to think about this is that Hellinger distance (see [6] for the details) implies that the information revealed is at least $\left(N_{00}-N_{11}\right)^{2}$. Our main insight is:

- The information revealed is at least $\Omega\left(N_{10}+N_{01}-N_{00}-N_{11}\right)$.
- This does not directly imply a better than $\Omega\left(\epsilon^{2}\right)$ lower bound for one-bit AND (afterall, how could it since there is such a protocol?) but nevertheless any protocol for disjointness can be "smoothed" so that the communication complexity does not increase and yet, in expectation $N_{00}=N_{10}=N_{01}$.

To put it another way, after this "smoothing" operation changing a pair of bits $\left(a_{j}, b_{j}\right)$ from (say) $(0,0)$ to $(1,0)$ will not change the probability that the protocol outputs one. Hence the onebit AND problems that we get by considering just a single pair of bits $\left(a_{j}, b_{j}\right)$ is not an arbitrary problem where the goal is to get advantage $\epsilon$ in any way possible but rather to do so in such a way that the probability of outputting a one is the same, independent of whether the input is $(0,0)$, $(0,1)$ or $(1,0)$. Yet for protocols that meet this extra restriction, getting an advantage of $\epsilon$ implies

[^3]that at least $\Omega(\epsilon)$ bits are revealed. Hence using the direct sum results for information complexity, we obtain an $\Omega(\epsilon n)$ lower bound for disjointness, thus resolving the asymptotic complexity of disjointness for any $\epsilon>0$ (not just $\epsilon=\Omega(1)$ ).

How does this relate to nonnegative rank? As we mentioned, Braun et al [7] gave an interesting generalization of Yannakakis's factorization theorem to the case of approximate extended formulations. And it turns out that if you want to prove a lower bound on the extension complexity of any polytope that approximates $P$ within a multiplicative factor of $C$, then this quantity is equal to the nonnegative rank of $M+(C-1) J$, where $M$ is the slack matrix of $P$ and $J$ is the all ones matrix. (And indeed nonnegative rank lower bounds are proven using communication complexity lower bounds for disjointness). Then we can think of the factors $A_{i}\left(\right.$ in $\left.M+(C-1) J=\sum_{i} A_{i}\right)$ intuitively as the "output" matrix $N$ in the preceding discussion, and hence proving a nonnegative rank lower bound for $C=n^{1-\epsilon}$ is akin to proving a lower bound for disjointness when the target advantage $\epsilon \approx 1 / C$. This is precisely why the proof of [7] breaks down for $C=\sqrt{n}$ - because it was not known whether or not getting advantage $1 / \sqrt{n}$ for disjointness requires a polynomial amount of communication!

However, lower bounds for disjointness do not immediately yield lower bounds for nonnegative rank. Intuitively, many techniques in communication complexity aim to show that a player must reveal too much information about his input in order to solve the communication problem. Yet in the setting of nonnegative rank ${ }^{5}$, a player can reveal a nonnegative real value (which can contain a great deal of information about his input).

## 2 A New Lower Bound for Disjointness

Here we prove that any protocol for disjointness that gets advantage $\epsilon$ over random guessing must communicate at least $\Omega(\epsilon n)$ bits between the two players. This improves the $\Omega\left(\epsilon^{2} n\right)$ lower bound due to Razborov [27] and Bar-Yossef et al [6]. The known lower bounds for disjointness were optimal only for $\epsilon=\Omega(1)$ - it was a priori possible that a protocol that gets advantage $1 / \sqrt{n}$ requires as much as $\Omega(\sqrt{n})$ communication or as little as $O(1)$ communication. Our lower bound completely resolves the asymptotic complexity of disjointness for any $\epsilon>0$.

The standard (information complexity) approach for proving lower bounds is to reduce an $n$-bit problem to a one-bit problem (which in the case of disjoints is a one-bit AND). The difficulty is that there is indeed a protocol that gets advantage $\epsilon$ for one-bit AND, but reveals $O\left(\epsilon^{2}\right)$ bits of information to an observer:

Lemma 2.1. There is a communication protocol for one-bit AND that gets advantage $\epsilon$ over random guessing but reveals $O\left(\epsilon^{2}\right)$ bits of information to an observer.

Proof. We first describe a simpler protocol that almost works: Alice and Bob behave identically, and on input one Alice sends a one with probability $1 / 2+4 \epsilon$ and otherwise sends a zero. On input zero Alice sends a zero with probability $1 / 2+4 \epsilon$ and otherwise sends a one. (Bob behaves identically). If both players send a one, the protocol outputs one. If both players send a zero, the protocol outputs zero. Otherwise the protocol flips a coin and outputs either one or zero with equal probability. (This is not quite the protocol in the lemma, but a simple modification will complete the proof). The probability that the protocol outputs a zero, for each pair of inputs, is given by the matrix:

$$
N=\left[\begin{array}{l}
1 / 2+4 \epsilon, 1 / 2 \\
1 / 2,1 / 2-4 \epsilon
\end{array}\right]
$$

[^4]This protocol $\Gamma$ does not have an advantage over random guessing for all inputs, but we can obtain one that does with a simple modification. The final protocol $\Pi$ will run $\Gamma$ with probability $1-2 \epsilon$, and with probability $2 \epsilon$ will output zero. Then $\Pi$ has advantage $\epsilon$ and a simple calculation shows that it reveals $O\left(\epsilon^{2}\right)$ bits of information to an observer.

Notice that the probability $\Pi$ outputs zero is different when the input to Alice and Bob are 0 and 0 or instead 1 and 0 respectively (even though the answer in both of these cases is the same). The standard approach to proving lower bounds (though information complexity) is to use a $n$-bit protocol to get a one-bit protocol. In fact, we can "symmetrize" any $n$-bit protocol for disjointness in such a way that the one-bit protocol we obtain has the same probability of outputting zero whether the input to Alice and Bob is 0 and 0,1 and 0 or 0 and 1 respectively. Our main insight is that this one bit communication problem (with an additional constraint that the top-left, top-right and bottom-left values of $N$ be the same) must reveal $\Omega(\epsilon)$ bits of information.

Theorem 2.2. Any protocol for disjointness that gets advantage $\epsilon$ over random guessing must reveal at least $\Omega(\epsilon n)$ bits of information to the participants.

Consider a protocol $\Gamma$. We will consider the information complexity of $\Gamma$ measured with respect to a particular distribution on inputs: We will group the $n$ bits into blocks of size exactly three, and for each pair of three bits we will generate $a_{j}, b_{j} \in\{0,1\}^{3}$ uniformly at random from the pairs of length three strings where $a_{j}$ and $b_{j}$ have exactly one 1 and two 0 's, and $a_{j}$ and $b_{j}$ are disjoint. Thus the location of the 1 must be different in $a$ and $b$, and therefore there are six such pairs.

Given the protocol $\Gamma$ we will construct a new protocol $\Pi$ which takes each block of size three and applies a random element $\pi$ of $S_{3}$ to both $a_{j}$ and $b_{j}$. This does not change whether or not $a_{j}$ or $b_{j}$ are disjoint, and furthermore has the effect of "smoothing" the protocol $\Gamma$ : For any pair $a_{j}, b_{j}$ (that are disjoint), the probability that $\Pi$ outputs one is exactly the same. We will abuse notation and let $B_{j}=i$ be the event that the only one in $b_{j}$ is the $i^{t h}$ bit in the block.

We will be interested in $I(A ; \Pi \mid B)+I(B ; \Pi \mid A)$. The following is well-known:
Fact 2.3. $\frac{1}{2}\left[\sum_{j} I\left(A_{j} ; \Pi \mid A_{1 \ldots j-1}, B_{j \ldots n}\right)+I\left(B_{j} ; \Pi \mid A_{1 \ldots j}, B_{j+1 \ldots n}\right)\right] \leq H(\Pi)$
Proof. Using the chain rule for mutual information (see e.g. [8]) we have

$$
\frac{1}{2}\left[\sum_{j} I\left(A_{j} ; \Pi \mid A_{1 \ldots j-1}, B_{1 \ldots n}\right)+I\left(B_{j} ; \Pi \mid A_{1 \ldots n}, B_{j+1 \ldots n}\right)\right] \leq \frac{1}{2}[I(A ; \Pi \mid B)+I(B ; \Pi \mid A)] \leq H(\Pi)
$$

Moreover we will use the following inequality for mutual information (again, see [8]): $I(W ; X \mid Y) \leq$ $I(W ; X \mid Y Z)$ if $I(X ; Z \mid Y)=0$. We can apply the inequality above for example by setting $W=\Pi$, $X=A_{j}, Y=A_{1 \ldots j-1} B_{j \ldots n}$ and $Z=B_{1 \ldots j-1}$. Then it is easy to see that $I(X ; Z \mid Y)=0$ in this case, and we conclude:

$$
I\left(A_{j} ; \Pi \mid A_{1 \ldots j-1}, B_{j \ldots n}\right) \leq I\left(A_{j} ; \Pi \mid A_{1 \ldots j-1}, B_{1 \ldots n}\right)
$$

Similarly we can set $W=\Pi, X=B_{j}, Y=A_{1 \ldots j} B_{j+1 \ldots n}$ and $Z=A_{j+1 \ldots n}$ and again we have $I(X ; Z \mid Y)=0$ and so

$$
I\left(B_{j} ; \Pi \mid A_{1 \ldots j}, B_{j+1 \ldots n}\right) \leq I\left(B_{j} ; \Pi \mid A_{1 \ldots n}, B_{j+1 \ldots n}\right)
$$

and this completes the proof.

Let $C_{j}=A_{1 \ldots j-1}, B_{j+1 \ldots n}$. Then we can write $I\left(A_{j} ; \Pi \mid A_{1 \ldots j-1} B_{j . . n}\right)$ as:

$$
\sum_{c, i} \sum_{t} \operatorname{Pr}\left[\Pi=t, C_{j}=c, B_{j}=i\right] D\left(A_{j}\left|C_{j}=c, B_{j}=i, \Pi=t \| A_{j}\right| C_{j}=c, B_{j}=i\right),
$$

where $D(\bullet \| \bullet)$ is the KL-divergence of the two distributions - see [11] for more background on information theory. Note that $A_{j}$ conditioned on $B_{j}=i$ is uniform on the set $\{1,2,3\}-\{i\}$.

Let us choose $C_{j}=A_{1 \ldots j-1}, B_{j+1 \ldots n}$ according to the distribution on inputs. Then, $I\left(A_{j} ; \Pi \mid C_{j}, B_{j}\right)+$ $I\left(B_{j} ; \Pi \mid C_{j}, A_{j}\right)=\sum_{t} \mathbf{E}\left[\operatorname{adv}\left(t, C_{j}\right)\right]$ where the expectation is over $C_{j}$ and $\operatorname{adv}\left(t, C_{j}\right)$ defined as:

$$
\begin{aligned}
\operatorname{adv}\left(t, C_{j}\right)=\quad \sum_{i=1,2,3} & \operatorname{Pr}\left[\Pi=t, B_{j}=i \mid C_{j}\right] D\left(A_{j}\left|B_{j}=i, \Pi=t, C_{j} \| A_{j}\right| B_{j}=i, C_{j}\right) \\
& +\operatorname{Pr}\left[\Pi=t, A_{j}=i \mid C_{j}\right] D\left(B_{j}\left|A_{j}=i, \Pi=t, C_{j} \| B_{j}\right| A_{j}=i, C_{j}\right) .
\end{aligned}
$$

Definition 2.4. Let $I C(a, b, c)=(a b+a c) D\left(\mathcal{B}_{a b /(a b+a c)} \| \mathcal{B}_{1 / 2}\right)$.
Lemma 2.5. $I C\left(\alpha_{0}, \beta_{0}, \beta_{1}\right)+I C\left(\beta_{0}, \alpha_{0}, \alpha_{1}\right) \geq \Omega\left(\alpha_{0} \beta_{1}+\alpha_{1} \beta_{0}-\alpha_{0} \beta_{0}-\alpha_{1} \beta_{1}\right)$
Proof. We have that $D\left(\mathcal{B}_{\alpha_{0} \beta_{0} /\left(\alpha_{0} \beta_{0}+\alpha_{0} \beta_{1}\right)} \| \mathcal{B}_{1 / 2}\right)=\Omega\left(\frac{\left(\beta_{0}-\beta_{1}\right)^{2}}{\left(\beta_{0}+\beta_{1}\right)^{2}}\right)$ and hence it suffices to show that:

$$
\frac{2 \alpha_{0}\left(\beta_{0}-\beta_{1}\right)^{2}}{\beta_{0}+\beta_{1}}+\frac{2 \beta_{0}\left(\alpha_{0}-\alpha_{1}\right)^{2}}{\alpha_{0}+\alpha_{1}} \geq \alpha_{0} \beta_{1}+\alpha_{1} \beta_{0}-\alpha_{0} \beta_{0}-\alpha_{1} \beta_{1} .
$$

We multiply both sides by $\left(\alpha_{0}+\alpha_{1}\right)\left(\beta_{0}+\beta_{1}\right)$ and the left hand side is:

$$
4 \alpha_{0}^{2} \beta_{0}^{2}+2 \alpha_{0}^{2} \beta_{1}^{2}+2 \alpha_{1}^{2} \beta_{0}^{2}+2 \alpha_{0} \alpha_{1} \beta_{1}^{2}+2 \alpha_{1}^{2} \beta_{0} \beta_{1}-2 \alpha_{0}^{2} \beta_{0} \beta_{1}-2 \alpha_{0} \alpha_{1} \beta_{0}^{2}-8 \alpha_{0} \alpha_{1} \beta_{0} \beta_{1}
$$

The right hand side is:

$$
\left(\alpha_{0}+\alpha_{1}\right)\left(\alpha_{0}-\alpha_{1}\right)\left(\beta_{1}-\beta_{0}\right)\left(\beta_{0}+\beta_{1}\right)=\alpha_{0}^{2} \beta_{1}^{2}+\alpha_{1}^{2} \beta_{0}^{2}-\alpha_{0}^{2} \beta_{0}^{2}-\alpha_{1}^{2} \beta_{1}^{2}
$$

and subtracting the right hand side from the left hand side we get:
$\left(\alpha_{0}^{2} \beta_{0}^{2}+\alpha_{0}^{2} \beta_{1}^{2}-2 \alpha_{0}^{2} \beta_{0} \beta_{1}\right)+\left(\alpha_{0}^{2} \beta_{0}^{2}+\alpha_{1}^{2} \beta_{0}^{2}-2 \alpha_{0} \alpha_{1} \beta_{0}^{2}\right)+\left(3 \alpha_{0}^{2} \beta_{0}^{2}+\alpha_{1}^{2} \beta_{1}^{2}+2 \alpha_{0} \alpha_{1} \beta_{1}^{2}+2 \alpha_{1}^{2} \beta_{0} \beta_{1}-8 \alpha_{0} \alpha_{1} \beta_{0} \beta_{1}\right)$
which is at least zero, using the weighted AMGM inequality.
We will consider a fixed block $j$, and the matrix $N^{t}\left(C_{j}\right)$ that gives the probability of $\Pi=t$ (where the output is one) for each pair of inputs for Alice and Bob conditioned on the parts of their input $C_{j}$ that we have already fixed. (Note that the expectation here is taken over the randomness of the protocol and the remaining bits in the input to Alice and Bob). To simplify notation we will abbreviate $N^{t}\left(C_{j}\right)$ as $N^{t}$. Let us write:

$$
N^{t}=\left[\begin{array}{l}
N_{11}^{t}, N_{12}^{t}, N_{13}^{t} \\
N_{21}^{t}, N_{22}^{t}, N_{23}^{t} \\
N_{31}^{t}, N_{32}^{t}, N_{33}^{t}
\end{array}\right]
$$

Note that $N^{t}$ is a nonnegative. Since $\Pi$ is a protocol and Alice and Bob can privately sample their remaining bits conditioned on $C_{j}, A_{j}$ and $B_{j}$ we conclude that $N^{t}$ is a rank one matrix ${ }^{6}$. So we can write $N^{t}=\left[a_{1}, a_{2}, a_{3}\right]\left[b_{1}, b_{2}, b_{3}\right]^{T}$. In particular $b_{i}$ is the probability over $B_{1 \ldots j-1}$ that the string $B=B_{1 \ldots j-1}, B_{j}=i, B_{j+1 \ldots n}$ is in the rectangle for $\Pi=t$.

[^5]Lemma 2.6. $a d v\left(t, C_{j}\right)=\Omega\left(\sum_{i \neq i^{\prime}} N_{i i^{\prime}}^{t}\left(C_{j}\right)-2 \sum_{i=1,2,3} N_{i i}^{t}\left(C_{j}\right)\right)$
Proof. We will prove this lemma by repeatedly applying Lemma 2.5 . For example, set $\alpha_{0}=a_{3}$, $\alpha_{1}=a_{1}$ and $\beta_{0}=b_{2}, \beta_{1}=b_{1}$. We can symbolically think of this as "covering" a set of entries in the matrix $N^{t}$ :

$$
\left[\begin{array}{ccc}
\alpha_{1} \beta_{1}, \alpha_{1} \beta_{0}, & - \\
-, & -, & - \\
\alpha_{0} \beta_{1}, & \alpha_{0} \beta_{0}, & -
\end{array}\right]
$$

where we have chosen $\alpha_{1}$ and $\beta_{1}$ so that the corresponding entry in $N^{t}$ is on the diagonal. The term $I C\left(\alpha_{0}, \beta_{0}, \beta_{1}\right)$ equals $\operatorname{Pr}\left[\Pi=t \mid A_{j}=3, C_{j}\right] D\left(B_{j}\left|A_{j}=3, \Pi=t, C_{j} \| B_{j}\right| A_{j}=3, C_{j}\right)$ and we can think of this as the bottom row in this covering. Similarly, the term $\operatorname{IC}\left(\beta_{0}, \alpha_{0}, \alpha_{1}\right)$ equals $\operatorname{Pr}\left[\Pi=t \mid B_{j}=2, C_{j}\right] D\left(A_{j}\left|B_{j}=2, \Pi=t, C_{j} \| A_{j}\right| B_{j}=2, C_{j}\right)$ and we can think of this as the middle column in the covering. And applying Lemma 2.5 has the effect of adding $N_{1,2}^{t}$ and $N_{3,1}^{t}$ and subtracting $N_{1,1}^{t}$ and $N_{3,2}^{t}$. We can apply Lemma 2.5 to the following covering scheme:

$$
\left[\begin{array}{c}
*, *,- \\
-,-,- \\
*, *,-
\end{array}\right],\left[\begin{array}{c}
-, *, * \\
-, *, * \\
-,-,--
\end{array}\right],\left[\begin{array}{c}
-, *, * \\
-,-,- \\
-, *, *
\end{array}\right],\left[\begin{array}{c}
*,-, * \\
*,-, * \\
-,-,-
\end{array}\right],\left[\begin{array}{c}
-,-,- \\
*, *,- \\
*, *,-
\end{array}\right],\left[\begin{array}{c}
-,-,- \\
*,-, * \\
*,-, *
\end{array}\right]
$$

always choosing $\alpha_{1}$ and $\beta_{1}$ so that the corresponding entry in $N^{t}$ is on the diagonal. The left hand side double-counts each term in $a d v\left(t, C_{j}\right)$ and the right hand side is $\sum_{i \neq i^{\prime}} N_{i i^{\prime}}^{t}-2 \sum_{i=1,2,3} N_{i i}^{t}$ and this completes the proof.

We can now complete the proof of the main theorem in this section.
Proof. If we sum $\sum_{i \neq i^{\prime}} N_{i i^{\prime}}^{t}\left(C_{j}\right)-2 \sum_{i=1,2,3} N_{i i}^{t}\left(C_{j}\right)$ over all transcripts $t$ for which the output of $\Pi$ is one and take the expectation over $C_{j}$, the value is $\Omega(\epsilon)$ by the assumption that (1) the advantage of the protocol is $\epsilon$ and (2) for any pair of disjoint $a_{j}, b_{j}$ where $a_{j}$ and $b_{j}$ each have exactly one one, the probability that the protocol outputs one is the same. Hence we conclude that $I\left(A_{j} ; \Pi \mid B_{j}, C_{j}\right)+I\left(B_{j} ; \Pi \mid A_{j}, C_{j}\right)=\Omega(\epsilon)$ for each block $j$, and combining this with Fact 2.3 we conclude that $H(\Pi)=\Omega(\epsilon n)$ since the input restricted to each block is mutually independent of the other blocks. And so the communication complexity of any protocol for disjointness that gets advantage $\epsilon$ is $\Omega(\epsilon n)$.

## 3 Extended Formulations for Clique

Recall that the nonnegative rank is defined as follows:
Definition 3.1. The nonnegative rank $\operatorname{rank}^{+}(M)$ of a nonnegative matrix $M$ is the smallest $r$ so that $M$ can be written as $M=\sum_{i=1}^{r} M_{i}$ where each $M_{i}$ is a rank-one nonnegative matrix.

The nonnegative rank plays a central role in lower bounds for extended formulations, and here we will explain this connection in more detail. The central polytope in the recent breakthrough lower bounds of Fiorini et al [14] and of Braun et al [7] is the correlation polytope.
Definition 3.2. $P=\operatorname{COR}(n)=\operatorname{conv}\left\{b b^{T} \mid b \in\{0,1\}^{n}\right\}$
Recall that the smallest extended formulation (for a given polytope) is exactly the nonnegative rank of the slack matrix. However one can prove lower bounds on extended formulations (for a given polytope) by considering only a subset of the constraints. This corresponds to proving a nonnegative rank lower bound on a submatrix of the full slack matrix. Fiorini et al [14] identified a set of constraints on the correlation polytope that are enough to prove strong lower bounds:

Lemma 3.3. [14] For any $a \in\{0,1\}^{n}$, the inequality $\left\langle 2 \operatorname{diag}(a)-a a^{T}, x\right\rangle \leq 1$ is valid for $C O R(n)$ and the slack of a vertex $x=b b^{T}$ is $\left(1-a^{T} b\right)^{2}$.

Hence these constraints define a polytope $Q$ for which $P \subset Q$. In fact, Fiorini et al [14] prove that the extension complexity of $\operatorname{COR}(n)$ is $2^{\Omega(n)}$ by considering only pairs of vertices $b b^{T}$ and constraints (for $a \in\{0,1\}^{n}$ ) such that $b^{T} a \in\{0,1\}$. Whether or not a vertex is on the facet is exactly the unique disjointness problem. An important definition from [7] (and a generalization from the one in [31]) is:
Definition 3.4. [7] Let $P \subset Q$, then the slack matrix $S M(P, Q)$ is a nonnegative matrix where each row corresponds to a vertex $v_{i}$ in $P$ and a column corresponds to a constraint $\left\langle a_{j}, x\right\rangle \leq b_{j}$ in $Q$ and the corresponding entry in $S M(P, Q)$ is $b_{j}-\left\langle a_{j}, v_{i}\right\rangle$.

Braun et al [7] gave a generalization of Yannakakis's factorization theorem [31]. This is the connection between nonnegative rank and extension complexity (that is relevant when our goal is to show lower bounds for approximate extended formulations) that we will use here:

Theorem 3.5. [7] The minimum extension complexity of any polytope $P \subset K \subset Q$ is exactly the nonnegative rank of $S M(P, Q)$.

In fact, the correlation polytope naturally defines a linear encoding for clique: Given a graph $G$ on $n$ nodes, one can choose an objective function $w(G)$ (which is an $n \times n$ matrix that is the direction we are trying to maximize over $\operatorname{COR}(n))$ where $w(G)$ is one on each diagonal entry and is zero on $i, j$ and $j, i$ if $(i, j)$ is an edge, and otherwise is -1 . The maximum value of $w(G)$ over $\operatorname{COR}(n)$ is exactly the maximum clique value of $G$. (And in fact, these directions are admissible in the sense of $[7]$ in that any lower bounds for the nonnegative $\operatorname{rank}$ of $S M(P, Q)$ imply that any extended formulation for this linear encoding is large).

What about approximate extended formulations for clique? Clearly, we would like to prove lower bounds on the nonnegative rank of $S M(P,(s+1) Q)$ in order to show that there is no small extended formulation (for clique) that has an approximation factor of $s$. In this case, the slack matrix $S M(P,(s+1) Q)$ has a particularly simple structure: Restricted to the pairs of strings that have at most one intersection, it is $S M(P, Q)+s J$ where $J$ is the all ones matrix! Hence, what we need to prove lower bounds on approximate extended formulations of clique is just to prove lower bounds for the nonnegative rank of the slack matrix (generated by the correlation polytope and the constraints defined above) when a large value (say $n^{1-\epsilon}$ ) is added to each entry that corresponds to a pair of strings with at most one intersection.

## 4 A New Method to Lower Bound Nonnegative Rank

Let $M$ be a $2^{n} \times 2^{n}$ nonnegative matrix. We will associate rows and columns of $M$ with strings $a, b \in\{0,1\}^{n}$ and furthermore if $a^{T} b \in\{0,1\}$ then $M_{a, b}=s+1-a^{T} b$. If $s=0$, Fiorini et al [14] proved that $\operatorname{rank}^{+}(M) \geq 2^{\Omega(n)}$. If $s=O\left(n^{\beta}\right)$ with $\beta<1 / 2$, Braun et al [7] proved that $\operatorname{rank}^{+}(M) \geq 2^{\Omega\left(n^{1-2 \beta}\right)}$. Here we prove that even if $s=n^{1-\epsilon}$, then $\operatorname{rank}^{+}(M) \geq 2^{\Omega\left(n^{\epsilon}\right)}$.

In fact, we will prove such a lower bound by constructing a sampling procedure (from a nonnegative matrix factorization of $M$ ) and deriving an entropy-theoretic contradiction. Suppose that $f_{1}, f_{2}, \ldots f_{r}$ and $g_{1}, g_{2}, \ldots g_{r}$ are nonnegative functions defined on $\{0,1\}^{n}$ and that furthermore for each $a, b \in\{0,1\}^{n}$ with $a^{T} b \in\{0,1\}$ we have

$$
\sum_{i=1}^{r} f_{i}(a) g_{i}(b)=s+1-a^{T} b
$$

## Algorithm 1. Generate, Output: $(a, b)$

1. Choose an index $i \in[r]$ from $\gamma$
2. Choose $(a, b) \in S$ from $\mu_{i}$
3. Output $(a, b)$

Definition 4.1. Let $S=\left\{a, b \mid a^{T} b=0\right\} \subset\{0,1\}^{n} \times\{0,1\}^{n}$.
We will next define a sampling procedure (based on $f_{i}$ and $g_{i}$ ) that generates a uniformly random pair of strings $(a, b) \in S$ :

Definition 4.2. Let $W_{i}=\sum_{(a, b) \in S} f_{i}(a) g_{i}(b)$. Let $\gamma(i)=\frac{W_{i}}{\sum_{i=1}^{r} W_{i}}$ and let $\mu_{i}(a, b)=\frac{f_{i}(a) g_{i}(b)}{W_{i}}$ for all $(a, b) \in S$ and $\mu_{i}=0$ elsewhere.

Claim 4.3. The procedure GENERATE samples a $(a, b) \in S$ uniformly at random.

Proof. The probability of Generate outputting $(a, b) \in S$ is

$$
\sum_{i=1}^{r} \gamma(i) \mu_{i}(a, b)=\sum_{i=1}^{r} \frac{W_{i}}{\sum_{i^{\prime}=1}^{r} W_{i^{\prime}}} \frac{f_{i}(a) g_{i}(b)}{W_{i}}=\frac{s+1}{\sum_{i^{\prime}=1}^{r} W_{i^{\prime}}}
$$

Throughout this section, we will let $A_{1}, A_{2}, \ldots A_{n}$ and $B_{1}, B_{2}, \ldots B_{n}$ denote the random variables associated with the output $(a, b)$ and additionally we will let $I$ denote the random variable associated with the intermediate random variable $i$ in Generate.

Corollary 4.4. $H\left(A_{1 \ldots n}, B_{1 \ldots n}\right)=\log |S|=(\log 3) n$ and $n H\left(A_{j}\right)=H\left(A_{1 \ldots n}\right)=H\left(B_{1 \ldots n}\right)=$ $n H\left(B_{j}\right)=n H(1 / 3)$

We proceed to invoke the chain rule (and we will analyze each term using properties of $f_{i}$ and $g_{i}$ ):

$$
\begin{aligned}
H\left(A_{1 \ldots n}, B_{1 \ldots n}\right) \leq & H\left(A_{1 \ldots n}, B_{1 \ldots n}, I\right)=H(I)+H\left(A_{1 \ldots n}, B_{1 \ldots n} \mid I\right) \\
= & H(I)+\frac{1}{2}\left[H\left(B_{1 \ldots n} \mid I\right)+\sum_{j=1}^{n} H\left(A_{j} \mid A_{1 \ldots j-1}, B_{1 \ldots n}, I\right)\right] \\
& \quad+\frac{1}{2}\left[H\left(A_{1 \ldots n} \mid I\right)+\sum_{j=1}^{n} H\left(B_{j} \mid A_{1 \ldots n}, B_{j+1 \ldots n}, I\right)\right] \\
& \quad \log r+n H(1 / 3)+\frac{1}{2} \sum_{j=1}^{n} H\left(A_{j} \mid A_{1 \ldots j-1}, B_{j \ldots n}, I\right)+\frac{1}{2} \sum_{j=1}^{n} H\left(B_{j} \mid A_{1 \ldots j}, B_{j+1 \ldots n}, I\right)
\end{aligned}
$$

where in the last inequality we have used the fact that $H(X \mid Y) \leq H(X)$.
Definition 4.5. Let $a d v_{j}^{A}(i)=1-H\left(A_{j} \mid A_{1 \ldots j-1}, B_{j}=0, B_{j+1 \ldots n}, I=i\right)$ and similarly let $a d v_{j}^{B}(i)=$ $1-H\left(B_{j} \mid A_{1 \ldots j-1}, A_{j}=0, B_{j+1 \ldots n}, I=i\right)$.

Note that if $B_{j}=1$, then $A_{j}=0$ and the entropy of $A_{j}$ (conditioned on the given events) is zero. We will use this in conjunction with the following fact about conditional entropy:

Fact 4.6. $H(W \mid X=x, Y, Z)=\sum_{z} \operatorname{Pr}[Z=z \mid X=x] H(W \mid X=x, Y, Z=z)$
Then:

$$
\begin{aligned}
H\left(A_{1 \ldots n}, B_{1 \ldots n}\right) \leq & \log r
\end{aligned}+n H(1 / 3)+\frac{1}{2} \sum_{j=1}^{n} \underset{I}{\mathbf{E}}\left[H\left(A_{j} \mid A_{1 \ldots j-1}, B_{j \ldots n}, I=i\right)\right] \quad \begin{aligned}
&+\frac{1}{2} \sum_{j=1}^{n} \underset{I}{\mathbf{E}}\left[H\left(B_{j} \mid A_{1 \ldots j}, B_{j+1 \ldots n}, I=i\right)\right] \\
&= \log r+n H(1 / 3)+\frac{1}{2} \sum_{j=1}^{n} \underset{I}{\mathbf{E}}\left[\operatorname{Pr}\left[B_{j}=0 \mid I=i\right] H\left(A_{j} \mid A_{1 \ldots j-1}, B_{j}=0, B_{j+1 \ldots n}, I=i\right)\right] \\
&+\frac{1}{2} \sum_{j=1}^{n} \underset{I}{\mathbf{E}}\left[\operatorname{Pr}\left[A_{j}=0 \mid I=i\right] H\left(B_{j} \mid A_{1 \ldots j-1}, A_{j}=0, B_{j+1 \ldots n}, I=i\right)\right] \\
&=\log r+n H(1 / 3)+\frac{1}{2} \sum_{j=1}^{n} \underset{I}{\mathbf{E}}\left[\operatorname{Pr}\left[B_{j}=0 \mid I=i\right]\left(1-a d v_{j}^{A}(i)\right)\right] \\
&+\frac{1}{2} \sum_{j=1}^{n} \underset{I}{\mathbf{E}}\left[\operatorname{Pr}\left[A_{j}=0 \mid I=i\right]\left(1-a d v_{j}^{B}(i)\right)\right] \\
&=\log r+n H(1 / 3)+2 n / 3-\frac{1}{2} \sum_{j=1}^{n} \underset{I}{\mathbf{E}}\left[\operatorname{Pr}\left[B_{j}=0 \mid I=i\right] a d v_{j}^{A}(i)\right] \\
& \quad-\frac{1}{2} \sum_{j=1}^{n} \underset{I}{\mathbf{E}}\left[\operatorname{Pr}\left[A_{j}=0 \mid I=i\right] a d v_{j}^{B}(i)\right]
\end{aligned}
$$

Note that $H(1 / 3)+2 / 3=\log 3$.
Definition 4.7. Let $L_{i, j}=a d v_{j}^{A}(i) \operatorname{Pr}\left[B_{j}=0 \mid I=i\right]$.
It will be easier to analyze $a d v_{j}^{A}(i)$ directly rather than the quantity $L_{i, j}$ and the following lemma will provide the means for translating lower bounds on $a d v_{j}^{A}(i)$ into lower bounds on $L_{i, j}$ :

Lemma 4.8. Suppose for all $j \in[n], \mathbf{E}_{I}\left[a d v_{j}^{A}(i)\right] \geq \frac{1}{s}$, then $\mathbf{E}_{I}\left[\sum_{j=1}^{n} L_{i, j}\right]=\Omega\left(\frac{n}{s}\right)-O(\log r)$.
Proof. Since $a d v_{j}^{A}(i) \leq 1$, we get:

$$
\begin{aligned}
\underset{I}{\mathbf{E}}\left[\sum_{j=1}^{n} L_{i, j}\right] & \geq \frac{1}{4} \underset{I}{\mathbf{E}}\left[\sum_{j=1}^{n} a d v_{j}^{A}(i)\right]-\frac{1}{4} \underset{I}{\mathbf{E}}\left[\left|\left\{j \left\lvert\, \operatorname{Pr}\left[B_{j}=0 \mid I=i\right] \leq \frac{1}{4}\right.\right\}\right|\right] \\
& \geq \frac{n}{4 s}-\frac{1}{4} \mathbf{E}\left[\left|\left\{j \left\lvert\, \operatorname{Pr}\left[B_{j}=0 \mid I=i\right] \leq \frac{1}{4}\right.\right\}\right|\right]
\end{aligned}
$$

Our immediate goal is to bound $\mathbf{E}_{I}\left[\left|\left\{j \left\lvert\, \operatorname{Pr}\left[B_{j}=0 \mid I=i\right] \leq \frac{1}{4}\right.\right\}\right|\right]$. To accomplish this, we will use the following claim:
Claim 4.9. $H\left(B_{j} \mid I\right) \leq H(1 / 3)-\Omega\left(\operatorname{Pr}_{I}\left[\operatorname{Pr}\left[B_{j}=0 \mid I=i\right]<\frac{1}{4}\right]\right)$
Proof.

Fact 4.10. $H(c+x) \leq H(c)+H^{\prime}(c) x-\Omega\left(x^{2}\right)$
Hence:

$$
\begin{aligned}
H\left(B_{j} \mid I\right) & =\underset{I}{\mathbf{E}}\left[H\left(B_{j} \mid I=i\right)\right] \\
& \leq H(2 / 3)+H^{\prime}(2 / 3) \underset{I}{\mathbf{E}}\left[\operatorname{Pr}\left[B_{j}=0 \mid I=i\right]-2 / 3\right]-\Omega\left(\underset{I}{\mathbf{E}}\left[\left(\operatorname{Pr}\left[B_{j}=0 \mid I=i\right]-2 / 3\right)^{2}\right]\right) \\
& =H(1 / 3)-\Omega\left(\underset{I}{\mathbf{E}}\left[\left(\operatorname{Pr}\left[B_{j}=0 \mid I=i\right]-2 / 3\right)^{2}\right]\right) \\
& \leq H(1 / 3)-\Omega\left(\operatorname{Pr}_{I}\left[\operatorname{Pr}\left[B_{j}=0 \mid I=i\right]<\frac{1}{4}\right)\right.
\end{aligned}
$$

Next, we can use this claim to prove Lemma 4.8:

$$
\begin{aligned}
H\left(B_{1 \ldots n}\right)-H(I) & =H\left(B_{1 \ldots n} \mid I\right) \leq \sum_{j=1}^{n} H\left(B_{j} \mid I\right) \\
& \leq \sum_{j=1}^{n}\left(H\left(B_{j}\right)-\Omega\left(\operatorname{Pr} r_{I}\left[\operatorname{Pr}\left[B_{j}=0 \mid I=i\right]<\frac{1}{4}\right]\right)\right) \\
& =n H(1 / 3)-\Omega\left(\underset{I}{\mathbf{E}}\left[\left|\left\{j \left\lvert\, \operatorname{Pr}\left[B_{j}=0 \mid I=i\right] \leq \frac{1}{4}\right.\right\}\right|\right]\right)
\end{aligned}
$$

However, $H\left(B_{1 \ldots n}\right)=n H(1 / 3)$ and so $\mathbf{E}_{I}\left[\left|\left\{j \left\lvert\, \operatorname{Pr}\left[B_{j}=0 \mid I=i\right] \leq \frac{1}{4}\right.\right\}\right|\right]=O(H(I))=O(\log r)$.

An analogous lemma holds with $A$ and $B$ exchanged. We will prove the following lemma in the next section:

Lemma 4.11 (Main). For all $j \in[n], \mathbf{E}_{I}\left[a d v_{j}^{A}(i)+a d v_{j}^{B}(i)\right] \geq \Omega\left(\frac{1}{s+1}\right)$
In fact in the exact case $(s=0)$, it is much easier to prove an $\Omega(1)$ lower bound (see the discussion after Lemma 5.10). We can now put these pieces together:

$$
(\log 3) n=H\left(A_{1 \ldots n}, B_{1 \ldots n}\right) \leq O(\log r)+(\log 3) n-\Omega\left(\frac{n}{s+1}\right)
$$

and hence $r \geq 2^{\Omega\left(\frac{n}{s+1}\right)}$, and this proves our main theorem:
Theorem 4.12. $\operatorname{rank}^{+}(M) \geq 2^{\Omega\left(\frac{n}{s+1}\right)}$

## 5 The Typical Advantage

Here we prove the main lemma, and this will complete our main theorem. Throughout this section, we will fix $j \in[n]$.

Definition 5.1. Let $S_{j}=\left\{a, b \mid a_{-j}^{T} b_{-j}=0\right\} \subset\{0,1\}^{n} \times\{0,1\}^{n}$
In particular, $S_{j}$ is the set of pairs of strings $a, b$ so that $a$ and $b$ are disjoint except possibly on the $j^{\text {th }}$ coordinate. Note that $\left|S_{j}\right|=4 \times 3^{n-1}$. Our proof will be based on a single data structure that will allow us to sample from either $A_{j}$ conditioned on $A_{1 \ldots j-1}, B_{j}=0, B_{j+1 \ldots n}, I=i$ or from $B_{j}$ conditioned on $A_{1 \ldots j-1}, A_{j}=0, B_{j+1 \ldots n}, I=i$. We now proceed to define this structure.

Definition 5.2. Let $E \subset S_{j}$ and let $K \subset[r]$, then we will denote $N_{E, K}^{j}$ as a $2 \times 2$ matrix, which for each $a_{j}^{\prime}, b_{j}^{\prime} \in\{0,1\}$ the $\left(a_{j}^{\prime}, b_{j}^{\prime}\right)$ entry of $N_{E, K}^{j}$ is equal to

$$
\sum_{(a, b) \in E, a_{j}=a_{j}^{\prime}, b_{j}=b_{j}^{\prime}} \sum_{i \in K} f_{i}(a) g_{i}(b)
$$

Claim 5.3. $N_{S_{j},[r]}^{j}=\frac{\left|S_{j}\right|}{4}\left[\begin{array}{l}s+1, s+1 \\ s+1, \quad s\end{array}\right]$
Proof. Consider $(a, b) \in S_{j}$. Suppose, for example, that $a_{j}=0, b_{j}=0$ then this pair contributes $s+1$ to the $(0,0)$ entry in $N_{S_{j},[r]}^{j}$.

Next we define the data structure that will be crucial in our proof. The nodes will correspond to $2 \times 2$ matrices given by $N_{E, H}$ :

Definition 5.4. The sampling tree $T$ is a depth four tree in which:

- The root node corresponds to $E=S_{j}$ and $K=[r]$
- The first layer corresponds to nodes $E=S_{j}$, and $K=\{i\}$
- The second layer corresponds to $E=\left\{(a, b) \in S_{j} \mid a_{1}=a_{1}^{\prime}, \ldots a_{j-1}=a_{j-1}^{\prime}, b_{j+1}=b_{j+1}^{\prime}, \ldots b_{n}=\right.$ $\left.b_{n}^{\prime}\right\}$ and $K=\{i\}$
- The third layer corresponds to $E=\left\{(a, b) \in S_{j} \mid a_{1}=a_{1}^{\prime}, \ldots a_{j-1}=a_{j-1}^{\prime}, a_{j}=a_{j}^{\prime}, b_{j}=\right.$ $\left.b_{j}^{\prime}, b_{j+1}=b_{j+1}^{\prime}, \ldots b_{n}=b_{n}^{\prime}\right\}$ and $K=\{i\}$

Furthermore a node defined by $E, K$ is connected to a node in a lower layer defined by $E^{\prime}, K^{\prime}$ if and only if $E^{\prime} \subset E$ and $K^{\prime} \subset K$.

Claim 5.5. For any node $u \in T$, we have that $N_{u}=\sum_{v \in \operatorname{child}(u)} N_{v}$.
Proof. In fact, by construction if a node $u$ corresponds to $E, K$ then the children $v$ either have $E^{\prime}=E$ or $K^{\prime}=K$ and in the first case, the children $v$ define a partition of $K$ (and vice-versa in the other case).

An important point is that in any leaf node, the corresponding matrix $N$ has only one non-zero entry. Next, we define three different methods for sampling the pair $\left(a_{j}, b_{j}\right)$ using this tree:

Definition 5.6. Given a $2 \times 2$ matrix $N$, we will define $N(F)$ to be the sum of the first column of $N, N(G)$ to be the sum of the first row of $N$ and $N(T)$ to be the total sum of the entries in $N$.

Using this definition, we can define $F$-sampling, $G$-sampling and $T$-sampling:
Definition 5.7. $F$-sampling generates a $\left(a_{j}, b_{j}\right)$ pair as follows:

- Start at the root node.
- While the current node $u$ is not a leaf node, choose a child $v$ of $u$ with probability $\frac{N_{v}(F)}{N_{u}(F)}$.
- Output the $\left(a_{j}, b_{j}\right)$ pair corresponding to the non-zero entry in the leaf node at termination.
$G$-sampling and $T$-sampling are defined analogously with $F$ replaced by $G$ or $T$ respectively. Also if instead we start the $F$ sampling procedure at a first layer node with $K=\{i\}$ we will call it $F, i$-sampling and similarly for $G$ and $T$.

The crucial observations are:
Observation 1. The F-sampling and G-sampling procedures represent a faithful method to sample from the distribution $A_{j}$ conditioned on $A_{1 \ldots j-1}, B_{j}=0, B_{j+1 \ldots n}, I$ and $B_{j}$ conditioned on $A_{1 \ldots j-1}, A_{j}=0, B_{j+1 \ldots n}, I$ respectively.
Observation 2. The $F, i$-sampling and $G, i$-sampling procedures represent a faithful method to sample from the distribution $A_{j}$ conditioned on $A_{1 \ldots j-1}, B_{j}=0, B_{j+1 \ldots n}, I=i$ and $B_{j}$ conditioned on $A_{1 \ldots j-1}, A_{j}=0, B_{j+1 \ldots n}, I=i$ respectively.

Note that the last choice in an $F$-sampling procedure represents choosing the random variable $A_{j}$ conditioned on $A_{1 \ldots j-1}, B_{j}=0, B_{j+1 \ldots n}, I$. And similarly the last choice in a $G$-sampling procedure represents choosing the random variable $B_{j}$ conditioned on $A_{1 \ldots j-1}, A_{j}=0, B_{j+1 \ldots n}, I$. Hence we can interpret the advantage through these processes:
Corollary 5.8. $\quad$ adv $v_{j}^{A}(i)$ is equal to $1-H$, where $H$ is the expected entropy of the last random choice when performing $F$, $i$-sampling. Also adv $v_{j}^{B}(i)$ is equal to $1-H$, where $H$ is the expected entropy of the last random choice when performing $G, i$-sampling.

Finally, we can use the $T$-sampling procedure to simulate either $F$-sampling or $G$-sampling:
Claim 5.9. The following procedure is equivalent to $F$-sampling: Use $T$-sampling to obtain $\left(a_{j}, b_{j}\right)$ and if $b_{j}=0$, output $\left(a_{j}, b_{j}\right)$ and otherwise restart the process. Also, the following procedure is equivalent to $G$-sampling: Use $T$-sampling to obtain $\left(a_{j}, b_{j}\right)$ and if $a_{j}=0$, output ( $a_{j}, b_{j}$ ) and otherwise restart the process.

We are now ready to prove the main lemma. We will let $N$ denote the random variable corresponding to the matrix generated by running the $T$-sampling procedure until reaching the parent of a leaf. Similarly, let $N_{00}, N_{10}, N_{10}$ and $N_{11}$ denote the entries of this matrix. Then:

Lemma 5.10. $N$ is always a rank one matrix.
Proof. Since $N$ is the matrix corresponding to the parent of a leaf, we have that $E$ corresponds to a fixed choice $a_{1 \ldots j-1}$ for $A_{1 \ldots j-1}, b_{j+1 \ldots n}$ for $B_{j+1 \ldots n}$ and $I=i$. Hence the ( $a_{j}^{\prime}, b_{j}^{\prime}$ ) entry of $N$ is $F_{i}\left(a_{j}^{\prime}\right) G_{i}\left(b_{j}^{\prime}\right)$, where $F_{i}\left(a_{j}^{\prime}\right)$ is the sum of $f_{i}(a)$ over all $a$ with $a_{j}=a_{j}^{\prime}$ and $a_{j+1 \ldots n}^{T} b_{j+1 \ldots n}=0$ and similarly $G_{i}\left(b_{j}^{\prime}\right)$ is the sum of $g_{i}(b)$ over all $b$ with $b_{j}=b_{j}^{\prime}$ and $b_{1 . . . j-1}^{T} a_{1 \ldots j-1}=0$.

In the exact case ( $s=0$ ), our proof is particularly simple. Since $N$ is a rank one matrix and must have $N_{11}=0$, there must be another zero in either the same row or column. Hence choosing an entry in $N$ proportional to its value has entropy at most one bit, and this is already $\log 3-1$ bits smaller than the uniform distribution on the top-left, bottom-left and top-right entries, and this already implies an exponential lower bound in the exact case. Since our goal here is to give lower bounds even for large values of $s$ that approach $n$, we must be more careful in accounting for the entropy lost when the entry $N_{11}$ is allowed to be non-zero.
Definition 5.11. Let $\operatorname{adv}(N)=\frac{N_{00}+N_{10}}{N_{00}+N_{10}+N_{10}+N_{11}}\left[1-H\left(\frac{N_{00}}{N_{00}+N_{10}}\right)\right]+\frac{N_{00}+N_{01}}{N_{00}+N_{10}+N_{10}+N_{11}}\left[1-H\left(\frac{N_{00}}{N_{00}+N_{01}}\right)\right]$.
Lemma 5.12. $\quad \mathbf{E}_{I}\left[a d v_{j}^{A}(i)+a d v_{j}^{B}(i)\right] \geq \mathbf{E}_{N}[a d v(N)]$
Proof. We can use Corollary 5.8 to get an interpretation of $\mathbf{E}_{I}\left[a d v_{j}^{A}(i)\right]$ as $1-H$, where $H$ is the entropy of the last random choice in $F$-sampling. And using Claim 5.9 we have that $\mathbf{E}_{N}\left[\frac{N_{00}+N_{10}}{N_{00}+N_{10}+N_{10}+N_{11}}\right]$ is the probability that $T$-sampling results in a sample coupled with one from $F$-sampling. Hence, the expected entropy lost in the last step of $F$-sampling is lower bounded by the entropy lost in the last step of $T$-sampling if the sample satisfies $B_{j}=0$. An identical argument for $a d v_{j}^{B}(i)$ implies the lemma.

Furthermore $N$ is nonnegative so we can always write

$$
N=\left(N_{01}+N_{10}+N_{00}+N_{11}\right)\left[\begin{array}{c}
\left(\frac{1}{2}+\beta\right)\left(\frac{1}{2}-\gamma\right),\left(\frac{1}{2}+\beta\right)\left(\frac{1}{2}+\gamma\right) \\
\left(\frac{1}{2}-\beta\right)\left(\frac{1}{2}-\gamma\right),\left(\frac{1}{2}-\beta\right)\left(\frac{1}{2}+\gamma\right)
\end{array}\right]
$$

Note that $\frac{N_{01}+N_{10}-N_{00}-N_{11}}{N_{01}+N_{10}+N_{00}+N_{11}}=4 \beta \gamma$.
Lemma 5.13. $\operatorname{adv}(N)=\Omega(\beta \gamma)$
Proof. Using the bound $1-H\left(\frac{1}{2}+x\right) \geq x^{2}$ :

$$
\begin{aligned}
\operatorname{adv}(N) & =\left(\frac{1}{2}-\gamma\right) H\left(\frac{1}{2}+\beta\right)+\left(\frac{1}{2}+\beta\right) H\left(\frac{1}{2}+\gamma\right) \\
& \geq\left(\frac{1}{2}-\gamma\right) \beta^{2}+\left(\frac{1}{2}+\beta\right) \gamma^{2}
\end{aligned}
$$

Clearly we can assume that $\beta \gamma>0$ otherwise the lemma is trivial. We can prove the lemma through a simple case analysis:

- Case: $\beta>-1 / 3$ and $\gamma<1 / 3$, then $\operatorname{adv}(N)=\Omega\left(\beta^{2}+\gamma^{2}\right)=\Omega(\beta \gamma)$.
- Case: $\beta \leq-1 / 3$, then $\operatorname{adv}(N) \geq\left(\frac{1}{2}-\gamma\right) \beta^{2}$ because the other term is non-negative, and furthermore $\beta^{2}=\Omega(1)$ and $\left(\frac{1}{2}-\gamma\right)=\Omega(1)$ since $\gamma \leq 0$.
- Case: $\gamma \geq 1 / 3$, then $\operatorname{adv}(N) \geq\left(\frac{1}{2}+\beta\right) \gamma^{2}$ and again both terms are $\Omega(1)$ because $\beta \geq 0$.

Lemma 5.14. $\quad \mathbf{E}_{N}\left[\frac{N_{01}+N_{10}-N_{00}-N_{11}}{N_{01}+N_{10}+N_{00}+N_{11}}\right]=\frac{1}{4 s+3}$
Proof. The quantity $\mathbf{E}_{N}\left[\frac{N_{01}+N_{10}-N_{00}-N_{11}}{N_{01}+N_{10}+N_{00}+N_{11}}\right]$ is exactly the probability that under $T$-sampling the output has $\left(a_{j}, b_{j}\right) \in\{(0,1),(1,0)\}$ minus the probability that it has $\left(a_{j}, b_{j}\right) \in\{(0,0),(1,1)\}$. We can analyze this probability directly by considering the matrix corresponding to the root node, which has $(2 s+2) \frac{\left|S_{j}\right|}{4}$ total weight in the $(1,0)$ and $(0,1)$ entries and has $(2 s+1) \frac{\left|S_{j}\right|}{4}$ total weight in the $(0,0)$ and $(1,1)$ entries. Hence the difference of the probabilities of these two events is exactly $\frac{1}{4 s+3}$.

And using Lemma 5.12, Lemma 5.13 and Lemma 5.14 we conclude that $\mathbf{E}_{I}\left[a d v_{j}^{A}(i)+a d v_{j}^{B}(i)\right]=$ $\Omega\left(\frac{1}{s+1}\right)$. This concludes the proof of the main lemma.

## Acknowledgements

We would like to thank Avi Wigderson for many helpful discussions at an early stage of this work.

## References

[1] S. Arora, B. Bollobás and L. Lovász. Proving integrality gaps without knowing the linear program. FOCS, pp. 313-322, 2002.
[2] S. Arora, R. Ge, R. Kannan and A. Moitra. Computing a nonnegative matrix factorization provably. STOC, pp. 145-162, 2012.
[3] A. Asadpour, M. Goemans, A. Madry, S. Oveis Gharan and A. Saberi. An $O(\log n / \log \log n)-$ approximation algorithm for the asymmetric traveling salesman problem. SODA, pp. 379-389, 2010.
[4] E. Balas. Disjunctive programming and a hierarchy of relaxations for discrete optimization problems. SIAM J. Algebraic Discrete Methods, pp. 466-486, 1985.
[5] B. Barak, M. Braverman, X. Chen and A. Rao. How to compress interactive communication. STOC, pp. 67-76, 2010.
[6] Z. Bar-Yossef, T.S. Jayram, R. Kumar and D. Sivakumar. An information statistics approach to data stream and communication complexity. JCSS, pp. 702-732, 2004.
[7] G. Braun, S. Fiorini, S. Pokutta and D. Steurer. Approximation limits of linear programs (beyond hierarchies). FOCS, pp. 480-489, 2012.
[8] M. Braverman. Interactive information complexity. STOC pp. 505-524, 2012.
[9] M. Braverman and A. Rao. Information Equals Amortized Communication. FOCS, pp. 748757, 2011.
[10] A. Chattopadhyay and T. Pitassi. The story of set disjointness. SIGACT News, pp. 59-85, 2010.
[11] T. Cover and J. Thomas. Elements of Information Theory. J. Wiley and Sons, 1991.
[12] J. Edmonds. Maximum matching and a polyhedron and 0, 1-vertices. Journal of Research of the National Bureau of Standards, pp. 125-130, 1965.
[13] U. Feige. Approximating maximum clique by removing subgraphs. SIAM J. Discrete Math, pp. 219-225, 2004.
[14] S. Fiorini, S. Massar, S. Pokutta, H. Tiwary and R. de Wolf. Linear vs semidefinite extended formulations: exponential separations and strong lower bounds. STOC, pp. 95-106, 2012.
[15] S. Fiorini, T. Rothvoß and H. Tiwary. Extended formulations for polygons. Arxiv, 2011.
[16] J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$ Acta Mathematica, pp. 105-142, 1999. Preliminary version in FOCS 1996.
[17] V. Kaibel. Extended formulations in combinatorial optimization. Optima, pp. 2-7, 2011.
[18] V. Kaibel, K. Pashkovich and D. Theis. Symmetry matters for the sizes of extended formulations. $I P C O$, pp. 135-148, 2010.
[19] B. Kalyanasundaram and G. Schnitger. The probabilistic communication complexity of set intersection. SIAM J. Discrete Math, pp. 545-557, 1992.
[20] I. Kerenidis, S. Laplante, V. Lerays, J. Roland and D. Xiao. Lower bounds on information complexity via zero-communication protocols and applications. FOCS 2012, to appear.
[21] S. Khot. Improved inapproximability results for MaxClique, chromatic number and approximate graph coloring. FOCS, pp. 600-609, 2001.
[22] E. Kushilevitz and N. Nisan. Communication Complexity. Cambridge University Press, 1997.
[23] K. Martin. Using separation algorithms to generate mixed integer model reformulations. $O p$ erations Research Letters, pp. 119-128, 1991.
[24] A. Moitra. An almost optimal algorithm for computing nonnegative rank. SODA 2013, to appear.
[25] S. Oveis Gharan, A. Saberi and M. Singh. A randomized rounding approach to the traveling salesman problem. FOCS, pp. 550-559, 2011.
[26] M. Padberg and M. Rao. Odd minimum cut-sets and b-matchings. Mathematics of Operations Research, pp. 67-80, 1982.
[27] A. Razborov. On the distributional complexity of disjointness. Theoretical Computer Science, pp. 385-390, 1992.
[28] T. Rothvoss. Some 0/1 polytopes need exponential size extended formulations. Arxiv, 2011.
[29] S. Vavasis. On the complexity of nonnegative matrix factorization. SIAM Journal on Optimization, pp. 1364-1377, 2009.
[30] L. Wolsey. Using extended formulations in practice. Optima, pp. 7-9, 2011.
[31] M. Yannakakis. Expressing combinatorial optimization problems by linear programs. JCSS, pp. 441-466, 1991. Preliminary version in STOC 1988.
[32] G. Ziegler. Lectures on Polytopes. Springer-Verlag, 1995.
[33] D. Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. STOC, pp. 681-690, 2006.


[^0]:    ${ }^{*}$ Princeton University．Email：mbraverm＠cs．princeton．edu．Research supported in part by an Alfred P．Sloan Fellowship，an NSF CAREER award，and a Turing Centenary Fellowship．
    ${ }^{\dagger}$ Institute for Advanced Study and Princeton University．Email：moitra＠ias．edu．Research supported in part by NSF grant No．DMS－0835373 and by an NSF Computing and Innovation Fellowship．

[^1]:    ${ }^{1}$ Note that even if a polynomial-sized extended formulation did exist, it is not necessarily easy to compute what it is. So in general the existence of a polynomial-sized extended formulation would only imply that if we are given advice about what it is, we could solve $N P$-hard problems.

[^2]:    ${ }^{2}$ Yannakakis [31] proved that there is no polynomial-sized symmetric extended formulation for matching.
    ${ }^{3}$ There is a combinatorial algorithm for this problem. And in fact there is also a polynomial time algorithm based on a linear program: even though the characterization of the matching polytope given by Edmonds [12] has exponential size, there is a polynomial time separation oracle for these constraints [26] and hence matching with general weights can be solved by the ellipsoid algorithm. This illustrates a subtlety in not only our lower bounds for extended formulationss but all such works; these lower bounds show that there is no polynomial-sized extended formulation, but it still could be the case that there is an exponential-sized extended formulation that has an efficient separation oracle.

[^3]:    ${ }^{4}$ If a protocol gets $\epsilon$-advantage and requires $o\left(\epsilon^{2} n\right)$ bits of communication, then we can run the protocol $O\left(1 / \epsilon^{2}\right)$ times and take the majority vote to get constant advantage, but this would violate the linear lower bound for disjointness.

[^4]:    ${ }^{5}$ Here, one player is given a row and one player is given a column, and their goal is to compute the value of corresponding entry in expectation, using only nonnegative values.

[^5]:    ${ }^{6}$ We would like to thank Thomas Watson for pointing out an oversight in our earlier version. We had asserted that $N^{t}$ is rank one, but had neglected to condition on $C_{j}$ as we had intended to do. Indeed, this is the same conditioning trick that we make use of in our later applications to extended formulations, and is by now standard in information complexity (see e.g. [9]).

