

On the Separation of Correlation-Assisted Sum Capacities of Multiple Access Channels

Akshay Seshadri¹, Felix Leditzky¹, Vikesh Siddhu¹, and Graeme Smith

Abstract—The capacity of a channel characterizes the maximum rate at which information can be transmitted through the channel asymptotically faithfully. For a channel with multiple senders and a single receiver, computing its sum capacity is possible in theory, but challenging in practice because of the nonconvex optimization involved. To address this challenge, we investigate three topics in our study. In the first part, we study the sum capacity of a family of multiple access channels (MACs) obtained from nonlocal games. For any MAC in this family, we obtain an upper bound on the sum rate that depends only on the properties of the game when allowing assistance from an arbitrary set of correlations between the senders. This approach can be used to prove separations between sum capacities when the senders are allowed to share different sets of correlations, such as classical, quantum or no-signalling correlations. We also construct a specific nonlocal game to show that the approach of bounding the sum capacity by relaxing the nonconvex optimization can give arbitrarily loose bounds. Owing to this result, in the second part, we study algorithms for non-convex optimization of a class of functions we call Lipschitz-like functions. This class includes entropic quantities, and hence these results may be of independent interest in information theory. Subsequently, in the third part, we show that one can use these techniques to compute the sum capacity of an arbitrary two-sender MACs to a fixed additive precision in quasi-polynomial time. We showcase our method by efficiently computing the sum capacity of a family of two-sender MACs for which one of the input alphabets has size two. Furthermore, we demonstrate with an example that our algorithm may compute the sum capacity to a higher precision than using the convex relaxation.

Manuscript received 21 September 2022; revised 31 March 2023; accepted 22 April 2023. Date of publication 15 May 2023; date of current version 18 August 2023. This work was supported in part by NSF Grant CCF 1652560, in part by NSF Grant Physics (PHY) 1915407, in part by NSF Quantum Interconnect Challenges for Transformational Advances in Quantum Systems (QuIC-TAQS) under Grant 2137984, and in part by NSF QuIC-TAQS under Grant 2137953. Akshay Seshadri acknowledges support from the Professional Research Experience Program (PREP) operated jointly by the National Institute of Standards and Technology (NIST) and the University of Colorado. An earlier version of this paper was presented at the 2022 IEEE International Symposium on Information Theory [DOI: 10.1109/ISIT50566.2022.9834459]. (Corresponding author: Akshay Seshadri.)

Akshay Seshadri is with the Department of Physics, University of Colorado Boulder, Boulder, CO 80309 USA, and also with the JILA, University of Colorado/NIST, Boulder, CO 80309 USA (e-mail: akshay.seshadri@colorado.edu).

Felix Leditzky is with the Department of Mathematics and IQUIST, University of Illinois Urbana-Champaign, Urbana, IL 61801 USA.

Vikesh Siddhu was with the JILA, University of Colorado/NIST, Boulder, CO 80309 USA.

Graeme Smith is with the Department of Physics, University of Colorado Boulder, Boulder, CO 80309 USA, also with the JILA, University of Colorado/NIST, Boulder, CO 80309 USA, and also with the Center for Theory of Quantum Matter, University of Colorado Boulder, Boulder, CO 80309 USA.

Communicated by S. Mancini, Associate Editor for Quantum.

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TIT.2023.3274434>.

Digital Object Identifier 10.1109/TIT.2023.3274434

Index Terms—Channel capacity, multiuser channels, optimization, quantum entanglement, quantum information science.

I. INTRODUCTION

STUDYING information transmission over a noisy channel is of fundamental importance in communication theory. In his landmark paper, Shannon studied the rate of transmission over a channel having a single sender and a single receiver [2]. The maximum number of bits per use of channel that can be transmitted through the channel with asymptotically vanishing error is known as the capacity of the channel [3]. Shannon showed that the channel capacity can be calculated by maximizing the mutual information between the input and the output of the channel over all possible input probability distributions [2], [3]. This is a convex optimization problem that can be solved using standard optimization techniques [4] or more specialized methods like the Arimoto-Blahut algorithm [5], [6].

The importance of Shannon's work was soon recognized, and consequently, Shannon's ideas were generalized in various ways. In this study, we focus on one such generalization, namely a channel that has multiple senders and a single receiver, commonly known as a multiple access channel (MAC). For such a channel, a tuple of rates is called achievable if each sender can send information through their respective input at their respective rate such that the total error of transmission vanishes asymptotically. The set of all such achievable rate tuples is called the capacity region, and Ahlswede [7] and Liao [8] were the first to give an entropic characterization of it. The total rate at which asymptotically error-free transmission through a MAC is possible is called the sum capacity of the MAC.

While there has been much research on MACs since the work of Ahlswede & Liao, there are no efficient (polynomial time) algorithms to date that compute the sum capacity of a MAC. This difficulty stems from the fact that, unlike the computation of the capacity of a point-to-point channel, the optimization involved in computing the sum capacity is non-convex [9], [10]. Proposals to solve this nonconvex problem efficiently were found to be unsuitable. On the contrary, it was recently shown that computing the sum capacity to a precision that scales inversely with the cube of the input alphabet size is an NP-hard problem [11]. Therefore, one should not expect efficient general-purpose algorithms for computing the capacity region or the sum capacity of a MAC.

Since computing the sum capacity of a MAC is a hard nonconvex problem, a common approach that is adopted to

circumvent this optimization is to relax the optimization to obtain a convex problem. The relaxation gives an upper bound on the sum capacity which can be efficiently computed. Such an approach was adopted, for example, by Calvo et al. [9]. Since we generally expect that there are no efficient algorithms to compute the sum capacity, quantifying the performance of such an upper bound becomes important and essential.

We undertake this task of elucidating the performance of the upper bound on sum capacity obtained through a convex relaxation. For this purpose, we consider MACs that are constructed from nonlocal games. Such MACs were introduced by Leditzky et al. [11] based on previous work of Quek & Shor [12], and subsequently generalized in [13]. We present an analytical upper bound on the sum capacity of such MACs. Our bound extends to cases where the senders of these MACs can share arbitrary correlations, such as classical, quantum, and no-signalling correlations. Our upper bound depends only on the number of question tuples in the nonlocal game and the maximum winning probability of the game when the questions are drawn uniformly and answers are obtained using the strategies allowed by the shared correlations. Using these bounds, we obtain separations between the sum rate obtained from using different sets of correlations. These separations help distinguish the ability of these correlations to assist communication in MAC coding scenarios. In particular, our bound gives a separation between the sum capacity and the entanglement-assisted sum rate for the MAC obtained from the magic square game. The separation found here is roughly 5 times larger than the previously reported 1% separation in [11]. Furthermore, using our bounds, we show how prior bounds on the sum capacity obtained using convex relaxation [9] can be arbitrarily loose.

Our result highlights the need to find better techniques to bound the sum capacity from above. We take a step in this direction by showing that computing the sum capacity is equivalent to optimizing a Lipschitz-like function. Thereupon, we present some algorithms for optimization of Lipschitz-like functions, which may be of independent interest. We show that these algorithms can compute the sum capacity of two-sender MACs to a given additive precision in quasi-polynomial time. Instead of a fixed precision, one of our algorithms can also accept a fixed number of iterations as an input and output an upper bound on the sum capacity. In particular, for a specific family of two-sender MACs that includes binary MACs, the number of iterations required for convergence grows at most polynomially with the dimensions and inverse precision. Thus, while it might not be possible to efficiently compute the sum capacity for an arbitrary MAC in practice, we can nevertheless efficiently compute the sum capacity for a large family of MACs.

A. Organization of the Paper

We organize the paper in three parts, each one presenting different contributions of this work that may be of independent interest. Focus of the first part (Sec. III) is MACs obtained from non-local games. We investigate the advantages of sharing correlations between various senders of a nonlocal games MAC. In Sec. III-A, we obtain an upper bound on the sum

capacity of MACs constructed from nonlocal games, and show separations between different sets of correlations. In Sec. III-B we accomplish another primary goal of our study, showing that the convex relaxation of the sum capacity can be arbitrarily loose.

Motivated by this result, in the second part (Sec. IV) we develop methods to solve non-convex problems like sum capacity computation. To this end, we define and study optimization of Lipschitz-like functions in Sec. IV-A. We then give an overview of the optimization algorithms for Lipschitz-like functions featured in this work in Sec. IV-B, and discuss them in greater detail in the following sections. That is, in Sec. IV-C and Sec. IV-D we develop algorithms for global optimization of such functions over a closed interval and over the standard simplex respectively. Results on optimization of Lipschitz-like functions over arbitrary compact and convex domains can be found in App. C-D. We also present relevant convergence guarantees and complexity analysis for our algorithms. In this way, we generalize certain prior algorithms for optimizing Lipschitz continuous functions. Our generalized algorithm may be of independent interest for non-convex optimization in information theory. For instance, we show that some entropic quantities that may be derived from Shannon and von Neumann entropies are Lipschitz-like functions.

The third part (Sec. V) applies algorithms developed in the second part for computing the sum capacity of an arbitrary two-sender MAC. In Sec. V-A, we prove that the sum capacity computation can be viewed as an optimization of a Lipschitz-like function. In Sec. V-B, we develop an efficient algorithm to compute the sum capacity of a large family of two-sender MACs, where one of the input alphabets is of size 2. Subsequently, in Sec. V-C, we present algorithms that can compute or upper-bound the sum capacity of an arbitrary two-sender MAC, along with a detailed complexity analysis. Finally, in Sec. V-D, we construct examples which demonstrate that our algorithm performs provably better than convex relaxation in computing the sum capacity.

We have tried to minimize the overlap between these three parts of our study, so that any one of these parts may be read independently without loss of continuity. In the next section, we provide a brief overview of notations, quantum states and measurements, nonlocal games, and multiple access channels. We recommend readers mainly interested in parts I and III to read Sec. II-C on multiple access channels as it introduces definitions used in these parts.

II. PRELIMINARIES

We briefly review some concepts that are used later. A short summary of the notation used throughout is given below.

We denote the set of natural numbers as $\mathbb{N} = \{0, 1, \dots\}$ and the set of positive integers as $\mathbb{N}_+ = \{1, 2, \dots\}$. For any integer $N \geq 1$, let $[N] := \{1, \dots, N\}$. We denote the $(n-1)$ -dimensional standard simplex by $\Delta_n = \{x \in \mathbb{R}^n \mid x \geq 0, \sum_{i=1}^n x_i = 1\}$. When $x \in \mathbb{R}^n$ is a vector, we interpret the inequality $x \geq 0$ component-wise, i.e., $x_i \geq 0$ for all $i \in [n]$. We denote the non-negative orthant in n -dimensional Euclidean space by $\mathbb{R}_+^n = \{x \in \mathbb{R}^n \mid x \geq 0\}$. The Euclidean inner product between two vectors $x, y \in \mathbb{R}^n$ is denoted by

$\langle x, y \rangle$. The space of $(m \times n)$ -matrices with entries in \mathbb{C} is denoted by $\mathbb{C}^{m \times n}$. The Kronecker product of two matrices A and B is denoted by $A \otimes B$.

For a random variable X taking values in a finite alphabet \mathcal{X} , we denote by $p_X(x)$ the probability of X taking the value $x \in \mathcal{X}$. The Shannon entropy of the random variable X is denoted by $H(X) = -\sum_{x \in \mathcal{X}} p_X(x) \log(p_X(x))$. The unit of entropy is referred to as bits when the base of the logarithm in $H(X)$ is 2, whereas the unit is referred to as nats when the base of the logarithm is e . The mutual information between two random variables X and Y is defined as $I(X; Y) = H(X) + H(Y) - H(X, Y)$. A single-input single-output channel is a triple $(\mathcal{X}, \mathcal{Z}, \mathcal{N})$ consisting of an input alphabet \mathcal{X} , an output alphabet \mathcal{Z} and a probability transition matrix $\mathcal{N}(z|x)$ giving the probability of transmitting $z \in \mathcal{Z}$ given the input $x \in \mathcal{X}$. The capacity of the channel \mathcal{N} is given by the single-letter formula $C(\mathcal{N}) = \max_{p_X} I(X; Z)$, where X and Z are the random variables describing the input and output of \mathcal{N} , respectively [3].

A. Quantum States and Measurements

A quantum state (or density matrix) $\rho \in \mathbb{C}^{d \times d}$ is a self-adjoint, positive semi-definite (PSD) matrix with unit trace [14]. A measurement of this quantum state can be described by a positive operator-valued measure (POVM), a collection of PSD matrices $\{F_1, \dots, F_M\}$ of size $d \times d$ satisfying $\sum_{m=1}^M F_m = \mathbb{I}$, where \mathbb{I} is the identity matrix on the Hilbert space \mathbb{C}^d . Each POVM element F_m is associated to a measurement outcome $m \in [M]$, which is obtained with probability $\text{Prob}(m) = \text{Tr}(\rho F_m)$ [14].

In quantum mechanics, a Hermitian operator $\mathcal{O} \in \mathbb{C}^{d \times d}$ is often called an *observable*. It can be measured in the following sense. First, we write the spectral decomposition of \mathcal{O} as $\mathcal{O} = \sum_{m=1}^M \lambda_m \mathbb{P}_m$, where $\mathbb{P}_m \in \mathbb{C}^{d \times d}$ is the projector onto the eigenspace of \mathcal{O} corresponding to the eigenvalue $\lambda_m \in \mathbb{R}$ for $m = 1, \dots, M$. Then, the POVM associated with measuring the observable \mathcal{O} is given by $\{\mathbb{P}_m \mid m \in [M]\}$ [14]. For qubits, i.e., two-dimensional quantum systems described by quantum states $\rho \in \mathbb{C}^{2 \times 2}$, the Pauli matrices

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \text{ and } \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad (1)$$

represent three commonly used observables associated to the spin of an electron along different axes.

Let two parties, Alice and Bob, share a quantum state ρ . If they perform local measurements with POVMs $\{A_1, \dots, A_I\}$ and $\{B_1, \dots, B_J\}$, respectively, then the probability that Alice observes the outcome i and Bob observes the outcome j is given by $\text{Prob}(i, j) = \text{Tr}[\rho(A_i \otimes B_j)]$. In other words, the overall POVM for the measurement is given by $\{A_i \otimes B_j \mid i \in [I], j \in [J]\}$.

B. Nonlocal Games

A nonlocal game is played between two players who each receive a question from a referee that they need to answer. The players are not allowed to communicate with each other during the game. Prior to starting the game, one fixes the set

of questions and answers and a winning criterion, which are known to everyone. A referee then randomly draws questions and hands them out to the players. If the answers of the players satisfy the winning condition, they win the game.

Formally, an N -player promise-free nonlocal game G is a tuple $G = (\mathcal{X}_1, \dots, \mathcal{X}_N, \mathcal{Y}_1, \dots, \mathcal{Y}_N, \mathcal{W})$, where \mathcal{X}_i and \mathcal{Y}_i are the question and answer set for the i th player, respectively, and the winning condition $\mathcal{W} \subseteq (\mathcal{X}_1 \times \dots \times \mathcal{X}_N) \times (\mathcal{Y}_1 \times \dots \times \mathcal{Y}_N)$ determines the tuples of questions and answers that win the game [15]. Throughout this study, we restrict our attention to the case when \mathcal{X}_i and \mathcal{Y}_i are finite sets for all $i \in [N]$. Unless stated otherwise, we will always refer to promise-free nonlocal games as nonlocal games.¹

For convenience, we denote the question set as $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_N$ and the answer set as $\mathcal{Y} = \mathcal{Y}_1 \times \dots \times \mathcal{Y}_N$. Given any question tuple $\mathbf{x} = (x_1, \dots, x_N) \in \mathcal{X}$, the i th question is given by $x_i \in \mathcal{X}_i$, and a similar notation is used for answers.

A strategy for the game $G = (\mathcal{X}, \mathcal{Y}, \mathcal{W})$ is any conditional probability distribution $p_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{x})$ on the answer tuples $\mathbf{y} \in \mathcal{Y}$ given a question tuple $\mathbf{x} \in \mathcal{X}$. There are different types of strategies that one can consider, depending on whether the players are allowed to use shared correlations after the game starts. We list a few strategies that are central to our study.

- 1) *Classical strategy*: This is the typical setup of a nonlocal game. Given a question $x_i \in \mathcal{X}_i$, the i th player decides an answer as per the probability distribution $p_{Y_i|X_i}(y_i|x_i)$, $i \in [N]$. This gives rise to a classical strategy

$$\begin{aligned} p_{\mathbf{Y}|\mathbf{X}}(y_1, \dots, y_N | x_1, \dots, x_N) \\ = p_{Y_1|X_1}(y_1|x_1) \cdots p_{Y_N|X_N}(y_N|x_N) \end{aligned} \quad (2)$$

In the general setting the players are allowed to choose such a strategy on the basis of the outcome of a random variable shared by all players, which is called a probabilistic strategy. Therefore, the set of classical strategies \mathfrak{S}_{cl} corresponds to the convex hull of conditional probability distributions of the form given in Eq. (2) [16]. Operationally, convex combinations of product distributions correspond to shared randomness, and hence the set of classical correlations we use is similar to those defined by local hidden variable theories [16], [17]. However, when the questions are drawn uniformly at random, the classical strategy maximizing the probability of winning the game is of the form given in Eq. (2), with $p_{Y_i|X_i}(y_i|x_i) = \delta_{y_i, f_i(x_i)}$ for $i \in [N]$ [15]. Here, $f_i: \mathcal{X}_i \rightarrow \mathcal{Y}_i$ is a function that outputs an answer given a question. Such a strategy is called a deterministic strategy.

- 2) *Quantum strategy*: The players have access to a shared quantum state ρ , but cannot communicate otherwise. Given a question $x_i \in \mathcal{X}_i$, the i th player performs a local measurement with some POVM $\{E_y^{(x_i)} \mid y \in \mathcal{Y}_i\}$. Subsequently, a quantum strategy is described by the

¹One may consider nonlocal games for which the possible question tuples are restricted to a subset $P \subseteq \mathcal{X}_1 \times \dots \times \mathcal{X}_N$ called a *promise*. A promise-free nonlocal game as defined above is one with $P = \mathcal{X}_1 \times \dots \times \mathcal{X}_N$. Any nonlocal game $G = (\mathcal{X}_1, \dots, \mathcal{X}_N, \mathcal{Y}_1, \dots, \mathcal{Y}_N, \mathcal{W})$ with promise P can be turned into a promise-free game by defining a new winning condition $\mathcal{W}' = \mathcal{W} \cup [(\mathcal{X}_1 \times \dots \times \mathcal{X}_N) \setminus P] \times (\mathcal{Y}_1 \times \dots \times \mathcal{Y}_N)$, i.e., the players automatically win on question tuples not contained in the promise.

probability distribution

$$\begin{aligned} p_{\mathbf{Y}|\mathbf{X}}(y_1, \dots, y_N | x_1, \dots, x_N) \\ = \text{Tr} \left[\rho \left(E_{y_1}^{(x_1)} \otimes \dots \otimes E_{y_N}^{(x_N)} \right) \right] \end{aligned} \quad (3)$$

We denote the set of quantum strategies by \mathfrak{S}_Q . For any given classical strategy, one can construct a quantum state and POVMs so that the quantum strategy reduces to the given classical strategy, and therefore, $\mathfrak{S}_{\text{cl}} \subseteq \mathfrak{S}_Q$ [16].

3) *No-signalling strategy*: A strategy $p_{\mathbf{Y}|\mathbf{X}}(y_1, \dots, y_N | x_1, \dots, x_N)$ is said to be no-signalling if

$$p_{Y_i|\mathbf{X}}(y_i | x_1, \dots, x_N) = p_{Y_i|X_i}(y_i | x_i), \quad i \in [N] \quad (4)$$

for all x_1, \dots, x_N [18]. Informally, this means that players must respect locality, i.e., no information can be transmitted between players “faster than light”. As a consequence, the strategy used by each player cannot depend on the questions received by the other players. We denote the set of no-signalling strategies by \mathfrak{S}_{NS} . Both classical and quantum strategies are no-signalling, and therefore, $\mathfrak{S}_{\text{cl}} \subseteq \mathfrak{S}_Q \subseteq \mathfrak{S}_{\text{NS}}$.

4) *Full communication*: When we impose no restriction on the distribution $p_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{x})$, we are implicitly allowing for communication between the players *after* they have received the questions. The set of all possible conditional probability distributions $p_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{x})$ is denoted by $\mathfrak{S}_{\text{all}}$, which corresponds to allowing full communication between the players. This contains all classical, quantum and no-signalling strategies. Communication between the players is usually not allowed in nonlocal games, but we introduce this setting here for later use nevertheless.

In summary, we have the following hierarchy of correlations: $\mathfrak{S}_{\text{cl}} \subseteq \mathfrak{S}_Q \subseteq \mathfrak{S}_{\text{NS}} \subseteq \mathfrak{S}_{\text{all}}$.

Suppose that the questions are drawn randomly from the set \mathcal{X} as per the probability distribution $\pi(\mathbf{x})$. Say the players obtain their answers using strategies in some set \mathfrak{S} . Then the maximum winning probability of the game G is given by

$$\omega_{\pi}^{\mathfrak{S}}(G) = \sup_{p_{\mathbf{Y}|\mathbf{X}} \in \mathfrak{S}} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{W}} \pi(\mathbf{x}) p_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{x}). \quad (5)$$

Notice in particular that, if $\mathfrak{S} \subseteq \mathfrak{S}'$, then $\omega_{\pi}^{\mathfrak{S}}(G) \leq \omega_{\pi}^{\mathfrak{S}'}(G)$. Hence, as we go from classical to quantum to no-signalling strategies, the winning probability never decreases.

We will mainly be concerned with the scenario where the questions are drawn uniformly at random, and so $\pi = \pi_U$ is the uniform distribution on \mathcal{X} with $\pi_U(\mathbf{x}) = 1/|\mathcal{X}|$ for all $\mathbf{x} \in \mathcal{X}$. In this case, we use the notation $\omega_{\pi_U}^{\mathfrak{S}}(G) = \omega^{\mathfrak{S}}(G)$. The maximum winning probability when the questions are drawn uniformly and answers are generated using classical strategies is written as $\omega^{\text{cl}}(G)$. Similarly, when using quantum strategies, we write $\omega^Q(G)$, and when using no-signalling strategies, we write $\omega^{\text{NS}}(G)$.

C. Multiple Access Channels

A discrete memoryless multiple access channel without feedback, simply referred to as multiple access channel in this

study, is a tuple $(\mathcal{B}_1, \dots, \mathcal{B}_N, \mathcal{Z}, \mathcal{N}(z|b_1, \dots, b_N))$ consisting of input alphabets $\mathcal{B}_1, \dots, \mathcal{B}_N$, an output alphabet \mathcal{Z} , and a probability transition matrix $\mathcal{N}(z|b_1, \dots, b_N)$ that gives the probability that the channel output is $z \in \mathcal{Z}$ when the input is $b_1 \in \mathcal{B}_1, \dots, b_N \in \mathcal{B}_N$ [3]. For simplicity of notation, we will denote a MAC by the probability transition matrix \mathcal{N} when the input and output alphabets are understood.

Since transmission over the channel is error-prone, one usually encodes the messages and transmits them over multiple uses of the channel, and subsequently, the transmitted symbols are decoded to reconstruct the original message. Formally, an $((M_1, \dots, M_N), n)$ -code for a multiple access channel consists of message sets $\mathcal{M}_i = \{1, \dots, M_i\}$ for $i \in [N]$, encoding functions $g_i^{(e)}: \mathcal{M}_i \rightarrow \mathcal{B}_i^{n \times n}$ for $i \in [N]$, and a decoding function $g^{(d)}: \mathcal{Z}^n \rightarrow \mathcal{M}_1 \times \dots \times \mathcal{M}_N \cup \{\text{err}\}$, where err is an error symbol [3]. For convenience, we denote the message set as $\mathcal{M} = \mathcal{M}_1 \times \dots \times \mathcal{M}_N$. The performance of the code can be quantified by the average probability of error in reconstructing the message,

$$P_e^{(n)} = \frac{1}{|\mathcal{M}|} \sum_{m \in \mathcal{M}} \text{Prob}\{g^{(d)}(\mathcal{Z}^n) \neq m \mid \text{message } m \text{ was sent}\}.$$

Here we assume that the messages are chosen uniformly at random and transmitted through the channel. Note that the above code uses the channel n times. We say that a rate tuple (R_1, \dots, R_N) is *achievable* if there exists a sequence of codes $(([2^{nR_1}], \dots, [2^{nR_N}]), n)$ such that $P_e^{(n)} \rightarrow 0$ as $n \rightarrow \infty$ [3].

The *capacity region* of a multiple access channel is defined as the closure of the set of achievable rate tuples (R_1, \dots, R_N) . For $N = 2$ one obtains a two-sender MAC. Ahlswede [7] and Liao [8] were the first to give a so-called single-letter characterization of the capacity region of a two-sender MAC. The capacity region for an N -sender MAC \mathcal{N} can be written as

$$\begin{aligned} \text{Cap}(\mathcal{N}) = \text{conv} \left\{ (R_1, \dots, R_N) \mid \right. \\ \left. 0 \leq \sum_{j \in J} R_j \leq I(B_J; Z|B_{J^c}) \quad \forall \emptyset \neq J \subseteq [N], \right. \\ \left. p(b_1, \dots, b_N) = p(b_1) \dots p(b_N) \right\} \end{aligned} \quad (6)$$

where $p(b_1, \dots, b_N)$ is a probability distribution over the joint random variable B_1, \dots, B_N describing the channel input, $p(b_i)$ is a probability distribution for the random variable B_i corresponding to the i -th sender’s input, random variable Z describes the channel output, and for any set $J \subseteq [N]$, $B_J = \{B_j \mid j \in J\}$ [9]. Crucially, the optimization in Eq. (6) restricts the joint input random variable to have a product distribution, i.e., $p(b_1, \dots, b_N) = p(b_1) \dots p(b_N)$. The main quantity of interest in our study is the *sum capacity* of a MAC [9],

$$S(\mathcal{N}) = \sup \left\{ \sum_{i=1}^N R_i \mid (R_1, \dots, R_N) \in \text{Cap}(\mathcal{N}) \right\}$$

$$\begin{aligned}
&= \max_{p(b_1, \dots, b_N)} I(B_1, \dots, B_N; Z) \\
\text{s.t. } & p(b_1, \dots, b_N) = p(b_1) \cdots p(b_N). \quad (7)
\end{aligned}$$

Informally, the sum capacity represents the maximum sum rate at which the senders can send information through the MAC such that the transmission error vanishes asymptotically. Because the maximization involved in computing the sum capacity is constrained to be over product distributions on the input, the resulting optimization problem is nonconvex. This nonconvexity is the main source of difficulty in computing the sum capacity of a MAC in practice [10], [11]. A common approach to avoid this difficult is to relax the nonconvex constraint and maximize over all possible probability distributions on the input. Such an approach was adopted by Calvo et al. [9], leading to the upper bound

$$C(\mathcal{N}) = \max_{p(b_1, \dots, b_N)} I(B_1, \dots, B_N; Z) \quad (8)$$

on the sum capacity $S(\mathcal{N})$, where we now maximize over *arbitrary* input probability distributions. Note that $C(\mathcal{N})$ is the capacity of the channel \mathcal{N} when we think about it as a single-input single-output channel. Since it is a relaxation of the sum capacity corresponding to the capacity region, we call C the *relaxed sum capacity*.

III. NONLOCAL GAMES MAC AND RELAXED SUM CAPACITY

A. Bounding the Sum Capacity of MACs From Nonlocal Games

Suppose we are given a promise-free nonlocal game G with question sets $\mathcal{X}_1, \dots, \mathcal{X}_N$, answer sets $\mathcal{Y}_1, \dots, \mathcal{Y}_N$, and a winning condition $\mathcal{W} \subseteq (\mathcal{X}_1 \times \dots \times \mathcal{X}_N) \times (\mathcal{Y}_1 \times \dots \times \mathcal{Y}_N)$. Following Leditzky et al. [11], we construct a MAC \mathcal{N}_G with input alphabets $\mathcal{B}_i = \mathcal{X}_i \times \mathcal{Y}_i$ for $i \in [N]$, output alphabet $\mathcal{Z} = \mathcal{X}_1 \times \dots \times \mathcal{X}_N$, and a probability transition matrix

$$\begin{aligned}
&\mathcal{N}_G(\hat{x}_1, \dots, \hat{x}_N | x_1, y_1, \dots, x_N, y_N) \\
&= \begin{cases} \delta_{\hat{x}_1, x_1} \cdots \delta_{\hat{x}_N, x_N} & (x_1, \dots, x_N, y_1, \dots, y_N) \in \mathcal{W} \\ \frac{1}{d} & (x_1, \dots, x_N, y_1, \dots, y_N) \notin \mathcal{W}, \end{cases} \quad (9)
\end{aligned}$$

where $d = |\mathcal{Z}| = |\mathcal{X}_1 \times \dots \times \mathcal{X}_N|$.

In words, the channel \mathcal{N}_G takes question-answer pairs from each player as input, and if they win the game, then the questions are transmitted without any noise. However, if they lose the game, a question tuple chosen uniformly at random is output by the channel. For convenience, we will denote the input to the MAC \mathcal{N}_G as $\mathcal{XY} = (\mathcal{X}_1 \times \mathcal{Y}_1) \times \dots \times (\mathcal{X}_N \times \mathcal{Y}_N)$, and write $\mathbf{xy} = (x_1, y_1, \dots, x_N, y_N)$. We will usually abbreviate the phrase “MAC obtained from a nonlocal game” to “nonlocal game MAC” or NG-MAC.

Before diving into any technical details, we explain why such MACs are suitable for obtaining bounds on the sum capacity that are better than the relaxed sum capacity. We begin by noting that the sum capacity of the NG-MAC \mathcal{N}_G can be written as

$$S(\mathcal{N}_G) = \max_{p^{(1)}(x_1, y_1) \cdots p^{(N)}(x_N, y_N)} I(X_1, Y_1, \dots, X_N, Y_N; Z)$$

where (X_i, Y_i) is the random variable (with distribution $p^{(i)}$) describing the i th input and Z is the random variable describing the output of \mathcal{N}_G . Note that $p^{(i)}$ is a probability distribution over the question-answer pairs of the i th player. By writing $p^{(i)}(x_i, y_i) = \pi^{(i)}(x_i) p_{Y_i|X_i}(y_i|x_i)$, we can break $p^{(i)}$ into a distribution $\pi^{(i)}$ over questions and a strategy $p_{Y_i|X_i}$ chosen by the i th player. As a result, we can write $p^{(1)} \cdots p^{(N)} = \pi p_{\mathbf{Y}|\mathbf{X}}$, where $\pi = \pi^{(1)} \cdots \pi^{(N)}$ is some distribution over the questions and $p_{\mathbf{Y}|\mathbf{X}} = p_{Y_1|X_1} \cdots p_{Y_N|X_N}$ is a classical strategy chosen by the players. Such a decomposition allows us to optimize separately over questions and strategies. By performing suitable relaxations, we can obtain a bound on the sum capacity that depends only on the winning probability of the game (see Thm. (5) for a precise statement). In fact, such a proof technique allows us to bound the sum capacity even when the players are allowed to use different sets of strategies such as those obtained from quantum or no-signalling correlations. The resulting bound is helpful in obtaining separations between the communication capabilities of different sets of correlations.

In contrast, the relaxed sum capacity $C(\mathcal{N}_G)$ is computed by maximizing over all possible probability distributions. For a nonlocal game MAC, this amounts to maximizing over all distributions over the questions and all possible strategies (allowing full communication between the players). Using such strategies, the players will always win the game, assuming that the game has at least one correct answer for every question. This results in the trivial bound $C(\mathcal{N}_G) = \ln(d)$ on the sum capacity, since the players can always win the game. On the other extreme, if the winning probability of the game is zero, then we have $S(\mathcal{N}_G) = 0$. Following this line of thinking, we construct a game in Section III-B that allows us to give an arbitrarily large separation between $S(\mathcal{N}_G)$ and $C(\mathcal{N}_G)$.

The above discussion suggests that the sum capacity of the NG-MAC \mathcal{N}_G increases with the winning probability of the game G , an observation that was also noted by Leditzky et al. [11]. On the other hand, we know that the winning probability of the game can increase if we allow the players to use a larger set of strategies. This motivates us to allow the senders to share some set of correlations to play the game so as to increase the sum capacity of NG-MAC.

1) *Correlation Assistance:* By a correlation, we mean a probability distribution $P(y'|xy)$, where $xy \in \mathcal{XY}$ is the input to the NG-MAC, while $y' \in \mathcal{Y}$ is an answer to the nonlocal game G .² We allow the N senders to share the correlation P and perform local post-processing of the answer generated by P to obtain their final answer for the game for the input questions. This post-processing can be expressed as a probability distribution $f_i(\bar{y}_i|x_i, y_i, y'_i)$ over the answers $\bar{y}_i \in \mathcal{Y}_i$ given the input question-answer pair $(x_i, y_i) \in \mathcal{X}_i \times \mathcal{Y}_i$ and the answer $y'_i \in \mathcal{Y}_i$ generated by the correlation P . For

²The correlation P may depend on the answers $y \in \mathcal{Y}$ in addition to the questions $x \in \mathcal{X}$ because we wish to construct a larger MAC that has the same structure as the NG-MAC \mathcal{N}_G and has assistance from the correlation P . See Fig. 1 for example. In practice, the correlation will often produce the answer tuple y' solely from the question tuple x , using some strategy for the non-local game G .

convenience, denote

$$f(\bar{\mathbf{y}}|\mathbf{xy}, \mathbf{y}') = \prod_{i=1}^N f_i(\bar{y}_i|x_i, y_i, y'_i), \quad (10)$$

which is a distribution over the answers $\bar{\mathbf{y}} \in \mathcal{Y}$ given input question-answer pairs $\mathbf{xy} \in \mathcal{XY}$ and answers $\mathbf{y}' \in \mathcal{Y}$ generated by the correlation P .

This gives rise to the channel $\mathcal{A}_{P,f}$ having input and output alphabets \mathcal{XY} and the probability transition matrix

$$\mathcal{A}_{P,f}(\bar{\mathbf{xy}}|\mathbf{xy}) = \delta_{\bar{\mathbf{x}}, \mathbf{x}} \sum_{\mathbf{y}' \in \mathcal{Y}} f(\bar{\mathbf{y}}|\mathbf{xy}, \mathbf{y}') P(\mathbf{y}'|\mathbf{xy}). \quad (11)$$

This probability transition matrix gives a probability distribution over the question-answer pairs $\bar{\mathbf{xy}} \in \mathcal{XY}$ given some input question-answer pair $\mathbf{xy} \in \mathcal{XY}$. The probability distributions f_i denote local post-processing by the i th sender to generate the final answer \bar{y}_i . Note that the definition of $\mathcal{A}_{P,f}$ in Eq. (11) ensures that the input questions $\mathbf{x} \in \mathcal{X}$ are transmitted without any change, i.e., $\bar{\mathbf{x}} = \mathbf{x}$. The post-processings are only used to generate the final answers. This definition is slightly more restrictive than the one given in Leditzky et al. [11], where post-processing of both questions and answers is considered. We impose such a restriction on post-processing so as to ensure that the set of strategies induced by the channel $\mathcal{A}_{P,f}$ has a close relation to the set of correlations \mathcal{C} shared by the senders. For example, classical or quantum correlations shared by the senders give rise to classical or quantum strategies to play the nonlocal game G . We make this idea precise in the following discussion.

An input probability distribution $p(\mathbf{xy})$ over the question-answer pairs \mathcal{XY} can be written as $p(\mathbf{xy}) = \pi(\mathbf{x})p_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{x})$, a product of probability distribution over the questions $\pi(\mathbf{x})$ and a strategy $p_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{x})$. This input probability distribution is modified by the channel $\mathcal{A}_{P,f}$ to give

$$\bar{p}(\bar{\mathbf{xy}}) = \sum_{\mathbf{xy} \in \mathcal{XY}} \mathcal{A}_{P,f}(\bar{\mathbf{xy}}|\mathbf{xy}) p(\mathbf{xy}) = \pi(\bar{\mathbf{x}}) \bar{p}_{\mathbf{Y}|\mathbf{X}}(\bar{\mathbf{y}}|\bar{\mathbf{x}}), \quad (12)$$

where the new strategy $\bar{p}_{\mathbf{Y}|\mathbf{X}}$ is given by

$$\begin{aligned} \mathcal{A}_{P,f}(p_{\mathbf{Y}|\mathbf{X}}) &:= \bar{p}_{\mathbf{Y}|\mathbf{X}}(\bar{\mathbf{y}}|\bar{\mathbf{x}}) \\ &= \sum_{\mathbf{y} \in \mathcal{Y}} \sum_{\mathbf{y}' \in \mathcal{Y}} f(\bar{\mathbf{y}}|\bar{\mathbf{xy}}, \mathbf{y}') P(\mathbf{y}'|\bar{\mathbf{xy}}) p_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\bar{\mathbf{x}}). \end{aligned} \quad (13)$$

We use the notation $\mathcal{A}_{P,f}(p_{\mathbf{Y}|\mathbf{X}})$ to emphasize that $\bar{p}_{\mathbf{Y}|\mathbf{X}}$ is the strategy induced by the action of $\mathcal{A}_{P,f}$ on $p_{\mathbf{Y}|\mathbf{X}}$. Therefore, the channel $\mathcal{A}_{P,f}$ modifies the input strategy $p_{\mathbf{Y}|\mathbf{X}}$ by incorporating assistance from the correlation P and local post-processings f . For this reason, we call $\mathcal{A}_{P,f}$ the *correlation-assistance channel*.

If the senders have access to some set of correlations \mathcal{C} , we can define the set of strategies induced by these correlations as

$$\mathfrak{S}_{\mathcal{C}} = \left\{ \mathcal{A}_{P,f}(p_{\mathbf{Y}|\mathbf{X}}) \mid p_{\mathbf{Y}|\mathbf{X}} = \prod_{i=1}^N p_{Y_i|X_i}, P \in \mathcal{C}, f \in \text{PP} \right\}. \quad (14)$$

where PP is the set of all local post-processings of the answers generated by the correlation (as defined in Eq. (10)). We now show that there is a close relation between the correlations \mathcal{C} shared by the senders and the strategies $\mathfrak{S}_{\mathcal{C}}$ induced by these correlations.

If \mathcal{C}_{cl} is the set of classical correlations, then any $P \in \mathcal{C}_{\text{cl}}$ can be written as $P(\mathbf{y}'|\mathbf{xy}) = \prod_{i=1}^N P_i(y'_i|x_i, y_i)$ and convex combinations thereof, where P_i are some distributions over \mathcal{Y}_i given a question-answer pair from $(\mathcal{X}_i, \mathcal{Y}_i)$ for $i \in [N]$. Then, from Eq. (13), we find that the strategy $\mathcal{A}_{P,f}(p_{\mathbf{Y}|\mathbf{X}})$ induced by the correlation assistance channel is also a classical strategy (since the input $p_{\mathbf{Y}|\mathbf{X}}$ is a classical strategy). Therefore, the set of strategies induced by classical correlations (as defined in Eq. (14)) is the set of classical strategies \mathfrak{S}_{cl} .

On the other hand, if \mathcal{C}_{Q} is the set of quantum correlations, then any $P \in \mathcal{C}_{\text{Q}}$ can be written as

$$P(\mathbf{y}'|\mathbf{xy}) = \text{Tr} \left[\rho \left(E_{y'_1}^{(x_1, y_1)} \otimes \cdots \otimes E_{y'_N}^{(x_N, y_N)} \right) \right]$$

where ρ is the quantum state shared between the senders and $\{E_{y'}^{(x_i, y_i)} \mid y' \in \mathcal{Y}_i\}$ is the local measurement implemented by the i th player upon receiving the input (x_i, y_i) for each $i \in [N]$. Now, we define new local measurement operators

$$\bar{E}_{\bar{y}_i}^{(\bar{x}_i)} = \sum_{y_i, y'_i \in \mathcal{Y}_i} p_{Y_i|X_i}(y_i|\bar{x}_i) f_i(\bar{y}_i|\bar{x}_i, y_i y'_i) E_{y'_i}^{(\bar{x}_i, y_i)},$$

so that the strategy induced by the correlation P can be written using Eq. (13) as

$$\bar{p}_{\mathbf{Y}|\mathbf{X}}(\bar{\mathbf{x}}|\bar{\mathbf{y}}) = \text{Tr} \left[\rho \left(\bar{E}_{\bar{y}_1}^{(\bar{x}_1)} \otimes \cdots \otimes \bar{E}_{\bar{y}_N}^{(\bar{x}_N)} \right) \right].$$

This shows that the set of strategies induced by quantum correlations, as defined in Eq. (14), is the set of quantum strategies \mathfrak{S}_{Q} . Similarly, one can verify that the set of strategies induced by no-signalling correlations \mathcal{C}_{NS} is the set of no-signalling strategies \mathfrak{S}_{NS} .

We now elaborate on how one can use the correlation-assistance channel to boost the sum capacity of the nonlocal games MAC. Given the NG-MAC \mathcal{N}_G obtained from a nonlocal game G , some correlation P shared by the senders and local post-processings f , we define the correlation-assisted NG-MAC $\mathcal{N}_G \circ \mathcal{A}_{P,f}$. That is, the input question-answer pair is first passed through the correlation-assistance channel $\mathcal{A}_{P,f}$, which tries to improve the strategy for playing the game, and the modified question-answer pair is passed on to the NG-MAC \mathcal{N}_G . A schematic of this procedure for the case of two senders is shown in Fig. 1. If the local post-processing f discards information about the input questions as well as the answers generated by the correlation P , i.e., $f(\bar{\mathbf{y}}|\mathbf{xy}, \mathbf{y}') = \delta_{\bar{\mathbf{y}}, \mathbf{y}'}$, then $\mathcal{A}_{P,f}$ becomes the identity channel. Therefore, the correlation-assisted NG-MAC $\mathcal{N}_G \circ \mathcal{A}_{P,f}$ is at least as powerful as the NG-MAC \mathcal{N}_G if we allow the senders to perform any local post-processing.

Suppose that the senders share a set of correlations \mathcal{C} . The \mathcal{C} -assisted achievable rate region of the NG-MAC \mathcal{N}_G is defined as

$$\text{Cap}_{\mathcal{C}}^{(1)}(\mathcal{N}_G) = \bigcup_{\substack{P \in \mathcal{C}, \\ f \in \text{PP}}} \text{Cap}(\mathcal{N}_G \circ \mathcal{A}_{P,f}) \quad (15)$$

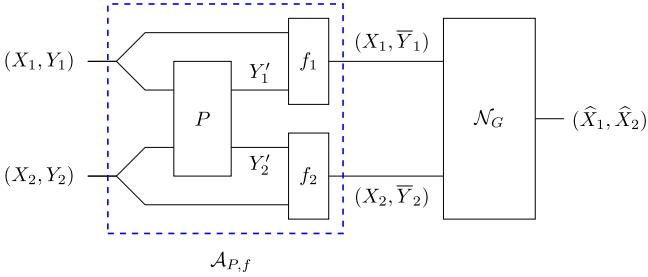


Fig. 1. Correlation-assisted NG-MAC $\mathcal{N}_G \circ \mathcal{A}_{P,f}$ for the case of two senders, obtained from the nonlocal games MAC \mathcal{N}_G defined in Eq. (9) and correlation-assistance channel $\mathcal{A}_{P,f}$ defined in Eq. (11).

where $\text{Cap}(\mathcal{N}_G \circ \mathcal{A}_{P,f})$ is the capacity region defined in Eq. (6) evaluated for the correlation-assisted NG-MAC $\mathcal{N}_G \circ \mathcal{A}_{P,f}$.³ The \mathcal{C} -assisted achievable sum rate of the NG-MAC \mathcal{N}_G is

$$S_{\mathcal{C}}(\mathcal{N}_G) = \sup \left\{ \sum_{i=1}^N R_i \mid (R_1, \dots, R_N) \in \text{Cap}_{\mathcal{C}}^{(1)}(\mathcal{N}_G) \right\}. \quad (16)$$

We now derive an alternate expression of $S_{\mathcal{C}}(\mathcal{N}_G)$ that is more convenient for computation. Prior to obtaining this expression, note that for any given family of sets $\{F_i\}_{i \in \mathcal{I}}$ indexed by some set \mathcal{I} and any function $g: \cup_{i \in \mathcal{I}} F_i \rightarrow \mathbb{R}$, we have

$$\sup_{r \in \cup_{i \in \mathcal{I}} F_i} g(r) = \sup_{i \in \mathcal{I}} \sup_{r \in F_i} g(r).$$

Using this equation along with Eq. (7), we can write

$$\begin{aligned} S_{\mathcal{C}}(\mathcal{N}_G) &= \sup_{P \in \mathcal{C}, f \in \text{PP}} S(\mathcal{N}_G \circ \mathcal{A}_{P,f}) \\ &= \sup_{P \in \mathcal{C}, f \in \text{PP}} \sup_{p^{(1)} \dots p^{(N)}} I(X_1, Y_1, \dots, X_N, Y_N; Z), \end{aligned} \quad (17)$$

where (X_i, Y_i) is the random variable (with distribution $p^{(i)}$) describing the input for $i \in [N]$, and Z is the random variable describing the output of the NG-MAC $\mathcal{N}_G \circ \mathcal{A}_{P,f}$.

Since $S_{\mathcal{C}}(\mathcal{N}_G)$ corresponds to a maximization over all possible local post-processings, we must have $S(\mathcal{N}_G) \leq S_{\mathcal{C}}(\mathcal{N}_G)$ for any set of correlations \mathcal{C} . Furthermore, if $\mathcal{C}_1 \subseteq \mathcal{C}_2$, then $\text{Cap}_{\mathcal{C}_1}^{(1)}(\mathcal{N}_G) \subseteq \text{Cap}_{\mathcal{C}_2}^{(1)}(\mathcal{N}_G)$, and consequently also $S_{\mathcal{C}_1}(\mathcal{N}_G) \leq S_{\mathcal{C}_2}(\mathcal{N}_G)$. Finally, note that we compute the relaxed sum capacity $C(\mathcal{N}_G)$ by maximizing over all possible distributions over the questions and all possible strategies $\mathfrak{S}_{\text{all}}$ that the players can use (since the maximization in Eq. (8) is over all input probability distributions). Because $\mathfrak{S}_{\mathcal{C}} \subseteq \mathfrak{S}_{\text{all}}$ for any set of correlations \mathcal{C} , we have $S_{\mathcal{C}}(\mathcal{N}_G) \leq C(\mathcal{N}_G)$. Therefore, we obtain a hierarchy,

$$S(\mathcal{N}_G) \leq S_{\text{cl}}(\mathcal{N}_G) \leq S_{\text{Q}}(\mathcal{N}_G) \leq S_{\text{NS}}(\mathcal{N}_G) \leq C(\mathcal{N}_G), \quad (18)$$

where ‘‘cl’’, ‘‘Q’’, and ‘‘NS’’ denote classical, quantum, and no-signalling correlations, respectively. Note that the sum capacity $S(\mathcal{N}_G)$ might not be equal to $S_{\text{cl}}(\mathcal{N}_G)$ because

³The superscript (1) signifies that $\text{Cap}_{\mathcal{C}}^{(1)}(\cdot)$ is merely a region consisting of achievable rate pairs, and hence contained in the full capacity region $\text{Cap}_{\mathcal{C}}(\cdot)$. To determine whether $\text{Cap}_{\mathcal{C}}^{(1)}(\cdot) = \text{Cap}_{\mathcal{C}}(\cdot)$ is outside the scope of this work.

classical correlations can be convex combinations of product distributions.

We now proceed to obtaining a bound on the \mathcal{C} -assisted achievable sum rate.

2) *Bounding the Correlation-Assisted Sum Rate*: Let \mathcal{C} be any set of correlations shared between the senders. In order to bound $S_{\mathcal{C}}(\mathcal{N}_G)$, we first obtain an optimization problem in terms of distributions over questions and strategies induced by the shared correlations. For a given correlation $P \in \mathcal{C}$, the input and output of $\mathcal{N}_G \circ \mathcal{A}_{P,f}$ can be described as follows. The channel $\mathcal{A}_{P,f}$ takes the input random variable $(X_1, Y_1, \dots, X_N, Y_N)$ and outputs $(\bar{X}_1, \bar{Y}_1, \dots, \bar{X}_N, \bar{Y}_N)$. The output of $\mathcal{A}_{P,f}$ becomes the input to \mathcal{N}_G that returns Z , i.e.,

$(X_1, Y_1, \dots, X_N, Y_N) \xrightarrow{\mathcal{A}_{P,f}} (\bar{X}_1, \bar{Y}_1, \dots, \bar{X}_N, \bar{Y}_N) \xrightarrow{\mathcal{N}_G} Z$ forms a Markov chain. From the data processing inequality [3], we obtain

$$I(X_1, Y_1, \dots, X_N, Y_N; Z) \leq I(\bar{X}_1, \bar{Y}_1, \dots, \bar{X}_N, \bar{Y}_N; Z). \quad (19)$$

Then, using Eq. (17) and Eq. (19), we get

$$S_{\mathcal{C}}(\mathcal{N}_G) \leq \sup_{\bar{p}} I(\bar{X}_1, \bar{Y}_1, \dots, \bar{X}_N, \bar{Y}_N; Z) \quad (20)$$

where the probability distribution \bar{p} defined in Eq. (12) is obtained by varying product distributions $p(x, y) = \prod_{i=1}^N p^{(i)}(x_i, y_i)$ input to $\mathcal{A}_{P,f}$, the correlation $P \in \mathcal{C}$, and the post-processing $f \in \text{PP}$.

We can reinterpret the above equation as a maximum over distributions over questions and strategies for playing the game G induced by the correlations \mathcal{C} . First, we write $p^{(i)} = \pi^{(i)} p_{Y_i|X_i}$, where $\pi^{(i)}$ is a distribution over the questions \mathcal{X}_i and $p_{Y_i|X_i}$ is a strategy chosen by i th player for $i \in [N]$. Therefore, the input probability distribution in Eq. (17) can be written as $p^{(1)} \dots p^{(N)} = \pi p_{Y|\mathbf{X}}$, where $\pi = \pi^{(1)} \dots \pi^{(N)}$ is a distribution over questions \mathcal{X} , and $p_{Y|\mathbf{X}} = \prod_{i=1}^N p_{Y_i|X_i}$ is a classical strategy chosen by the players. In particular, the input strategy $p_{Y|\mathbf{X}}$ is always a classical strategy. As noted in Eq. (13), the channel $\mathcal{A}_{P,f}$ takes this input strategy and returns a new strategy $\mathcal{A}_{P,f}(p_{Y|\mathbf{X}})$ that incorporates assistance from the shared correlation P . Since the senders have access to the set of correlations \mathcal{C} , we can write Eq. (20) as

$$S_{\mathcal{C}}(\mathcal{N}_G) \leq \sup_{\pi^{(1)} \dots \pi^{(N)}} \sup_{\bar{p}_{Y|\mathbf{X}} \in \mathfrak{S}_{\mathcal{C}}} I(\bar{X}_1, \bar{Y}_1, \dots, \bar{X}_N, \bar{Y}_N; Z), \quad (21)$$

where $\mathfrak{S}_{\mathcal{C}}$ is the set of strategies induced by \mathcal{C} as defined in Eq. (14). To obtain an upper bound, we will perform relaxations of the RHS of the above equation, and solve the resulting optimization problems.

We begin by writing the RHS of Eq. (21) in a form that is amenable for calculations. To that end, note that given an input probability distribution $\bar{p}(\bar{x}\bar{y}) = \pi(\bar{x})\bar{p}_{Y|\mathbf{X}}(\bar{y}|\bar{x})$, the probability distribution corresponding to the output of the channel \mathcal{N}_G is given by

$$p(\mathbf{z}) = \sum_{\bar{x}\bar{y} \in \mathcal{W}} \mathcal{N}_G(\mathbf{z}|\mathbf{x}\mathbf{y})\bar{p}(\bar{x}\bar{y}) + \sum_{\bar{x}\bar{y} \notin \mathcal{W}} \mathcal{N}_G(\mathbf{z}|\mathbf{x}\mathbf{y})\bar{p}(\bar{x}\bar{y})$$

$$= \pi(\mathbf{z}) \sum_{\mathbf{y}: \mathbf{z}, \mathbf{y} \in \mathcal{W}} \bar{p}_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{z}) + \frac{1}{d} p_L \quad (22)$$

where $d = |\mathcal{Z}|$ denotes the number of question pairs, while $p_L = \sum_{\overline{\mathbf{x}}\overline{\mathbf{y}} \notin \mathcal{W}} \bar{p}(\overline{\mathbf{x}}\overline{\mathbf{y}})$ denotes the probability of losing the game when questions are drawn as per the probability distribution π .

Note that $\mathcal{Z} = \mathcal{X}_1 \times \dots \times \mathcal{X}_N = \mathcal{X}$ is the set of question tuples output by the NG-MAC \mathcal{N}_G . Since \mathcal{Z} is a finite set of size d , we can fix a labelling for the elements of \mathcal{Z} and write $\mathcal{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_d\}$. Each \mathbf{z}_i corresponds to a particular question tuple. Then, we can define the contribution of a given strategy $\bar{p}_{\mathbf{Y}|\mathbf{X}}$ towards winning the game for each question tuple. We use \mathcal{Z} and \mathcal{X} interchangeably in the following discussion.

Definition 1 (Winning Vector): Given a strategy $\bar{p}_{\mathbf{Y}|\mathbf{X}}$ for playing the game $G = (\mathcal{X}, \mathcal{Y}, \mathcal{W})$, we let

$$w_i = \sum_{\mathbf{y}: \mathbf{z}_i, \mathbf{y} \in \mathcal{W}} \bar{p}_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{z}_i) \quad \text{for } i \in [d] \quad (23)$$

denote the contribution of the strategy towards winning the game G for question \mathbf{z}_i . We call the vector $\mathbf{w} = (w_1, \dots, w_d)$ the *winning vector* corresponding to the strategy $\bar{p}_{\mathbf{Y}|\mathbf{X}}$. Let \mathfrak{W}_C denote the set of winning vectors allowed by the strategies \mathfrak{S}_C ,

$$\mathfrak{W}_C = \left\{ \mathbf{w} \in [0, 1]^d \mid w_i = \sum_{\mathbf{y}: \mathbf{z}_i, \mathbf{y} \in \mathcal{W}} \bar{p}_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{z}_i), \text{ for } \bar{p}_{\mathbf{Y}|\mathbf{X}} \in \mathfrak{S}_C \text{ and } i \in [d] \right\}. \quad (24)$$

Observe that $w_i \in [0, 1]$ for all $i \in [d]$, so that \mathbf{w} is an element of the unit hypercube in \mathbb{R}^d . Note that we have $w_i = 1$ for a fixed strategy $\bar{p}_{\mathbf{Y}|\mathbf{X}}$ if and only if the players always win the game G when asked the question \mathbf{z}_i using the strategy $\bar{p}_{\mathbf{Y}|\mathbf{X}}$. On the other extreme, $w_i = 0$ if and only if the players always lose the game G when asked the question \mathbf{z}_i using the strategy $\bar{p}_{\mathbf{Y}|\mathbf{X}}$. Generally, questions are drawn with probability π over \mathcal{X} . The probability of winning the game for question i is $\pi_i w_i$, where $\pi_i = \pi(\mathbf{z}_i)$ is the probability of drawing the question tuple \mathbf{z}_i . The total probability of winning the game is $p_W = \sum_{i=1}^d \pi_i w_i$ and the probability of losing the game is

$$p_L = 1 - \sum_{i=1}^d \pi_i w_i = 1 - \langle \pi, \mathbf{w} \rangle. \quad (25)$$

Defining the matrix \overline{W} with components

$$\overline{W}_{ij} = w_i \delta_{ij} + \frac{1 - w_j}{d}, \quad (26)$$

one may write the output probability $p(\mathbf{z})$ in Eq. (22) as

$$p = \overline{W}\pi. \quad (27)$$

The mutual information $I(\overline{X}_1, \overline{Y}_1, \dots, \overline{X}_N, \overline{Y}_N; Z)$ can be written as

$$\begin{aligned} & I(\overline{X}_1, \overline{Y}_1, \dots, \overline{X}_N, \overline{Y}_N; Z) \\ &= H(Z) - H(Z|\overline{X}_1, \overline{Y}_1, \dots, \overline{X}_N, \overline{Y}_N) \\ &= H(Z) - p_L \ln(d), \end{aligned} \quad (28)$$

where we used the fact that $H(Z|\overline{\mathbf{x}}\overline{\mathbf{y}}) = 0$ when $\overline{\mathbf{x}}\overline{\mathbf{y}} \in \mathcal{W}$ whereas $H(Z|\overline{\mathbf{x}}\overline{\mathbf{y}}) = \ln(d)$ when $\overline{\mathbf{x}}\overline{\mathbf{y}} \notin \mathcal{W}$. Note that the

formula (28) was first derived in Ref. [11] for nonlocal games MAC with two senders. Using Eq. (27) and Eq. (25), we obtain

$$\begin{aligned} \mathcal{I}_w(\pi) &:= I(\overline{X}_1, \overline{Y}_1, \dots, \overline{X}_N, \overline{Y}_N; Z) \\ &= H(\overline{W}\pi) + \langle \pi, \mathbf{w} \rangle \ln(d) - \ln(d) \end{aligned} \quad (29)$$

where the notation, $\mathcal{I}_w(\pi)$, for the mutual information emphasizes that it is only a function of the distribution π over questions and the winning vector \mathbf{w} .

The RHS in Eq. (21) can be written as

$$\begin{aligned} \sup_{\pi^{(1)} \dots \pi^{(N)}} \sup_{\bar{p}_{\mathbf{Y}|\mathbf{X}}} I(\overline{X}_1, \overline{Y}_1, \dots, \overline{X}_N, \overline{Y}_N; Z) \\ &= \sup_{\pi^{(1)} \dots \pi^{(N)}} \sup_{\mathbf{w} \in \mathfrak{W}_C} \mathcal{I}_w(\pi^{(1)} \dots \pi^{(N)}) \\ &\leq \sup_{\pi \in \Delta_d} \sup_{\mathbf{w} \in \mathfrak{W}_C} \mathcal{I}_w(\pi), \end{aligned} \quad (30)$$

where \mathfrak{W}_C is the set of winning vectors defined in Eq. (24). To obtain Eq. (30), we relax the product distribution constraint $\pi^{(1)} \dots \pi^{(N)}$ over the questions to obtain a maximization over all distribution $\pi \in \Delta_d$ over the questions, where Δ_d denotes the $(d - 1)$ -dimensional standard simplex. This relaxation differs from that of Eq. (8) used in obtaining $C(\mathcal{N}_G)$ in that we only relax the distribution over the questions, but not the whole probability distribution.

For a fixed \mathbf{w} , the function $\mathcal{I}_w(\pi)$ is continuous in π over the compact set Δ_d . Thus, the maximization in Eq. (30) can be written as

$$\sup_{\mathbf{w} \in \mathfrak{W}_C} \max_{\pi \in \Delta_d} \mathcal{I}_w(\pi). \quad (31)$$

The inner optimization in Eq. (31) is a convex problem since $\mathcal{I}_w(\pi)$ is concave in π and Δ_d is a convex set. However, $\mathcal{I}_w(\pi)$ is not jointly concave in π and \mathbf{w} , and moreover, \mathfrak{W}_C need not be a convex set. Therefore, the optimization in Eq. (31) is generally nonconvex.

Our goal is to obtain an upper bound on the optimization in Eq. (31). To give a general idea of our approach to obtaining this bound, we list the main steps we will carry out.

Step 1: For a fixed \mathbf{w} , we obtain an upper bound on the inner optimization in Eq. (31). This bound is tight when either $\mathbf{w} \in \{0, 1\}^d$ or $\mathbf{w} > 0$ component-wise.

Step 2: We relax the set of allowed \mathbf{w} values to bound the outer optimization in Eq. (31) from above.

This procedure will result in the upper bound noted in Eq. (39). We explain the steps in detail in the following subsections.

Step 1: Bounding the inner optimization over question distributions: We obtain an upper bound on $\max_{\pi \in \Delta_d} \mathcal{I}_w(\pi)$ by considering two cases. First, we perform this optimization exactly for the case when $\mathbf{w} \in \{0, 1\}^d$. Next, for any $\mathbf{w} \in \mathfrak{W}_C$, we find an upper bound on $\max_{\pi \in \Delta_d} \mathcal{I}_w(\pi)$ using the result of case 1. The upper bound obtained in case 2 is tight when $\mathbf{w} > 0$.

Case 1: optimizing $\max_{\pi \in \Delta_d} \mathcal{I}_w(\pi)$ for fixed $\mathbf{w} \in \{0, 1\}^d$ Winning vectors $\mathbf{w} \in \{0, 1\}^d$ arise from strategies that either always win or always lose the game for any given question. Deterministic strategies, for example, give rise to such winning vectors. Recall that a classical deterministic

strategy corresponds to functions $f_i : \mathcal{X}_i \rightarrow \mathcal{Y}_i$, $i \in [N]$, chosen by the players. Such functions give rise to the classical strategy $p_{Y_i|X_i}(y|x) = \delta_{y,f_i(x)}$ that is 1 at $y = f_i(x)$ and zero elsewhere. It follows from Eq. (23) that $w_i \in \{0, 1\}$ for all $i \in \{1, \dots, d\}$, where \mathbf{w} is the winning vector defined by such a deterministic strategy.

The following proposition gives the result of the optimization $\max_{\pi \in \Delta_d} \mathcal{I}_{\mathbf{w}}(\pi)$ when $\mathbf{w} \in \{0, 1\}^d$. Note that \mathcal{W}_C need not contain such winning vectors. Computing $\max_{\pi \in \Delta_d} \mathcal{I}_{\mathbf{w}}(\pi)$ for $\mathbf{w} \in \{0, 1\}^d$ is a means to providing a bound for Eq. (31).

Proposition 2: Let G be a nonlocal game, and let \mathcal{N}_G be the MAC obtained from this nonlocal game. Let \mathbf{w} denote a winning vector as defined in Eq. (23), such that $w_i \in \{0, 1\}$ for all $i \in [d]$. Let $\mathcal{I}_{\mathbf{w}}$ denote mutual information between input and outputs of \mathcal{N}_G , as defined in Eq. (29). Let $\mathcal{K} = \{i \in [d] \mid w_i = 1\}$ denote the set of questions for which the strategy gives a correct answer. Denoting $K = |\mathcal{K}|$ and Δ_d to be the $(d-1)$ -dimensional standard simplex, we have

$$\max_{\pi \in \Delta_d} \mathcal{I}_{\mathbf{w}}(\pi) = \begin{cases} 0 & K = 0 \\ \mathcal{I}_K^* & 0 < K < d \\ \ln d & K = d. \end{cases} \quad (32)$$

The quantity \mathcal{I}_K^* is given by the expression

$$\mathcal{I}_K^* = \ln \left(K + (d - K) d^{-\frac{d}{d-K}} \right). \quad (33)$$

Proof: See Appendix A. \square

Observe that the maximum only depends on the total number d of questions, as well as the number K of questions that can be answered correctly using the deterministic strategy.

Case 2: Bounding $\max_{\pi \in \Delta_d} \mathcal{I}_{\mathbf{w}}(\pi)$ for fixed $\mathbf{w} \in \mathcal{W}_C$
When we work with arbitrary winning vectors, it is more challenging to maximize the mutual information over distributions on the questions. To make this maximization easier, we first show that the maximum mutual information \mathcal{I}_{d-1}^* corresponding to a winning vector that can answer exactly $d-1$ questions correctly will always be larger than the maximum mutual information $\max_{\pi \in \Delta_d} \mathcal{I}_{\mathbf{w}}(\pi)$ for any winning vector \mathbf{w} that answers no more than $d-1$ questions correctly. We therefore turn our attention to \mathbf{w} that doesn't necessarily do worse than this case, and obtain an expression for the maximum mutual information in terms of such \mathbf{w} .

Proposition 3: Let G be a nonlocal game, and let \mathcal{N}_G be the MAC obtained from this nonlocal game. Suppose that the senders of \mathcal{N}_G share a set of correlations \mathcal{C} . Let $\mathbf{w} \in \mathcal{W}_C$ be any winning vector allowed by the correlations \mathcal{C} as defined in Eq. (24). Let $\mathcal{K} = \{i \in [d] \mid w_i \neq 0\}$ be the set of questions with non-zero probability of winning the game using this strategy, and denote $K = |\mathcal{K}|$. Then, the following statements hold.

- 1) Suppose that $\max_{\pi \in \Delta_d} \mathcal{I}_{\mathbf{w}}(\pi)$ is achieved at π^* . Denote $\mathcal{K}^* = \{i \in [d] \mid w_i \pi_i^* \neq 0\}$ and $K^* = |\mathcal{K}^*|$ (we have $\mathcal{K}^* \subseteq \mathcal{K}$). Then, if $K^* < d$, we have $\max_{\pi \in \Delta_d} \mathcal{I}_{\mathbf{w}}(\pi) \leq \mathcal{I}_{d-1}^*$, where \mathcal{I}_{d-1}^* is given by Eq. (33).
- 2) As a consequence of the above result, we restrict our attention to strategies with $K^* = K = d$. In that case,

we have

$$\begin{aligned} \max_{\pi \in \Delta_d} \mathcal{I}_{\mathbf{w}}(\pi) &\equiv \mathcal{I}^*(\mathbf{w}) \\ &= \ln \left(\sum_{j=1}^d \exp \left[d w_{\text{eff}} \ln d \left(1 - \frac{1}{w_j} \right) \right] \right) \end{aligned} \quad (34)$$

where

$$w_{\text{eff}} = \left(\sum_{i=1}^d \frac{1}{w_i} \right)^{-1}. \quad (35)$$

Proof: See Appendix A. \square

Owing to the above result, we only need to focus on maximizing $\mathcal{I}^*(\mathbf{w})$ for those $\mathbf{w} \in \mathcal{W}_C$ with $w_i > 0$ for all $i \in [d]$. This is done in the next step.

Step 2: Bounding the outer optimization over winning vectors: As noted in the previous step, our goal is to maximize $\mathcal{I}^*(\mathbf{w})$ with respect to the feasible winning vectors $\mathbf{w} \in \mathcal{W}_C$ with $w_i > 0$ for all $i \in [d]$. The set of (feasible) winning vectors \mathcal{W}_C was defined in Eq. (24). Note that \mathcal{W}_C depends on the winning condition \mathcal{W} of the game as well as the set of correlations \mathcal{C} shared by the senders. Since we make no assumptions about the game or the set of correlations, it is difficult to optimize over \mathcal{W}_C . For this reason, we obtain a relaxation of the set \mathcal{W}_C , over which we optimize $\mathcal{I}^*(\mathbf{w})$. We will do this in two steps: (1) relate \mathbf{w} to the winning probability when the questions are drawn uniformly, and (2) use the maximum winning probability $\omega^{\mathcal{S}_C}(G)$ of the game (assumed to be known) corresponding to the strategies \mathcal{S}_C when the questions are drawn uniformly in order to obtain a convex set containing \mathcal{W}_C .

(1) From the definition of winning vector given in Eq. (23), we know that

$$w_i = \sum_{\mathbf{y}, \mathbf{z}_i, \mathbf{y} \in \mathcal{W}} \bar{p}_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{z}_i).$$

Recall that the winning probability of the game can be written as $p_W = \sum_{i=1}^d \pi_i w_i$ when the questions are drawn as per probability $\pi \in \Delta_d$. If the questions are drawn uniformly, then $\pi_U(z) = 1/d$ for all questions $\mathbf{z}_1, \dots, \mathbf{z}_d$. Therefore, $p_W = \sum_{i=1}^d w_i/d$ is the winning probability determined by the winning vector \mathbf{w} when the questions are drawn uniformly.

(2) We now look for a convex relaxation $\overline{\mathcal{W}_C}$ of \mathcal{W}_C . We want to make $\overline{\mathcal{W}_C}$ fairly independent of the winning set, except for dependence on $\omega^{\mathcal{S}_C}(G)$ and the number of question tuples d in the game.

Since $\omega^{\mathcal{S}_C}(G)$ is the maximum winning probability using the set of strategies \mathcal{S}_C under consideration, we must have

$$\frac{1}{d} \sum_{i=1}^d w_i \leq \omega^{\mathcal{S}_C}(G) \quad (36)$$

where $\mathbf{w} \in \mathcal{W}_C$. Now we make the relaxation that we allow *any* winning vector that satisfies Eq. (36). Consequently, we define

$$\overline{\mathcal{W}_C} = \left\{ \mathbf{w} \in [0, 1]^d \mid \frac{1}{d} \sum_{i=1}^d w_i \leq \omega^{\mathcal{S}_C}(G) \right\}. \quad (37)$$

Since any $\mathbf{w} \in \mathfrak{W}_C$ will satisfy Eq. (36), we have $\mathfrak{W}_C \subseteq \overline{\mathfrak{W}_C}$, confirming that $\overline{\mathfrak{W}_C}$ is a relaxation of \mathfrak{W}_C . Such a relaxation may allow for strategies not described by \mathfrak{S}_C . Note that $\overline{\mathfrak{W}_C}$ is a compact and convex set, and it depends only on the maximum winning probability and the number of questions in the game. Using this relaxation, we compute an upper bound on $\mathcal{I}^*(\mathbf{w})$ maximized over $\mathbf{w} \in \overline{\mathfrak{W}_C}$ satisfying $\mathbf{w} > 0$ componentwise.

Proposition 4: Let G be a nonlocal game and let \mathcal{N}_G be the NG-MAC constructed from G , as defined in Eq. (9). Suppose that the senders share the set of correlations \mathcal{C} , and let \mathfrak{W}_C be the corresponding set of winning vectors as defined in Eq. (24). Let $\overline{\mathfrak{W}_C}$ be the convex relaxation of \mathfrak{W}_C defined in Eq. (37) that depends only on the number of question tuples d in the game and the maximum winning probability $\omega^{\mathfrak{S}_C}(G)$ when the questions are drawn uniformly and answers given using strategies in \mathfrak{S}_C . Let $\mathcal{I}^*(\mathbf{w})$ be the function defined in Eq. (34). Then the maximum of $\mathcal{I}^*(\mathbf{w})$ over winning vectors $\mathbf{w} > 0$ in $\overline{\mathfrak{W}_C}$ is bounded from above by

$$\sup_{\mathbf{w} \in \overline{\mathfrak{W}_C}, \mathbf{w} > 0} \mathcal{I}^*(\mathbf{w}) \leq \ln \left(d - 1 + d^{-(1-\omega^{\mathfrak{S}_C}(G))d} \right) \quad (38)$$

Proof: See Appendix A. \square

Bound on the correlation-assisted achievable sum rate: We put all the above steps together to obtain a bound on $S_C(\mathcal{N}_G)$.

Theorem 5: Let G be an N -player promise-free nonlocal game with d question tuples, and let \mathcal{N}_G be the MAC obtained from G as defined in Eq. (9). Suppose that the senders share a set of correlations \mathcal{C} . Let \mathfrak{S}_C be the set of strategies induced by the correlations as defined in Eq. (14). Let $\omega^{\mathfrak{S}_C}(G)$ denote the maximum winning probability of the game when the questions are drawn uniformly and answers are obtained using strategies in \mathfrak{S}_C . Let $S_C(\mathcal{N}_G)$ denote the \mathcal{C} -assisted achievable sum rate of the NG-MAC \mathcal{N}_G as defined in Eq. (17). Then, we have

$$S_C(\mathcal{N}_G) \leq \ln \left(d - 1 + d^{-(1-\omega^{\mathfrak{S}_C}(G))d} \right) \quad (39)$$

with entropy measured in nats.

Proof: To obtain an upper bound on $S_C(\mathcal{N}_G)$, we start with Eq. (21). The RHS of Eq. (21) can be bounded by performing the maximization $\sup_{\mathbf{w} \in \mathfrak{W}_C} \max_{\pi \in \Delta_d} \mathcal{I}_{\mathbf{w}}(\pi)$, where $\mathcal{I}_{\mathbf{w}}(\pi)$ is the mutual information defined in Eq. (29). The set Δ_d denote the $(d-1)$ -dimensional standard simple, while \mathfrak{W}_C denotes the set of winning vectors induced by the correlations \mathcal{C} as defined in Eq. (24).

In Prop. 3, we show that if $\mathbf{w} \in \mathfrak{W}_C$ has one or more zero entries, then $\max_{\pi \in \Delta_d} \mathcal{I}_{\mathbf{w}}(\pi) \leq \mathcal{I}_{d-1}^*$, where \mathcal{I}_{d-1}^* is given by Eq. (33). Therefore, we only maximize $\mathcal{I}^*(\mathbf{w}) = \max_{\pi \in \Delta_d} \mathcal{I}_{\mathbf{w}}(\pi)$ over winning vectors $\mathbf{w} \in \mathfrak{W}_C$ satisfying $\mathbf{w} > 0$. The expression for $\mathcal{I}^*(\mathbf{w})$ in this case is given by Eq. (34). We relax the set \mathfrak{W}_C to the compact and convex set $\overline{\mathfrak{W}_C}$ defined in Eq. (37). Then, we give an upper bound on $\sup_{\mathbf{w} \in \overline{\mathfrak{W}_C}, \mathbf{w} > 0} \mathcal{I}^*(\mathbf{w})$ in Eq. (38).

By preceding remarks, we have

$$S_C(\mathcal{N}_G) \leq \max \left\{ \mathcal{I}_{d-1}^*, \ln \left(d - 1 + d^{-(1-\omega^{\mathfrak{S}_C}(G))d} \right) \right\}, \quad (40)$$

while from Eq. (33) we have

$$\mathcal{I}_{d-1}^* = \ln \left(d - 1 + d^{-d} \right).$$

Using $\ln \left(d - 1 + d^{-(1-\omega^{\mathfrak{S}_C}(G))d} \right) \geq \mathcal{I}_{d-1}^*$ in Eq. (40), we obtain Eq. (39). \square

Corollary 6: Let G be an N -player promise-free nonlocal game with d question tuples, and let \mathcal{N}_G be the MAC obtained from G as defined in Eq. (9). Let $\omega^{\text{cl}}(G)$ denote the maximum winning probability of the game when the questions are drawn uniformly and answers are obtained using classical strategies. Let $S(\mathcal{N}_G)$ denote the sum capacity of the NG-MAC \mathcal{N}_G . Then, we have

$$S(\mathcal{N}_G) \leq \ln \left(d - 1 + d^{-(1-\omega^{\text{cl}}(G))d} \right) \quad (41)$$

with entropy measured in nats.

Proof: From Eq. (18), we know that $S(\mathcal{N}_G) \leq S_{\text{cl}}(\mathcal{N}_G)$. Then, using Thm. 5, we obtain Eq. (41). \square

Note that the bounds on $S_C(\mathcal{N}_G)$ and $S(\mathcal{N}_G)$ given by Eq. (39) and Eq. (41), respectively, lie between $\ln(d-1)$ and $\ln(d)$. For sufficiently large d , when $\omega^{\mathfrak{S}_C}(G)$ is not close to 1, the sum capacity is bounded above by $\approx \ln(d-1)$. On the other hand, for $\omega^{\mathfrak{S}_C}(G) = 1$, we obtain an upper bound of $\ln(d)$, which can be achieved by $\mathbf{w} = (1, \dots, 1)^T$ as seen from Prop. 2. Using this, we can obtain separations between the correlation-assisted achievable sum rate corresponding to two different sets of correlations.

3) Separation Between Sum Rates With Assistance From Different Sets of Correlations: If \mathcal{C}_1 and \mathcal{C}_2 are two sets of correlations such that $\omega^{\mathfrak{S}_{\mathcal{C}_1}}(G) < 1$ while $\omega^{\mathfrak{S}_{\mathcal{C}_2}}(G) = 1$, then $S_{\mathcal{C}_1}(\mathcal{N}_G) < S_{\mathcal{C}_2}(\mathcal{N}_G) = \ln(d)$. We use this idea to provide separations of correlation-assisted achievable sum rate using classical, quantum and no-signalling correlations.

Separating $S_Q(\mathcal{N}_G)$ from $S(\mathcal{N}_G)$ for two-sender MACs: Consider the Magic Square Game, G_{MS} , used previously in [19], [20], [21], and [15] to obtain a separation between $S(\mathcal{N}_G)$ and $S_Q(\mathcal{N}_G)$. In this game, the referee selects a row $r \in \{1, 2, 3\}$ and column $c \in \{1, 2, 3\}$ from a 3×3 grid uniformly at random. The row is handed over to Alice while the column is given to Bob. Without communicating with each other, Alice & Bob need to fill bits in the given row and column such that the total parity of bits in the row is even, total parity of bits in the column is odd, and the bit at the intersection of the given row and column match.

There are $d = 9$ possible question pairs corresponding to the indices (r, c) . Classically, Alice & Bob can win the game at least 8 out of 9 times by implementing for example the following strategy:

1	0	1
1	1	0
1	0	?

where the entry in each box indicates bits filled by Alice and Bob. It can be shown that this strategy is optimal, therefore $\omega^{\text{cl}}(g_{\text{ms}}) = 8/9$ [15].

On the other hand, if Alice & Bob are allowed to use a quantum strategy, then they can share two copies of a

TABLE I
MEASUREMENT TABLE

$\sigma_x \otimes \mathbb{I}$	$\sigma_x \otimes \sigma_x$	$\mathbb{I} \otimes \sigma_x$
$-\sigma_x \otimes \sigma_z$	$\sigma_y \otimes \sigma_y$	$-\sigma_z \otimes \sigma_x$
$\mathbb{I} \otimes \sigma_z$	$\sigma_z \otimes \sigma_z$	$\sigma_z \otimes \mathbb{I}$

maximally entangled Bell state,

$$\rho_{\text{Bell}} = \frac{1}{2} \begin{pmatrix} 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 \end{pmatrix},$$

and submit answers to the referee based on a set of measurements given in Table I. Alice and Bob answer 0 if their measurement yields an eigenvector with eigenvalue 1, else they answer 1. Answers obtained this way can be shown to always satisfy the winning condition of the magic square game, i.e., $\omega^Q(G_{\text{MS}}) = 1$ [19], [20].

As the MAC obtained from the Magic Square Game has received attention in a previous study, we summarize the separation between sum capacity and entanglement-assisted sum rate given by our method in the following corollary.

Corollary 7: Let $\mathcal{N}_{G_{\text{MS}}}$ denote the MAC obtained from the magic square game. Then, the sum capacity of this MAC is bounded above as

$$S(\mathcal{N}_{G_{\text{MS}}}) \leq 3.02 \text{ bits.}$$

On the other hand, using assistance from quantum correlations, we obtain $S_Q(\mathcal{N}_{G_{\text{MS}}}) = 3.17$ bits. This gives a separation of at least 0.15 bits between sum rate with and without entanglement assistance.

Proof: Since $\omega^{\text{cl}}(G_{\text{MS}}) = 8/9$, we may use Cor. (6) to obtain a bound $S(\mathcal{N}_{G_{\text{MS}}}) \leq 3.02$ bits. At the same time, a perfect quantum strategy is available, i.e., $\omega^Q(G_{\text{MS}}) = 1$, and thus $S_Q(\mathcal{N}_{G_{\text{MS}}}) = 3.17$ bits [11]. \square

Our bound of 3.02 bits on the sum capacity of the NG-MAC $\mathcal{N}_{G_{\text{MS}}}$ is tighter than the previously reported bound of 3.14 bits [11]. Thus, our bound shows that entanglement assistance increases the sum rate by at least 4.7%, in comparison with the previously known result of 0.9%.

Since every quantum strategy is also a no-signalling strategy, we automatically obtain a separation between $S(\mathcal{N}_{G_{\text{MS}}})$ and $S_{\text{NS}}(\mathcal{N}_{G_{\text{MS}}})$. However, $S_{\text{NS}}(\mathcal{N}_{G_{\text{MS}}}) = S_Q(\mathcal{N}_{G_{\text{MS}}})$. In the following section, we use a game different from G_{MS} to obtain a separation between the quantum and no-signalling assisted sum rates.

Separating $S_{\text{NS}}(\mathcal{N}_G)$ from $S(\mathcal{N}_G)$ and $S_Q(\mathcal{N}_G)$ for two-sender MACs: In order to obtain a separation between the quantum-assisted sum rate and the no-signalling assisted sum rate, we consider the Clauser-Horne-Shimony-Holt (CHSH) game G_{CHSH} [22], [23]. In this game, a referee selects bits $x_1, x_2 \in \{0, 1\}$ uniformly at random, and gives them to Alice and Bob, respectively. Upon receiving these question bits, Alice answers with the bit $y_1 \in \{0, 1\}$ and Bob with $y_2 \in \{0, 1\}$. Alice and Bob chose their answers without

communicating with each other. They win the game if

$$x_1 \wedge x_2 = y_1 \oplus y_2,$$

where \wedge and \oplus represent logical AND and bitwise addition modulo 2. This game has a total of $d = 4$ question pairs.

It is known that the best classical strategy can answer only 3 out of the 4 question pairs correctly, i.e., $\omega^{\text{cl}}(G_{\text{CHSH}}) = 3/4$ [23]. The optimal quantum strategy achieves a winning probability of $\omega^Q(G_{\text{CHSH}}) = (1 + 1/\sqrt{2})/2 \approx 85.4\%$ [23]. While there is no classical or quantum strategy that can always win the game, one can construct a no-signalling distribution,

$$P_{\text{PR}}(y_1, y_2 | x_1, x_2) = \frac{1}{2} \delta_{x_1 \wedge x_2, y_1 \oplus y_2},$$

usually called the Popescu-Rohrlich (PR) box [24], which represents a perfect strategy for winning the CHSH game. Therefore $\omega^{\text{NS}}(G_{\text{CHSH}}) = 1$.

Using $\omega^{\text{cl}}(G_{\text{CHSH}}) = 3/4$ in Cor. (6) gives $S(\mathcal{N}_{G_{\text{CHSH}}}) \leq 1.7$ in the classical case. On the other hand, using $\omega^Q(G_{\text{CHSH}}) = (1 + 1/\sqrt{2})/2$ in Thm. 5, gives an upper bound, $S_Q(\mathcal{N}_{G_{\text{CHSH}}}) \leq 1.78$ bits, in the quantum case. In the case of no-signalling, a perfect strategy is possible and thus we have $S_{\text{NS}}(\mathcal{N}_{G_{\text{CHSH}}}) = 2$ bits. In this way, we obtain a separation between the quantum and no-signalling assisted achievable sum rate.

Separating $S_Q(\mathcal{N}_G)$ from $S(\mathcal{N}_G)$ for N -sender MACs: We now consider a game G_{MPP} that we call the multiparty parity game, which was first introduced by Brassard et al. [25]. In this game, N players are each handed a bit and they each answer by returning a bit. The players have a *promise*: the total number of ones in the N -bit string handed to them is even. If this even number is divisible by 4, then the winning condition is that the total bit string returned by the players have an even number of ones. Otherwise, the winning condition is to return a string with an odd number of ones.

Formally, we have $\mathcal{X}_i = \{0, 1\}$ and $\mathcal{Y}_i = \{0, 1\}$ for $i \in [N]$. As before, we denote $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_N$ as the set of questions and $\mathcal{Y} = \mathcal{Y}_1 \times \dots \times \mathcal{Y}_N$ as the set of answers for the N players. The promise,

$$\mathcal{P} = \left\{ \mathbf{x} \in \mathcal{X} \mid \sum_i x_i = 0 \pmod{2} \right\}$$

is a subset of \mathcal{X} from which the questions are drawn. The winning condition for the game can be described by the set

$$\mathcal{W}^{\mathcal{P}} = \left\{ (\mathbf{x}, \mathbf{y}) \in \mathcal{P} \times \mathcal{Y} \mid \sum_i y_i - \frac{1}{2} \sum_i x_i = 0 \pmod{2} \right\}.$$

Brassard et al. [25] demonstrated that classical strategies can win this game with a probability of at most $\omega^{\text{cl}, \mathcal{P}}(G_{\text{MPP}}) = 1/2 + 2^{-\lceil N/2 \rceil}$ when the questions are drawn uniformly from the promise set. In contrast, a perfect quantum strategy is possible [25].

Now we consider the following promise-free version of this game. Herein, the question and answer set remain the same, but the winning condition is defined as the set $\mathcal{W} = \mathcal{W}^{\mathcal{P}} \cup (\mathcal{P}^c \times \mathcal{Y})$. That is, the players win automatically if a question from outside the promise set is presented to them. To apply the bound obtained in Thm. 5, we need to compute the maximum classical winning probability $\omega^{\text{cl}}(G_{\text{MPP}})$ for

the promise-free case when questions are drawn uniformly. To that end, we show how to compute the maximum winning probability $\omega^{\mathfrak{S}}(G)$ when we convert any game with a promise \mathcal{P} to a promise-free game, assuming that the questions are drawn uniformly and answers are given using strategies in \mathfrak{S} . Since $\mathcal{W}^{\mathcal{P}} \cap (\mathcal{P}^c \times \mathcal{Y}) = \emptyset$, we have

$$\begin{aligned} p_W &= \sum_{(x,y) \in \mathcal{W}^{\mathcal{P}}} \frac{1}{|\mathcal{X}|} p(y|x) + \sum_{x \in \mathcal{P}^c} p(x) \\ &= \frac{|\mathcal{P}|}{|\mathcal{X}|} p_W^{\mathcal{P}} + \left(1 - \frac{|\mathcal{P}|}{|\mathcal{X}|}\right). \end{aligned}$$

Since the set of strategies \mathfrak{S} chosen by the players have a maximum probability $\omega^{\mathfrak{S}, \mathcal{P}}(G)$ of winning when the questions are drawn uniformly from \mathcal{P} , we can infer that

$$\omega^{\mathfrak{S}}(G) = \frac{|\mathcal{P}|}{|\mathcal{X}|} \omega^{\mathfrak{S}, \mathcal{P}}(G) + \left(1 - \frac{|\mathcal{P}|}{|\mathcal{X}|}\right).$$

For the multiparty parity game, since half the N -bit strings are in \mathcal{P} and the other half are in \mathcal{P}^c , we get

$$\omega^{\text{cl}}(G_{\text{MPP}}) = \frac{3}{4} + 2^{-\lceil N/2 \rceil + 1}. \quad (42)$$

Then, by Cor. (6), we find that the sum capacity for the MAC obtained from the multiparty parity game is bounded above as $S(\mathcal{N}_{G_{\text{MPP}}}) \leq \log(d-1 + 2^{-(1-\omega^{\text{cl}}(G))d})$ bits, where $d = 2^N$ and $\omega^{\text{cl}}(G)$ is given by Eq. (42). In particular, $S(\mathcal{N}_{G_{\text{MPP}}}) < \log(d)$. In contrast, since a perfect quantum strategy is available, we have $S_Q(\mathcal{N}_{G_{\text{MPP}}}) = \log(d)$, thus giving a separation between the sum capacity and the quantum-assisted sum rate for N -sender MACs. For example, when we have $N = 3$ senders, we obtain $S(\mathcal{N}_{G_{\text{MPP}}}) \leq 2.84$ bits. In contrast, $S_Q(\mathcal{N}_{G_{\text{MPP}}}) = 3$ bits in the quantum case.

B. Looseness of Convex Relaxation of the Sum Capacity

In the previous section, we looked at separations between the sum rates with assistance from classical, quantum and no-signalling strategies. In this section, we construct a game such that one can obtain an arbitrarily large separation between the sum capacity and the relaxed sum capacity. Recall that the relaxed sum capacity corresponds to dropping the product distribution constraint in the maximization problem:

$$C(\mathcal{N}_G) = \max_{\bar{p}(\bar{x}\bar{y})} I(\bar{X}_1, \bar{Y}_1, \dots, \bar{X}_N, \bar{Y}_N; Z),$$

where $\bar{X}_1, \bar{Y}_1, \dots, \bar{X}_N, \bar{Y}_N$ are random variables describing the input to the NG-MAC \mathcal{N}_G , while Z is the random variable describing the output. We maximize over all possible input probability distributions \bar{p} , so that the resulting quantity is the capacity of \mathcal{N}_G when we think of it as a single-input single-output channel. As noted in Sec. III-A.1, we have $S(\mathcal{N}_G) \leq S_C(\mathcal{N}_G) \leq C(\mathcal{N}_G)$ for any set of correlations \mathcal{C} . Indeed, since we maximize over all probability distributions over the input to \mathcal{N}_G , we can write the relaxed sum capacity as

$$C(\mathcal{N}_G) = \max_{\bar{p}_{\mathbf{Y}|\mathbf{X}} \in \mathfrak{S}_{\text{all}}} \max_{\bar{\pi} \in \Delta_d} I(\bar{X}_1, \bar{Y}_1, \dots, \bar{X}_N, \bar{Y}_N; Z)$$

where $\mathfrak{S}_{\text{all}}$ denotes the set of all possible strategies that the players can use to play the game. In particular, this amounts to allowing the players to communicate after the questions have been handed over to them.

To analyze $C(\mathcal{N}_G)$, we study some properties of $\mathfrak{S}_{\text{all}}$. It can be verified that $\mathfrak{S}_{\text{all}}$ is a convex set. The extreme points of this set correspond to deterministic strategies $f: \mathcal{X} \rightarrow \mathcal{Y}$ that allow for communication between the players (see Prop. 13 in App. B). This implies that the maximum winning probability of the game $\omega^{\mathfrak{S}_{\text{all}}}(G)$, when the questions are drawn uniformly and answers are obtained using the strategies in $\mathfrak{S}_{\text{all}}$, is always achieved by a deterministic strategy of the form mentioned above. We now give an explicit description of a deterministic strategy (not necessarily unique) achieving the maximum winning probability.

The best deterministic strategy $f^{(D)}: \mathcal{X} \rightarrow \mathcal{Y}$ can be written as

$$f^{(D)}(\mathbf{x}) = \begin{cases} \mathbf{y} & \exists \mathbf{y} \in \mathcal{Y} \text{ such that } (\mathbf{x}, \mathbf{y}) \in \mathcal{W} \\ \mathbf{y}_0 & \text{otherwise,} \end{cases} \quad (43)$$

where $\mathbf{y}_0 \in \mathcal{Y}$ is an arbitrary element chosen beforehand. We note that for each $\mathbf{x} \in \mathcal{X}$, some element $\mathbf{y} \in \mathcal{Y}$ satisfying $(\mathbf{x}, \mathbf{y}) \in \mathcal{W}$ is chosen *a priori* (if it exists), so that the function is well-defined, though not necessarily unique. In other words, $f^{(D)}$ gives the correct answer if a correct answer for the given question exists, and if not, it gives an arbitrary answer that is necessarily incorrect. It can, therefore, be inferred that the maximum winning probability can be written as

$$\omega^{\mathfrak{S}_{\text{all}}}(G) = \frac{|\{\mathbf{x} \in \mathcal{X} \mid \exists \mathbf{y} \in \mathcal{Y} \text{ such that } (\mathbf{x}, \mathbf{y}) \in \mathcal{W}\}|}{d} \quad (44)$$

This is the best that one can do given any nonlocal game G . Note also that $\omega^{\mathfrak{S}_{\text{all}}}(G)$ can be directly computed from the winning condition \mathcal{W} .

The above observation directly leads to upper and lower bounds on $C(\mathcal{N}_G)$. The upper bound is obtained by using $\omega^{\mathfrak{S}_{\text{all}}}(G)$ from Eq. (44) in Thm. 5. Here, we implicitly use the fact that our upper bound is valid even when the questions are drawn arbitrarily. Let $\mathbf{w}^{(D)} \in \{0, 1\}^d$ be the winning vector corresponding to the best deterministic strategy $f^{(D)}$ given in Eq. (43). Note that we maximize over all distributions over the questions when computing $C(\mathcal{N}_G)$. Therefore, using $\mathbf{w}^{(D)}$ in Prop. 2, we obtain a lower bound on $C(\mathcal{N}_G)$. In particular, if there is at least one correct answer for every question, then $f^{(D)}$ is a perfect strategy and $C(\mathcal{N}_G) = \ln(d)$.

Now, we obtain a separation between $C(\mathcal{N}_G)$ and $S(\mathcal{N}_G)$. To that end, we construct a game called the signalling game.

1) *Signalling Game G_s and separation of $C(\mathcal{N}_{G_s})$ from $S(\mathcal{N}_{G_s})$ and $S_{\text{NS}}(\mathcal{N}_{G_s})$:* Consider a game where Alice & Bob are each given a question from some set of questions. They win the game if they can correctly guess the question handed over to the other person. Since the game can be won if Alice & Bob “signal” their question to each other, we call this the signalling game G_s .

Formally, we consider question sets $\mathcal{X}_1, \mathcal{X}_2$ and answer sets $\mathcal{Y}_1 = \mathcal{X}_2$ and $\mathcal{Y}_2 = \mathcal{X}_1$, and the winning condition is defined as

$$\begin{aligned} \mathcal{W} = \{(x_1, x_2, y_1, y_2) \in (\mathcal{X}_1 \times \mathcal{X}_2) \times (\mathcal{Y}_1 \times \mathcal{Y}_2) \mid \\ y_1 = x_2, y_2 = x_1\}. \end{aligned}$$

Note that there is exactly one correct answer for each question pair $(x_1, x_2) \in \mathcal{X}_1 \times \mathcal{X}_2$.

Since we bound the sum capacity of the NG-MAC \mathcal{N}_{G_s} obtained from the signaling game G_s using the maximum winning probability, we analyze the winning strategies for this game. To that end, consider some set of strategies \mathfrak{S} that Alice and Bob use to play the game. For simplicity, we assume that this is a compact set (thinking of strategies as vectors as in Prop. 13), which holds, for example, when \mathfrak{S} is the set of no-signalling strategies (see Prop. 14). Suppose that $\omega^{\mathfrak{S}}(G) = 1$ corresponding to this set of strategies, and let $p_{\mathbf{Y}|\mathbf{X}}^*$ be the strategy that achieves this maximum winning probability. When the questions are drawn uniformly, the winning probability for strategy $p_{\mathbf{Y}|\mathbf{X}}^*$ is given by $p_W^* = (\sum_i w_i^*)/d$, where $\mathbf{w}^* = (w_1^*, \dots, w_d^*)$ is the winning vector corresponding to $p_{\mathbf{Y}|\mathbf{X}}^*$, as defined in Eq. (23). Since $p_W^* = \omega^{\mathfrak{S}}(G) = 1$ by assumption and $\mathbf{w}^* \in [0, 1]^d$, we must have $w_i^* = 1$ for all $i \in [d]$.

For convenience, denote $\mathcal{X} = \mathcal{X}_1 \times \mathcal{X}_2$ as the set of questions. Using Def. 23, we can write the winning vector \mathbf{w}^* as

$$w_{\mathbf{x}}^* = \sum_{\mathbf{y}: \mathbf{x}\mathbf{y} \in \mathcal{W}} p_{\mathbf{Y}|\mathbf{X}}^*(\mathbf{y}|\mathbf{x})$$

where we label the components of \mathbf{w}^* using the questions $\mathbf{x} \in \mathcal{X}$. Since $w_{\mathbf{x}}^* = 1$ for all $\mathbf{x} \in \mathcal{X}$ and because there is exactly one correct answer for each question \mathbf{x} , we can infer that $p_{\mathbf{Y}|\mathbf{X}}^*$ must be a deterministic strategy. Written explicitly, we have

$$p_{\mathbf{Y}|\mathbf{X}}^*(y_1, y_2|x_1, x_2) = \delta_{y_1, x_2} \delta_{y_2, x_1}.$$

Note that $p_{\mathbf{Y}|\mathbf{X}}^*$ cannot be a no-signalling strategy. Indeed,

$$p_{\mathbf{Y}|\mathbf{X}}^*(y_1|x_1, x_2) = \sum_{y_2 \in \mathcal{Y}_2} p_{\mathbf{Y}|\mathbf{X}}^*(y_1, y_2|x_1, x_2) = \delta_{y_1, x_2}. \quad (45)$$

This cannot satisfy the no-signalling condition $p_{\mathbf{Y}|\mathbf{X}}^*(y_1|x_1, x_2) = p_{\mathbf{Y}|\mathbf{X}}^*(y_1|x_1)$ given in Eq. (4) because the RHS in Eq. (45) depends on x_2 . In other words, we cannot have $w_{\mathbf{x}} = 1$ for *any* question $\mathbf{x} \in \mathcal{X}$ using a no-signalling strategy. In particular, the perfect strategy is not no-signalling, and therefore, $\omega^{\text{NS}}(G_s) < 1$ for no-signalling strategies. Subsequently, we also have $\omega^Q(G_s) < 1$ and $\omega^{\text{cl}}(G_s) < 1$, because the set of classical and quantum strategies are contained in the set of no-signalling strategies. It then follows from Thm. 5 and Cor. 6 that each of $S(\mathcal{N}_{G_s}), S_Q(\mathcal{N}_{G_s}), S_{\text{NS}}(\mathcal{N}_{G_s})$ is strictly less than $\ln(d)$. On the other hand, since a perfect strategy is possible allowing communication between Alice & Bob, we have $C(\mathcal{N}_{G_s}) = \ln(d)$. Therefore, we have obtained a separation between $C(\mathcal{N}_{G_s})$ and $S(\mathcal{N}_{G_s}), S_Q(\mathcal{N}_{G_s}), S_{\text{NS}}(\mathcal{N}_{G_s})$.

Below, we argue that this separation becomes arbitrarily large as the number of questions increases. To that end, we compute $\omega^{\text{cl}}(G_s)$. Since the maximum winning probability obtained using classical strategies when questions are drawn uniformly is achieved by a deterministic strategy, it is sufficient to restrict our attention to classical deterministic strategies chosen by Alice & Bob. Recall that a classical deterministic

strategy corresponds to two functions $f_1: \mathcal{X}_1 \rightarrow \mathcal{Y}_1$ and $f_2: \mathcal{X}_2 \rightarrow \mathcal{Y}_2$ chosen by Alice & Bob, respectively. This translates to the probability distribution $p_{\mathbf{Y}|\mathbf{X}}^{(D)}(y_1, y_2|x_1, x_2) = \delta_{y_1, f_1(x_1)} \delta_{y_2, f_2(x_2)}$. Then, the winning probability using a classical deterministic strategy when the questions are drawn uniformly is given by

$$p_W^{(D)} = \frac{1}{d} \sum_{x_1 \in \mathcal{X}_1, x_2 \in \mathcal{X}_2} \delta_{x_2, f_1(x_1)} \delta_{x_1, f_2(x_2)} \quad (46)$$

where we used the fact that the signalling game has only one correct answer (x_2, x_1) corresponding to each question (x_1, x_2) . It can be seen from Eq. (46) that, for achieving maximum winning probability, the function f_1 must be able to invert the action of the function f_2 or vice-versa.

If $|\mathcal{X}_2| \leq |\mathcal{X}_1|$, then the set $f_2(\mathcal{X}_2)$ can cover at most $|\mathcal{X}_2|$ elements of \mathcal{X}_1 . Subsequently, $\delta_{x_1, f_2(x_2)} \neq 0$ for at most $|\mathcal{X}_2|$ elements of \mathcal{X}_1 . We can then infer from Eq. (46) that

$$p_W^{(D)} \leq \frac{|\mathcal{X}_2|}{d} = \frac{1}{|\mathcal{X}_1|}$$

since $d = |\mathcal{X}_1||\mathcal{X}_2|$. Using a similar reasoning when $|\mathcal{X}_1| \leq |\mathcal{X}_2|$,

$$\omega^{\text{cl}}(G_s) = \frac{1}{\max(|\mathcal{X}_1|, |\mathcal{X}_2|)}.$$

In particular, we have $d \rightarrow \infty$ if either of $|\mathcal{X}_1|$ or $|\mathcal{X}_2|$ diverges, while the winning probability $\omega^{\text{cl}}(G) \rightarrow 0$. Subsequently, $C(\mathcal{N}_{G_s}) \rightarrow \infty$ but $S(\mathcal{N}_{G_s}) \rightarrow 0$. Therefore, we get an arbitrarily large separation between $C(\mathcal{N}_{G_s})$ and $S(\mathcal{N}_{G_s})$.

In fact, we verify through numerical simulations that the situation is equally bad for the no-signalling assisted sum rate. To that end, we compute the maximum winning probability $\omega^{\text{NS}}(G_s)$ numerically. We show in Prop. 14 that the set of no-signalling strategies for N -player games is a compact and convex set (specifically, a convex polytope). Therefore, computing $\omega^{\text{NS}}(G_s)$ amounts to solving a linear program (this fact is well-known for 2-player games [26]). For $2 \leq |\mathcal{X}_1|, |\mathcal{X}_2| \leq 10$, we verify that the numerically computed value for $\omega^{\text{NS}}(G_s)$ matches $\omega^{\text{cl}}(G_s) = 1/\max(|\mathcal{X}_1|, |\mathcal{X}_2|)$. Thus, we expect $\omega^{\text{cl}}(G_s) = \omega^Q(G_s) = \omega^{\text{NS}}(G_s)$ for the signalling game. This would imply that $S_Q(\mathcal{N}_{G_s}), S_{\text{NS}}(\mathcal{N}_{G_s}) \rightarrow 0$ as $d \rightarrow \infty$ but $C(\mathcal{N}_{G_s}) \rightarrow \infty$. In other words, even with quantum or no-signalling assistance, the sum rate is far lower than the bound given by the relaxed sum capacity.

This example highlights the importance of finding better methods to upper bound the sum capacity of MACs. In the next part, we take a step in this direction by defining and studying a class of global optimization problems with relevance in information theory. This class of optimization problems is motivated from the non-convex problem encountered in sum capacity computation. In the third part of this study, we will show how to apply these algorithms for computing the sum capacity of arbitrary two-sender MACs.

IV. OPTIMIZATION OF LIPSCHITZ-LIKE FUNCTIONS

A. Lipschitz-Like Functions

The main object of our study is a Lipschitz-like function. Such functions are a generalization of Lipschitz-continuous

and Hölder continuous functions. Recall that a function $f: \mathcal{D} \rightarrow \mathcal{E}$ between two subsets of normed vector spaces is said to be Lipschitz continuous if for all $x, x' \in \mathcal{D}$, we have $\|f(x) - f(x')\|_{\mathcal{E}} \leq L \|x - x'\|_{\mathcal{D}}$ for some constant $L > 0$. The function f is said to be Hölder continuous if $\|f(x) - f(x')\|_{\mathcal{E}} \leq L \|x - x'\|_{\mathcal{D}}^{\gamma}$ for some constants $L, \gamma > 0$ [27]. Hölder continuity is a more general notion than Lipschitz continuity since $\gamma = 1$ gives the definition of Lipschitz continuity. Even so, this definition is not general enough to capture the continuity properties of entropic quantities. Shannon and von Neumann entropies, for example, satisfy a different continuity bound. Specifically, if $p, q \in \Delta_d$ are discrete probability distributions, then the Shannon entropy satisfies

$$|H(p) - H(q)| \leq \frac{1}{2} \log(d-1) \|p - q\|_1 + h\left(\frac{1}{2} \|p - q\|_1\right), \quad (47)$$

where $h(x) = -x \log(x) - (1-x) \log(1-x)$ is the binary entropy function [28]. Similarly, von Neumann entropy satisfies the Fannes-Audenaert inequality [29]. To encapsulate such behaviour of entropic quantities, we define Lipschitz-like functions as follows.

Definition 8 (Lipschitz-Like Function): Let $\beta: \mathbb{R}_+ \rightarrow \mathbb{R}$ be a non-negative, continuous, monotonically increasing function such that $\beta(0) = 0$. Let \mathcal{D} and \mathcal{E} be subsets of a normed vector space. Then a function $f: \mathcal{D} \rightarrow \mathcal{E}$ is said to be Lipschitz-like or β -Lipschitz-like if it satisfies

$$\|f(x) - f(x')\|_{\mathcal{E}} \leq \beta(\|x - x'\|_{\mathcal{D}}) \quad \forall x, x' \in \mathcal{D}. \quad (48)$$

Some remarks about this definition are in order. The definition of Lipschitz-like functions can be generalized to metric spaces in a straightforward manner. The reason we require β to be monotonically increasing is because it makes Lipschitz-like functions behave similar to Lipschitz continuous functions in the sense that the bound on $\|f(x) - f(x')\|_{\mathcal{E}}$ tightens or loosens with the value of $\|x - x'\|_{\mathcal{D}}$. Moreover, this assumption helps with optimization of Lipschitz-like functions. Similarly, we require continuity of β for simplicity, but this can be relaxed to right-continuity at 0. Since β is (right-)continuous at 0, it follows from Eq. (48) that f is a continuous function. If $\beta(x) = Lx$ for some $L > 0$, then f is a Lipschitz continuous function with Lipschitz constant L , and if $\beta(x) = Lx^{\gamma}$ for $L, \gamma > 0$, then f is a Hölder continuous function with constant L . Lipschitz-like functions are therefore a generalization of Lipschitz and Hölder continuous functions.

With a slight modification of the right-hand side of Eq. (47), we can show that entropy is a Lipschitz-like function. To that end, we define the modified binary entropy as follows.

$$\bar{h}(x) = \begin{cases} -x \log(x) - (1-x) \log(1-x) & \text{if } x \leq \frac{1}{2} \\ \log(2) & \text{if } x \geq \frac{1}{2} \end{cases} \quad (49)$$

Observe that \bar{h} is a non-negative, continuous, and monotonically increasing function that satisfies $\bar{h}(0) = 0$, and furthermore, we have $h(x) \leq \bar{h}(x)$ for all $x \in [0, 1]$. Thus,

defining $\beta_H(x) = \log(d-1)x/2 + \bar{h}(x/2)$, we obtain

$$|H(p) - H(q)| \leq \beta_H(\|p - q\|_1).$$

From this, we can conclude that Shannon entropy (and similarly, von Neumann entropy) is a Lipschitz-like function.

Moreover, linear combinations and compositions of Lipschitz-like functions is again a Lipschitz-like function. We summarize this observation in the following result.

Proposition 9: 1) If $f_1: \mathcal{D} \rightarrow \mathcal{E}$ and $f_2: \mathcal{D} \rightarrow \mathcal{E}$ are β_1 -Lipschitz-like and β_2 -Lipschitz-like functions respectively, then the linear combination $\alpha_1 f_1 + \alpha_2 f_2$ is a $(|\alpha_1|\beta_1 + |\alpha_2|\beta_2)$ -Lipschitz-like function, where α_1, α_2 are scalars.

2) If $f_1: \mathcal{D} \rightarrow \mathcal{E}$ and $f_2: \mathcal{E} \rightarrow \mathcal{F}$ are β_1 -Lipschitz-like and β_2 -Lipschitz-like functions respectively, then the composition $f_2 \circ f_1$ is a $(\beta_2 \circ \beta_1)$ -Lipschitz-like function.

Proof: 1. The function $g = \alpha_1 f_1 + \alpha_2 f_2$ satisfies $\|g(x) - g(y)\|_{\mathcal{E}} \leq |\alpha_1| \|f_1(x) - f_1(y)\|_{\mathcal{E}} + |\alpha_2| \|f_2(x) - f_2(y)\|_{\mathcal{E}}$ for $x, y \in \mathcal{D}$ by triangle inequality. Then, using the fact that f_1, f_2 are Lipschitz-like, we obtain $\|g(x) - g(y)\|_{\mathcal{E}} \leq |\alpha_1|\beta_1(\|x - y\|_{\mathcal{D}}) + |\alpha_2|\beta_2(\|x - y\|_{\mathcal{D}})$. Since $\beta = |\alpha_1|\beta_1 + |\alpha_2|\beta_2$ is non-negative, continuous and monotonically increasing with $\beta(0) = 0$, we can conclude that g is β -Lipschitz-like.

2. The function $g = f_2 \circ f_1$ satisfies $\|g(x) - g(y)\|_{\mathcal{F}} \leq \beta_2(\|f_1(x) - f_1(y)\|_{\mathcal{E}})$ for $x, y \in \mathcal{D}$ by using the fact that f_2 is β_2 -Lipschitz-like. Then, since f_1 is β_1 -Lipschitz-like and β_2 is monotonically increasing, we obtain $\|g(x) - g(y)\|_{\mathcal{F}} \leq \beta_2(\beta_1(\|x - y\|_{\mathcal{D}}))$. Since $\beta = \beta_2 \circ \beta_1$ is non-negative, continuous and monotonically increasing with $\beta(0) = 0$, we can conclude that g is β -Lipschitz-like. \square

In light of the above result, we can conclude that several entropic quantities derived from Shannon entropy and von Neumann entropy are Lipschitz-like functions. This lends support to our claim that the techniques developed in this section can potentially be useful for non-convex optimization problems in information theory.

For the purposes of optimization, we take the co-domain to be the real line, i.e., $\mathcal{E} = \mathbb{R}$ with the usual norm. Motivated by applications in information theory, we will mainly be focusing on the case where the domain $\mathcal{D} = \Delta_d$ is the standard simplex in \mathbb{R}^d and $\|\cdot\|_{\mathcal{D}} = \|\cdot\|_1$. Nevertheless, our techniques work more generally with any norm (for example, using equivalence of norms in finite dimensions).

In this section, we will develop algorithms for optimizing Lipschitz-like functions over a closed interval and a standard simplex. We also discuss techniques for optimizing Lipschitz-like functions over arbitrary compact and convex domains in App. C-D. We approach this problem by showing how to extend a Lipschitz-like function from a compact and convex domain to all of Euclidean space such that the global optimum is not affected, which may be of independent interest. Some of these algorithms we discuss are generalizations of existing algorithms for optimization of Lipschitz continuous functions. For this reason, we present a brief overview of some existing algorithms and our approach to generalizing them. In the next section, we focus on outlining the high-level ideas without delving into the technical details. Subsequently,

we will give detailed results on the optimization algorithms, along with some numerical examples.

B. Overview of the Optimization Algorithms

All the algorithms we propose in this study are designed to optimize any β -Lipschitz-like function f . The function β is assumed to be known beforehand, but the function f is unknown and we can only query it at a specified point. The algorithms can then use the knowledge of the domain, the function β , the queried points and corresponding values of the objective function f to approximate the maximum of f to an additive precision $\epsilon > 0$. Such a setting is commonly used to study the performance of optimization algorithms [27].

In order to make concrete statements about the complexity of optimization, we will assume that β does not explicitly depend on the dimension. Our algorithms and convergence analysis are valid even if this assumption does not hold. By an efficient algorithm for optimization, we mean an algorithm that computes the optimum of f to a given additive precision $\epsilon > 0$ in time polynomial in the dimension and inverse precision $1/\epsilon$. We sometimes informally use the phrase “practically efficient” and variations thereof to mean that the algorithm runs reasonably fast in practice, e.g., to exclude situations where the scaling of runtime with dimension is too large (for example, $O(d^{10})$).

Before discussing technical details of the algorithms presented in this study, we give a high-level overview of the main ideas. We begin by presenting an algorithm that can optimize any β -Lipschitz-like function f when the domain $\mathcal{D} = [a, b]$ is a closed interval. Our algorithm is a generalization of the Piyavskii-Shubert algorithm [30], [31], and it focuses on constructing successively better upper-bounding functions by using the Lipschitz-like property of the objective function. At each iteration, the maximum of the upper-bounding function is computed, which is an easier problem because we only need to maximize the function β that is known to be monotonically increasing. The computed maximum of the upper-bounding function at each iteration generates a sequence of points that partitions the interval. When the distance between any two of these points becomes sufficiently small, one can show that the maximum of the upper-bounding function is close to the maximum of the original objective function. Our algorithm is guaranteed to converge to the optimal solution within a precision of $\epsilon > 0$ in $\lceil (b - a)/\delta \rceil$ time steps in the worst-case, where $\delta > 0$ is the largest number satisfying $\beta(\delta) \leq \epsilon/2$ (see Prop. 15). We refer to this algorithm as modified Piyavskii-Shubert algorithm. We remark that several extensions of the Piyavskii-Shubert algorithm have been presented in the literature (see, for example, Ref. [32], [33], [34], [35]). It would be interesting to undertake a more detailed study comparing such algorithms with our method in the future.

In higher dimensions, we focus on the case where the objective f is a β -Lipschitz-like function over the standard simplex $\mathcal{D} = \Delta_d$. For this case, we present two algorithms for finding the global optimum of f . For the first algorithm, we resort to a straightforward grid search. Using the results of

Ref. [36], one can show that a grid of size $O(d^{\lceil 1/\delta^2 \rceil})$ suffices to converge to a specified precision of $\epsilon > 0$, where δ is the largest number satisfying $\beta(\delta) \leq \epsilon/2$ (see Prop. 17).⁴ These results stand to demonstrate that one can, in principle, compute the maximum of f in polynomial time for a fixed precision. Moreover, a simple observation about ordering the elements of the grid allows for efficient construction of the grid, along with the possibility of parallelizing the grid search. Despite the possibility of polynomial complexity (in dimension) and numerical improvements, grid search is still too inefficient to be of practical use except for very small dimensions.

Another strategy we propose is to construct a “dense” Lipschitz continuous curve that gets close to each point of \mathcal{D} to within some specified distance. Such a strategy was adopted by Ref. [37] for optimizing Lipschitz continuous functions over a hypercube, and is referred to as Alienor method in the literature. When \mathcal{D} is the standard simplex in d dimensions, we give a time and memory efficient algorithm to construct such a curve. This allows us to reduce the d -dimensional problem of optimizing f over \mathcal{D} to the one-dimensional problem of optimizing it over an interval using the generated curve. If $\alpha > 0$ is the largest number satisfying $\beta(\alpha) \leq \epsilon/2$, this method takes $O(\alpha^{1-d}/d)$ iterations in the worst case for large dimensions. While this is much worse than a grid search in large dimensions, in small dimensions this takes fewer iterations to converge than grid search when the tolerance ϵ is small. The dense curve algorithm also has the advantage that we can find an upper bound on the maximum by running the algorithm for a fixed number of time steps. Furthermore, similar to grid search, the dense curve algorithm can be parallelized. Nevertheless, we remark that both the grid search and dense curve algorithms are impractical for even moderately large dimensions. We detail these methods here in the hope that these engender the development of more practical algorithms for optimizing Lipschitz-like functions over the standard simplex in higher dimensions.

We also study the more general case when \mathcal{D} is a compact and convex set in \mathbb{R}^d . To handle optimization in this general case, we seek to reduce it to a case that can be solved using known techniques. To that end, we show how to extend the objective function f from the domain \mathcal{D} to the full Euclidean space \mathbb{R}^d while retaining the Lipschitz-like property when β is itself Lipschitz-like (see Def. 19 and Prop. 20). The extended function \bar{f} has the property that its optimum coincides with the optimum of f when optimized over any set containing \mathcal{D} . Since we have the freedom to choose which set to optimize \bar{f} over, there are different algorithms one can potentially use for this optimization. In this study, we find the maximum value of f by optimizing its extension \bar{f} over a hypercube containing \mathcal{D} . To perform this optimization, we resort to using dense curves because the convergence analysis is very similar to the case of the standard simplex $\mathcal{D} = \Delta_d$. The effective problem, as before, is one-dimensional and can be solved using the modified Piyavskii-Shubert algorithm. In general, this algorithm needs exponentially many iterations (with the

⁴Note the different scaling in δ in Prop. 17 compared to that of $O(d^{\lceil 1/\sqrt{\delta} \rceil})$ mentioned in [36], which we believe to be erroneous.

dimension) to find the optimal solution to within a specified precision of $\epsilon > 0$ (see Prop. 21 for a precise statement). This exponential scaling with the dimension stems from the fact that we do not use any structure of \mathcal{D} to construct the dense curve and instead rely on a curve generated for a hypercube. The $O(1/\epsilon^d)$ complexity cannot be improved in general without additional assumptions on \mathcal{D} or the class of functions we optimize (see Prop. 21). However, it might be possible to improve the other factors that scale exponentially in the algorithm as noted in App. C-D. An advantage of using this algorithm is that one can specify a fixed number of iterations to obtain an upper bound on the maximum. Because optimizing f over an arbitrary compact and convex domain is not directly relevant to our study of MACs, we relegate this discussion to App. C-D.

We now dive into details of the proposed algorithms.

C. Optimizing Lipschitz-Like Functions Over an Interval Using Modified Piyavskii-Shubert Algorithm

We begin our study by presenting an algorithm for computing the maximum of any β -Lipschitz-like function f over a closed interval $\mathcal{D} = [a, b]$. A pseudocode for this algorithm is given in Alg. 1.

Algorithm 1 Computing the Maximum of a Function f Satisfying Eq. (48) for $\mathcal{D} = [a, b]$, Given $\epsilon > 0$

```

1: function
  MAXIMIZE_LIPSCHITZ-LIKE_FUNCTION_1D( $\epsilon$ )
2:   Initialize  $q^{(0)} = a$ 
3:   Set  $F_0(q) = f(q^{(0)}) + \beta(|q - q^{(0)}|)$  for  $q \in [a, b]$ 
4:   Set  $F \leftarrow F_0$  and  $q^* \leftarrow b$ 
5:   Set  $q^{(1)} \leftarrow q^*$ ,  $k \leftarrow 1$ 
6:   while  $F(q^*) - f(q^*) > \epsilon$  do
7:     Sort  $\{q^{(0)}, \dots, q^{(k)}\}$  from smallest to largest
       and relabel the points in ascending order.
8:     Define  $F_i(q) = f(q^{(i)}) + \beta(|q - q^{(i)}|)$ 
       for  $0 \leq i \leq k$ 
9:     for  $i = 0, \dots, k - 1$  do
10:      Set  $g_i(q) = F_i(q) - F_{i+1}(q)$ 
11:      Find  $\bar{q}^{(i)} \in [q^{(i)}, q^{(i+1)}]$  such that
           $g_i(\bar{q}^{(i)}) = 0$  using any root finding method.
12:    end for
13:    Pick an index  $m \in \text{argmax}_{0 \leq i \leq k-1} F_i(\bar{q}^{(i)})$ 
       and set  $q^* = \bar{q}^{(m)}$ .
14:    Update  $F \leftarrow F_m$ 
15:    Set  $q^{(k+1)} \leftarrow q^*$ ,  $k \leftarrow k + 1$ 
16:  end while
17:  return  $f(q^*)$ 
18: end function

```

We refer to the function F_i defined in line 8 of Alg. 1 as a bounding function, since $f(q) \leq F_i(q)$ for all $q \in [a, b]$ and all i . Note that $F_i(q^{(i)}) = f(q^{(i)})$ for all i . The bounding function F_i depends non-trivially on the argument q only through the function $\beta(|q - q^{(i)}|)$. We compute the optimum of the function f by maximizing these bounding functions, which is an easier problem because β is continuous and monotonic.

Essentially, the algorithm does the following. Suppose that at the K th time step, we have the iterates $q^{(0)}, \dots, q^{(K)}$ sorted in the ascending order $a = q^{(0)} \leq q^{(1)} \leq \dots \leq q^{(K)} = b$. Then in each of the intervals $[q^{(k)}, q^{(k+1)}]$ for $k \in \{0, \dots, K-1\}$ we compute the point $\bar{q}^{(k)}$ attaining the maximum of the bounding function $\min\{F_k, F_{k+1}\}$ (see lines 10, 11 in Alg. 1). Our next iterate $q^{(K+1)}$ is chosen to be in $\text{argmax}\{\bar{q}^{(0)}, \dots, \bar{q}^{(K-1)}\}$. In the next iterate, since $F_{K+1}(q^{(K+1)}) = f(q^{(K+1)})$, we have essentially tightened the upper bound on the function f . Proceeding this way, one can verify that the algorithm eventually approximates the function f from above well enough. A schematic of this procedure is shown in Fig. 2.

Indeed, we show in Prop. 15 that Alg. 1 is guaranteed to converge to the global maximum within an error of ϵ . This convergence takes at most $\lceil (b-a)/\delta \rceil$ in the worst-case, where $\delta > 0$ is the largest number satisfying $\beta(\delta) \leq \epsilon/2$. Since we find successively better upper bounds on the objective, Alg. 1 can be modified so that it accepts a fixed number of iterations instead of a precision, and outputs the upper bound $F(q^*)$ on the maximum of f . This upper bound has an error of at most $F(q^*) - f(q^*)$.

When the domain $\mathcal{D} = \Delta_2$ is the standard simplex in 2-dimensions, we can parameterize any $x \in \Delta_2$ as $x = (q, 1-q)$ with $q \in [0, 1]$. Furthermore, we have $\|x - y\|_1 = 2|p - q|$ for $x, y \in \Delta_2$ parametrized as $x = (q, 1-q)$ and $y = (p, 1-p)$. Thus, taking $a = 0$, $b = 1$, and replacing $|q - q^{(i)}|$ with $2|q - q^{(i)}|$ in Alg. 1, we get an algorithm for optimizing f over Δ_2 .

Next, we present algorithms to optimize Lipschitz-like functions in higher dimensions. Our main focus will be on optimizing Lipschitz-like functions over the standard simplex Δ_d for $d \geq 3$.

D. Optimizing Lipschitz-Like Function Over the Standard Simplex

We present two algorithms to optimize Lipschitz-like functions over $\mathcal{D} = \Delta_d$. The first algorithm is a simple grid search, whereas the second algorithm uses dense curves to fill \mathcal{D} . We describe these algorithms, corresponding convergence guarantees and practical implementation in detail.

1) *Optimization Using Grid Search*: Our goal in this section is to optimize a β -Lipschitz-like function f over the standard simplex $\mathcal{D} = \Delta_d$. Based on the results of Ref. [36], we present a grid search method for finding the maximum of f over \mathcal{D} . We begin by defining the integer grid

$$\mathcal{I}_{d,N} = \left\{ n \in \mathbb{N}^d \mid \sum_{i=1}^d n_i = N \right\}, \quad (50)$$

where $d \in \mathbb{N}_+$ denotes the dimension and $N \in \mathbb{N}_+$ denotes the size of the integer grid. This grid has

$$N_{\text{grid}} = \binom{N+d-1}{d-1} \quad (51)$$

elements because each element of $\mathcal{I}_{d,N}$ can be obtained by arranging N ones into d coordinates. From

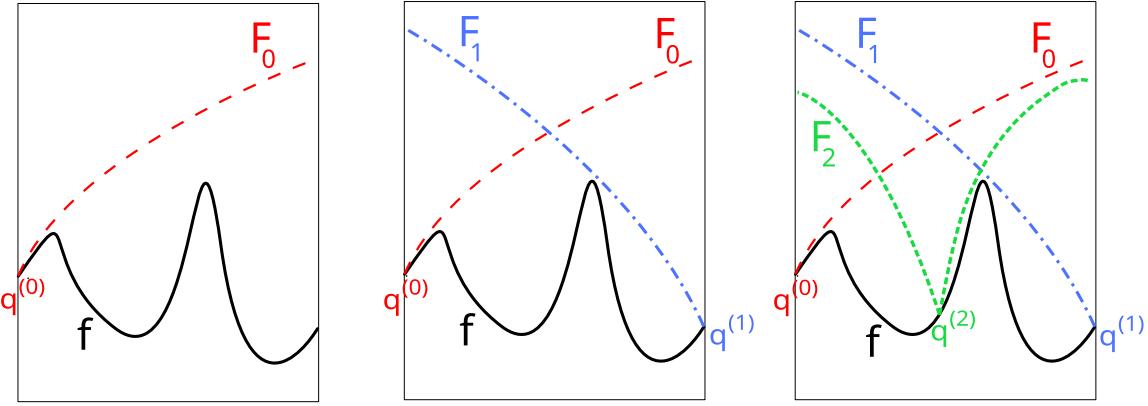


Fig. 2. A schematic of iterations of Alg. 1 for optimizing a Lipschitz-like function f .

$\mathcal{I}_{d,N}$, we obtain the grid

$$\Delta_{d,N} = \left\{ \frac{n}{N} \mid n \in \mathcal{I}_{d,N} \right\} \quad (52)$$

for the standard simplex Δ_d . Note that $\Delta_{d,N}$ is just a rescaled version of $\mathcal{I}_{d,N}$, i.e., $\Delta_{d,N} = \mathcal{I}_{d,N}/N$.

The authors of Ref. [36] propose to search the grid $\Delta_{d,N}$ in order to optimize Hölder continuous functions over the standard simplex. Their results are based on approximations of the function f using Bernstein polynomials, and are in fact general enough to handle the optimization of Lipschitz-like functions (see Prop. 17 for details). Such an approach to compute the maximum of f to a specified precision $\epsilon > 0$ is summarized in Alg. 2.

Algorithm 2 Computing the Maximum of a Function f Satisfying Eq. (48) for $\mathcal{D} = \Delta_d$, Given $\epsilon > 0$

```

1: function
  MAXIMIZE_LIPSCHITZ-LIKE_FUNC_SIMPLEX( $d, \beta, \epsilon$ )
2:   Set  $N = \lceil 1/\delta^2 \rceil$ , where  $\delta$  is the
      largest number satisfying  $\beta(\delta) \leq \epsilon/2$ 
3:   Construct the grid  $\Delta_{d,N}$ 
4:   return  $f^* = \max\{f(x) \mid x \in \Delta_{d,N}\}$ 
5: end function

```

In Prop. 17, we show that Alg. 2 computes the maximum of f to a precision $\epsilon > 0$ in $N_{\text{grid}} = \binom{N+d-1}{d-1}$ time steps, where $N = \lceil 1/\delta^2 \rceil$ and δ is the largest number satisfying $\beta(\delta) \leq \epsilon/2$. For fixed $\epsilon > 0$ and $d \gg N$, this amounts to $O(d^{1/\delta^2})$ iterations. In other words, we can find the optimum of f in polynomial time for a fixed precision. We note that if β achieves the value $\epsilon/2$, we can obtain δ by solving $\beta(\delta) = \epsilon/2$ using bisection (or any other root finding method) because β is continuous and monotonically increasing.

The crucial step in implementing the above algorithm is computing the grid $\Delta_{d,N}$ efficiently. For that purpose, note that we can write any element $x \in \Delta_{d,N}$ as $x = (N-n_{d-1}, n_{d-1}-n_{d-2}, \dots, n_2-n_1, n_1)/N$, where $0 \leq n_i \leq n_{i+1} \leq N$ for $i \in [d-2]$ (see Prop. 16 for a proof). Thus, the elements of $\Delta_{d,N}$ can be computed iteratively, with the total number of iterations equalling N_{grid} . This also allows for parallelizing

the search over the grid. The exact algorithm we use to query the elements of the grid orders the elements such that every consecutive element is equidistant with respect to l_1 -norm. This approach is explained in the next section and allows for easy parallelization.

2) *Optimization Using Dense Curves*: Next, we outline a method to optimize Lipschitz-like functions over $\mathcal{D} = \Delta_n$ by filling \mathcal{D} with an α -dense Lipschitz curve. We propose this method as a way to reduce the total number of iterations required for finding the optimum to a small additive error in practice compared to grid search. Such a strategy of using dense curves was outlined in Ref. [37] to optimize Lipschitz continuous functions over a hypercube.

Definition 10 (α -Dense Curve): Given a number $\alpha > 0$, numbers $a, b \in \mathbb{R}$, and a nonempty set $\mathcal{S} \subseteq \mathbb{R}^n$, a function $\gamma: [a, b] \rightarrow \mathcal{S}$ is said to be an α -dense curve in the norm $\|\cdot\|$ if for any $x \in \mathcal{S}$, we can find some $\theta \in [a, b]$ such that $\|\gamma(\theta) - x\| \leq \alpha$ [37].

The curve γ is said to be β_γ -Lipschitz-like if there is some non-negative, continuous, monotonically increasing function $\beta_\gamma: \mathbb{R}_+ \rightarrow \mathbb{R}$ with $\beta_\gamma(0) = 0$ such that for any $\theta, \theta' \in [a, b]$, we have $\|\gamma(\theta) - \gamma(\theta')\| \leq \beta_\gamma(|\theta - \theta'|)$.

The numbers $a, b \in \mathbb{R}$ are the end points of the interval over which the curve is defined. In this section, we will focus on constructing an α -dense curve for the standard simplex $\mathcal{S} = \Delta_d$ in l_1 -norm, where $\alpha = 2(d-1)/N$. Here, d is the dimension and N is a positive integer that controls the value of α . We achieve this by finding a way to efficiently connect the points of the grid $\mathcal{I}_{d,N}$ defined in Eq. (50) in the previous section. Towards this end, we define the following ordering of the grid $\mathcal{I}_{d,N}$.

Definition 11 (Equidistant Ordering of $\mathcal{I}_{d,N}$): Given $d, N \in \mathbb{N}_+$ with $d \geq 2$, let $\mathcal{I}_{d,N}$ be the grid defined in Eq. (52). For $d = 2$, define the forward ordering $\mathcal{I}_{2,N} = \{(N, 0), (N-1, 1), \dots, (1, N-1), (0, N)\}$ and the reverse ordering $\mathcal{I}_{2,N} = \{(0, N), (1, N-1), \dots, (N-1, 1), (N, 0)\}$. For $d \geq 3$, define the forward ordering inductively as follows. Start with $\mathcal{I}_{d,N} = \emptyset$. For each $n_{d-1} \in [N]$, forward order the elements of $\mathcal{I}_{d-1, n_{d-1}}$ if n_{d-1} is odd, and reverse order them if n_{d-1} is even. Append the elements $(N-n_{d-1}, n_{d-1}-n_{d-2}, \dots, n_2-n_1, n_1)$ with $(n_{d-1}-n_{d-2}, \dots, n_2-n_1, n_1) \in$

$\mathcal{I}_{d-1, n_{d-1}}$ as ordered above to the (ordered) set $\mathcal{I}_{d, N}$. Reverse ordering of $\mathcal{I}_{d, N}$ corresponds to writing the forward ordered set in the reverse order.

We remark that $d = 2$ is just a special case of the definition showing the basis step for induction. It can be seen that the first element of $\mathcal{I}_{d, N}$ is $(N, 0, \dots, 0)$ and the last element is $(0, \dots, 0, N)$ if $\mathcal{I}_{d, N}$ is forward ordered (and the opposite is true if the elements are reverse ordered). In Prop. 16, we show that $\mathcal{I}_{d, N}$ can be ordered in this manner and the distance between any two consecutive elements of $\mathcal{I}_{d, N}$ according to this ordering is 2 measured in l_1 -norm. By default, we work with forward ordering unless specified otherwise. As an example, the (forward) ordering of $\mathcal{I}_{3, 3}$ is given by

$$\mathcal{I}_{3, 3} = \{(3, 0, 0), (2, 1, 0), (2, 0, 1), (1, 0, 2), (1, 1, 1), (1, 2, 0), (0, 3, 0), (0, 2, 1), (0, 1, 2), (0, 0, 3)\}.$$

It can be seen that any two consecutive elements have an l_1 -norm distance of 2. Since the grid $\Delta_{d, N}$ defined in Eq. (52) is just a scaled version of $\mathcal{I}_{d, N}$, this ordering also applies to $\Delta_{d, N}$.

To generate a dense curve filling Δ_d , we make use of the grid $\Delta_{d, N}$ along with the ordering described above. However, $\Delta_{d, N}$ has $N_{\text{grid}} = \binom{N+d-1}{d-1}$ elements, and therefore, it can get very expensive to store this grid in memory even if N and d are only moderately large (for example, $N = d = 20$ gives $N_{\text{grid}} \approx 7 \times 10^{10}$). For this purpose, we develop an algorithm that can efficiently query the elements of ordered $\Delta_{d, N}$ without explicitly constructing the set. We skip the details of this algorithm in the study, but include an implementation of the algorithm on Github (see Sec. E-B). Owing to such an approach, the resulting construction is both time and memory efficient, i.e., $\gamma(\theta)$ can be computed efficiently for a curve γ . We describe the construction of the curve below.

Algorithm 3 Construct an α -Dense Curve That Fills the Simplex Δ_d for $\alpha = 2(d-1)/N$, $N \in \mathbb{N}_+$ and $\theta \in [0, L_{\text{curve}}]$, Where L_{curve} Is Given in Eq. (53)

```

1: function
  CONSTRUCT_DENSE_CURVE_STD_SIMPLEX( $d, N, \theta$ )
2:   Compute the index  $k = \lfloor \theta N/2 \rfloor$ 
3:   Compute  $t = 1 + k - \theta N/2$ 
4:   Obtain the grid points  $x_k, x_{k+1} \in \Delta_{d, N}$ , where
      the elements of  $\Delta_{d, N}$  are ordered as per Def. (11)
5:   return  $tx_k + (1-t)x_{k+1}$ 
6: end function

```

We now explain the reasoning behind Alg. 3. We construct the $(2(d-1)/N)$ -dense curve γ by joining the consecutive points of the ordered grid $\Delta_{d, N}$. Since every two consecutive points in the ordered grid $\Delta_{d, N}$ are $2/N$ distance apart in l_1 -norm, the total length of this curve measured in l_1 -norm is

$$L_{\text{curve}} = \frac{2}{N} N_{\text{grid}} = \frac{2}{N} \binom{N+d-1}{d-1}. \quad (53)$$

In order to obtain $\gamma(\theta)$ efficiently for a given $\theta \in [0, L_{\text{curve}}]$, we first compute the grid points x_k and x_{k+1} for which $\gamma(\theta) \in$

$[x_k, x_{k+1}]$. In particular, we have $\gamma(\theta) = tx_k + (1-t)x_{k+1}$ for some $t \in [0, 1]$. In order to compute t , we note that $\|\gamma(\theta) - x_k\|_1 = \theta - 2k/N$, where the RHS is the length of the curve at θ minus the length of the curve at grid point x_k (which is $2k/N$). Since $\|\gamma(\theta) - x_k\|_1 = (1-t) \|x_{k+1} - x_k\|_1 = 2(1-t)/N$, we obtain $t = 1 + k - \theta N/2$. Since $\theta N/2 - 1 \leq k \leq \theta N/2$, we have $0 \leq t \leq 1$. In Prop. 18, we show that the curve constructed using Alg. 3 is $(2(d-1)/N)$ -dense and satisfies the property

$$\|\gamma(\theta) - \gamma(\theta')\|_1 \leq \min\{|\theta - \theta'|, 2\}. \quad (54)$$

That is, the curve γ is β_γ -Lipschitz-like with $\beta_\gamma(x) = \min\{x, 2\}$. In particular, γ is Lipschitz continuous with Lipschitz constant 1.

The essential idea behind the optimization of Lipschitz-like functions using an α -dense curve is as follows. If the objective function $f: \mathcal{D} \rightarrow \mathbb{R}$ is Lipschitz-like and $\gamma: [a, b] \rightarrow \mathcal{D}$ is an α -dense Lipschitz-like curve, then the function $f \circ \gamma: [a, b] \rightarrow \mathbb{R}$ is also Lipschitz-like. Therefore, one can optimize the function $f \circ \gamma$ using Alg. 1. In order to converge to the optimum within a precision of $\epsilon > 0$, the constant α must be chosen appropriately. The exact procedure is outlined in Alg. 4.

Algorithm 4 Computing the Maximum of a β -Lipschitz-Like Function f Satisfying Eq. (48) for $\mathcal{D} = \Delta_d$, Given $\epsilon > 0$

```

1: function
  MAXIMIZE_LIPSCHITZ-LIKE_FUNC_SIMPLEX( $d, \beta, \epsilon$ )
2:   Compute the largest number  $\alpha > 0$  such that
     $\beta(\alpha) \leq \epsilon/2$ 
3:   Set  $N = \lceil 2(d-1)/\alpha \rceil$ 
4:   Construct the  $(2(d-1)/N)$ -dense curve
     $\gamma: [0, L_{\text{curve}}] \rightarrow \mathcal{D}$  as per Alg. 3
5:   Compute the maximum  $g^*$  of  $g = f \circ \gamma$  over  $[0, L_{\text{curve}}]$ 
    to a precision of  $\epsilon/2$  using Alg. 1
6:   return  $g^*$ 
7: end function

```

In Prop. 18, we show that Alg. 4 is guaranteed to converge to the maximum of f to within a precision of $\epsilon > 0$. This takes $\lceil 2(N+d-1)/d\alpha \rceil$ iterations in the worst case. For large dimensions and fixed precision, this number scales as $O(\alpha^{1-d}/d)$ with the dimension.

One can parallelize the above algorithm by breaking the interval $[0, L_{\text{curve}}]$ into finitely many sub-intervals and optimizing the function over each interval separately. The final maximum can be obtained by taking the maximum of the maximum values computed for each sub-interval. Since we only break the interval into finitely many sub-intervals, the guarantees given by Prop. 18 remain valid.

3) *Numerical Examples*: For verifying the performance of the grid search algorithm (Alg. 2) and the algorithm based on dense curves (Alg. 4), we present two numerical examples in Tab. II where the maximum over the simplex can be computed exactly. The functions considered in both these examples are Lipschitz continuous with respect to the l_1 -norm. It can be seen that both these methods compute the maximum to within the specified precision. For small dimensions and small values of

tolerance ϵ , we find that the algorithm based on dense curves requires far fewer iterations than grid search. As a result, the dense curve algorithm converges much faster than grid search for the examples presented in Tab. II. On the other hand, for larger values of tolerance and higher dimensions, grid search can converge faster since its complexity scales polynomially with the dimension whereas the complexity of dense curve algorithm scales exponentially in the worst case. Even so, the power of the polynomial in complexity of grid search can be large. Therefore, as the dimension d increases, both these algorithms become infeasible to implement in practice.

Details on how one can use α -dense curves to optimize a Lipschitz-like function over any compact and convex domain is given in App. C-D.

V. SUM CAPACITY COMPUTATION OF TWO-SENDER MACS

A. Sum Capacity Computation as Optimization of a Lipschitz-Like Function

Our focus in this part is computing the sum capacity of arbitrary 2-sender MACs. As noted in the preliminary Sec. II-C, the optimization involved in computing the sum capacity is non-convex. Unfortunately, the approach via convex relaxation commonly pursued is, in general, not a viable strategy for bounding the sum capacity of MACs. The family of two-sender MACs constructed in Sec. III-B provides a striking example of this claim, where we show that the difference between the convex relaxation and the actual sum capacity can be made arbitrarily large. In light of this result, developing algorithms to compute, or at least better approximate, the sum capacity becomes important, and we undertake this task here.

As discussed in Sec. II-C, a two-sender MAC with input alphabets \mathcal{B}_1 and \mathcal{B}_2 and output alphabet \mathcal{Z} is described by a transition matrix \mathcal{N} . Each entry of this matrix, $\mathcal{N}(z|b_1, b_2)$, where $z \in \mathcal{Z}$, $b_1 \in \mathcal{B}_1$, and $b_2 \in \mathcal{B}_2$, represents the probability that the channel output takes a value z when the first and second channel inputs take values b_1 and b_2 , respectively. We introduce new notation for this sub-section. Let d_1 , d_2 , and d_o denote $|\mathcal{B}_1|$, $|\mathcal{B}_2|$, and $|\mathcal{Z}|$ respectively. We write $p(b_1) \in \Delta_{d_1}$, but refer to $p(b_2)$ by $q(b_2)$ where $q(b_2) \in \Delta_{d_2}$. In this notation, the sum capacity (7) of the two-sender MAC \mathcal{N} takes the form

$$S(\mathcal{N}) = \max_{p(b_1, b_2)} I(B_1, B_2; Z) \quad \text{such that} \quad p(b_1, b_2) = p(b_1)q(b_2), \quad (55)$$

where B_1 , B_2 , and Z are random variables that describe the channel's first input, second input, and output respectively. For fixed $\mathcal{N}(z|b_1, b_2)$, the mutual information $I(B_1, B_2; Z)$ function is concave in the argument $p(b_1, b_2)$. On the other hand, the set of joint distributions which satisfy the product constraint, $p(b_1, b_2) = p(b_1)p(b_2)$, is not convex. This lack of convexity turns the maximization in Eq. (55) into a non-convex problem.

Our approach to solving the non-convex problem (55) is to move the non-convexity from the constraint to the objective function. Instead of maximizing over the set of product distributions, $p(b_1, b_2) = p(b_1)q(b_2)$, we maximize sequentially,

that is, we write

$$S(\mathcal{N}) = \max_{q \in \Delta_{d_2}} \max_{p \in \Delta_{d_1}} I(B_1, B_2; Z). \quad (56)$$

Now, both the inner and outer maximization in Eq. (56) are carried out over convex sets Δ_{d_1} and Δ_{d_2} , respectively. To carry out these optimizations we derive certain convenient expressions. The output probability distribution p^Z over \mathcal{Z} can be written as

$$p^Z(z) = \sum_{b_1 \in \mathcal{B}_1} A_q(z, b_1)p(b_1) \quad (57)$$

where

$$A_q(z, b_1) = \sum_{b_2 \in \mathcal{B}_2} \mathcal{N}(z|b_1, b_2)q(b_2) \quad (58)$$

for $z \in \mathcal{Z}$ and $b_1 \in \mathcal{B}_1$. Note that A_q can be considered as a left stochastic matrix of size $d_o \times d_1$, i.e., every entry of A_q is non-negative and the columns sum to 1. One can view Eq. (57) as a vector equation $p^Z = A_q p$, where $p \in \Delta_{d_1}$ and $p^Z \in \Delta_{d_o}$. The mutual information $I(B_1, B_2; Z)$ can be written as

$$I(B_1, B_2; Z) = H(Z) - \sum_{b_1, b_2} p(b_1)q(b_2)H(Z|B_1=b_1, B_2=b_2).$$

To express the mutual information in terms of vectors and matrices, we define a d_1 -dimensional vector b_q with non-negative components $b_q(b_1)$ for $b_1 = 1, \dots, d_1$, where

$$b_q(b_1) = - \sum_{b_2 \in \mathcal{B}_2} q(b_2) \sum_{z \in \mathcal{Z}} \mathcal{N}(z|b_1, b_2) \log(\mathcal{N}(z|b_1, b_2)). \quad (59)$$

This allows us to express the mutual information compactly as

$$I(p, q) \equiv I(B_1, B_2; Z) = H(A_q p) - \langle b_q, p \rangle, \quad (60)$$

and the sum capacity as

$$S(\mathcal{N}) = \max_{q \in \Delta_{d_2}} \max_{p \in \Delta_{d_1}} \{H(A_q p) - \langle b_q, p \rangle\}. \quad (61)$$

The inner optimization over $p \in \Delta_{d_1}$ is a convex optimization problem because $I(p, q)$ is a concave function of p for fixed q . Since $I(p, q)$ is not jointly concave over (p, q) , the function

$$I^*(q) = \max_{p \in \Delta_{d_1}} (H(A_q p) - \langle b_q, p \rangle) \quad (62)$$

is not concave in general. Therefore, the outer optimization of $I^*(q)$ over $q \in \Delta_{d_2}$, i.e., $S(\mathcal{N}) = \max_{q \in \Delta_{d_2}} I^*(q)$, is in general a nonconvex problem.

Nevertheless, we show that the non-concave function $I^*(q)$ is a Lipschitz-like function as defined in Eq. (48). To elaborate, this means there is some real-valued function β_I that is non-negative, continuous, and monotonically increasing with $\beta_I(0) = 0$, such that $|I^*(q) - I^*(q')| \leq \beta_I(\|q - q'\|_1)$. We show that $I^*(q)$ indeed satisfies such a property by proving an appropriate continuity bound for $I^*(q)$. This is summarized in the following result.

Proposition 12: Let \mathcal{N} be any two-sender MAC with input alphabets \mathcal{B}_1 , \mathcal{B}_2 of size d_1 , d_2 , and output alphabet \mathcal{Z} of size d_o . Assume that $d_1, d_2, d_o \geq 2$. Given input probability

TABLE II

COMPUTING THE MAXIMUM OF LIPSCHITZ-CONTINUOUS FUNCTIONS f OVER THE STANDARD SIMPLEX $\Delta_d \subseteq \mathbb{R}^d$ IN DIMENSION $d = 3$ TO A PRECISION OF $\epsilon = 0.15$. GRID SEARCH (ALG. 2) AND THE ALGORITHM BASED ON DENSE CURVES (ALG. 4) ARE USED TO COMPUTE THE MAXIMUM NUMERICALLY. NUMERICAL VALUES ARE ROUNDED TO THREE DECIMAL PLACES

function	algorithm	true maximum	computed maximum	iterations
$f(x) = \sin(\ x\ _2)$	grid search	0.841	0.841	16290
	dense curve	0.841	0.836	155
$f(x) = -\frac{\ x\ _2^3}{6} + \frac{\ x\ _2^2}{4} - \frac{\ x\ _2}{6\pi}$	grid search	0.033	0.033	19900
	dense curve	0.033	0.033	480

distributions p over \mathcal{B}_1 and q over \mathcal{B}_2 , the mutual information between the inputs and the output of the MAC can be written as

$$I(p, q) = H(A_q p) - \langle b_q, p \rangle$$

where the matrix A_q and the vector b_p are defined in Eq. (58) and Eq. (59) respectively. Define the function

$$I^*(q) = \max_{p \in \Delta_{d_1}} I(p, q).$$

Then, for any $p \in \Delta_{d_1}$, and any $q, q' \in \Delta_{d_2}$, we have

$$|I(p, q) - I(p, q')| \leq \beta_I(\|q - q'\|_1),$$

and subsequently,

$$|I^*(q) - I^*(q')| \leq \beta_I(\|q - q'\|_1). \quad (63)$$

The function β_I is defined as

$$\beta_I(x) = \left(\frac{1}{2} \log(d_o - 1) + H_{\mathcal{N}}^{\max} \right) x + \bar{h}\left(\frac{x}{2}\right), \quad (64)$$

where

$$H_{\mathcal{N}}^{\max} = \max_{a_1 \in \mathcal{A}_1, a_2 \in \mathcal{A}_2} \left\{ - \sum_{z \in \mathcal{Z}} \mathcal{N}(z|a_1, a_2) \log(\mathcal{N}(z|a_1, a_2)) \right\}, \quad (65)$$

and

$$\bar{h}(x) = \begin{cases} -x \log(x) - (1-x) \log(1-x) & \text{if } x \leq \frac{1}{2} \\ \log(2) & \text{if } x \geq \frac{1}{2} \end{cases}$$

is the modified binary entropy defined in Eq. (49).

Proof: See Appendix D. \square

This observation is important in the sense that it allows for off-the-shelf use of algorithms developed in Sec. IV for optimizing any Lipschitz-like function in order to compute the sum capacity of two-sender MACs. Following this line of approach, we will begin by developing an efficient algorithm for computing the sum capacity of any two-sender MAC where one of the input alphabets is of size 2. As a result, we can efficiently compute the sum capacity of a large family of MACs that includes all binary MACs. Next, we will develop two algorithms for computing the sum capacity of an arbitrary two-sender MAC. The first algorithm is important from a theoretical standpoint, and shows that sum capacity of an arbitrary two-sender MAC can be computed in quasi-polynomial time. However, it can be costly to implement in practice. Our second algorithm can be faster to run in practice, at least when one of the MAC input alphabet sizes is small, but it suffers from exponential complexity in the worst case.

B. Computing the Sum Capacity of a Two-Sender MAC With One Input Alphabet of Size 2

As before, we take \mathcal{N} to be a two-sender MAC with input alphabets $\mathcal{B}_1, \mathcal{B}_2$ of sizes d_1, d_2 and an output alphabet \mathcal{Z} of size d_o . In this section, we focus on the case where at least one of d_1 or d_2 is equal to 2. For concreteness, take $d_2 = 2$ (this choice is inconsequential for the algorithm).

For this case, any probability distribution $q_s \in \Delta_{d_2}$ can be expressed as $q_s = (s, 1-s)$ for some $0 \leq s \leq 1$. Thus, the maximization over $q_s \in \Delta_2$ in computing the sum capacity $S(\mathcal{N}) = \max_{q_s \in \Delta_{d_2}} I^*(q_s)$ is essentially one-dimensional. In other words, considering the objective $I^*(s) := I^*(q_s)$ as a function of s , we can write the sum capacity for any MAC \mathcal{N} with $d_2 = 2$ as

$$S(\mathcal{N}) = \max_{s \in [0, 1]} I^*(s).$$

We will show that the mutual information $I^*(s)$ considered as a function of s is still a Lipschitz-like function. First, observe that for $q_s = (s, 1-s)$, we have $\|q_s - q_{s'}\|_1 = 2|s - s'|$. Then, $|I^*(s) - I^*(s')| = |I^*(q_s) - I^*(q_{s'})| \leq \beta_I(\|q_s - q_{s'}\|_1) = \beta_I(2|s - s'|)$. Therefore, taking $\beta(x) = \beta_I(2x)$, we find that I^* as a function of s is β -Lipschitz-like. Subsequently, we can use modified Piyavskii-Shubert algorithm developed in Sec. IV-C to compute the maximum $\max_{s \in [0, 1]} I^*(s)$ to any given precision.

For the convenience of the reader, we rewrite Alg. 1 specifically for the purpose of computing the sum capacity of a two-sender MAC \mathcal{N} when $d_2 = 2$. This sum capacity computation algorithm is summarized in Alg. 5 below. We note that similar to Alg. 1, one can modify Alg. 5 so that it accepts a fixed number of iterations and outputs the value $F(s^*)$, which is an upper bound on the sum capacity. This upper bound exceeds the sum capacity by at most $F(s^*) - I^*(s^*)$.

We make a technical remark that is important to perform the optimization as per Alg. 5. The objective function I^* appearing in the optimization $S(\mathcal{N}) = \max_{s \in [0, 1]} I^*(s)$ is non-trivial to compute. More precisely, in order to compute the value of $I^*(s)$ for any given s , we need to solve another optimization problem. This follows from the definition of I^* given in Eq. (62). Fortunately, this optimization problem is convex, and hence, it can be solved by standard convex optimization techniques.

In Prop. 22, we show that the number of iterations required by the while loop in Alg. 5 to converge to the sum capacity within a tolerance of $0 < \epsilon \leq 3$ is bounded above as $O(\log(d_o)/\epsilon)$, where $d_o \geq 2$ is the size of the output alphabet

Algorithm 5 Computing the Sum Capacity of a Two-Sender MAC \mathcal{N} to a Precision $\epsilon > 0$, When One of the Input Alphabets has Size 2

```

1: function COMPUTE_SUM_CAPACITY( $\mathcal{N}, d_1, \epsilon$ )
2:   Initialize  $s^{(0)} = 0$ 
3:   Define  $F_0(s) = I^*(s^{(0)}) + \beta_I(2|s - s^{(0)}|)$ 
   for  $s \in [0, 1]$ 
4:   Set  $F \leftarrow F_0$  and  $s^* \leftarrow 1$ 
5:   Set  $s^{(1)} \leftarrow s^*, k \leftarrow 1$ 
6:   while  $F(s^*) - I^*(s^*) > \epsilon$  do
7:     Sort  $\{s^{(0)}, \dots, s^{(k)}\}$  from smallest to largest and
       relabel the points in ascending order.
8:     Set  $F_i(s) = I^*(s^{(i)}) + \beta_I(2|s - s^{(i)}|)$ 
     for  $0 \leq i \leq k$ 
9:     for  $i = 0, \dots, k - 1$  do
10:    Set  $g_i(s) = F_i(s) - F_{i+1}(s)$ 
11:    Find  $\bar{s}_i \in [s^{(i)}, s^{(i+1)}]$  such that  $g_i(\bar{s}_i) = 0$ 
       using any root finding method.
12:   end for
13:   Pick an index  $m \in \operatorname{argmax}_{0 \leq i \leq k-1} F_i(\bar{s}_i)$  and
     set  $s^* = \bar{s}_m$ .
14:   Update  $F \leftarrow F_m$ 
15:   Set  $s^{(k+1)} \leftarrow s^*, k \leftarrow k + 1$ 
16: end while
17: return  $I^*(s^*)$ 
18: end function

```

and ϵ is a fixed constant. Note that this bound does not account for the number of iterations required to compute I^* , for sorting or for root-finding. We show in Prop. 22 that the total cost involved is at most polynomial in d_1 , d_o , and $1/\epsilon$.

Next, we show how to compute the sum capacity of any two-sender MAC.

C. Computing the Sum Capacity of Any Two-Sender MAC

Let an arbitrary two-sender MAC \mathcal{N} with input alphabet sizes d_1, d_2 and output alphabet size d_o . Without loss of generality, we assume that $d_2 \leq d_1$. Then, the sum capacity computation can be expressed as $S(\mathcal{N}) = \max_{q \in \Delta_{d_2}} I^*(q)$. This amounts to optimizing the Lipschitz-like function $I^*(q)$ over the standard simplex Δ_{d_2} . As a result, the algorithms developed in Sec. IV-D.1 and Sec. IV-D.2 can both be used to compute the sum capacity of \mathcal{N} . For the convenience of the reader, we present these algorithms here again adapted specifically for sum capacity computation.

The first algorithm that we discuss is grid search. This algorithm is helpful in proving complexity results, but not as helpful from a practical implementation standpoint.

1) *Sum Capacity Computation Using Grid Search:* Our goal is to perform the optimization of $I^*(q)$ over the simplex $q \in \Delta_{d_2}$. We perform this optimization by computing the maximum over the grid

$$\Delta_{d_2, N} = \left\{ \frac{n}{N} \mid n \in \mathcal{I}_{d_2, N} \right\}$$

defined in Eq. (52). Thus, the grid search algorithm can be described as follows.

Algorithm 6 Computing the Sum Capacity of a Two-Sender MAC \mathcal{N} to a Precision $\epsilon > 0$ Using Grid Search. d_1, d_2 Are Input Alphabet Sizes With $d_2 \leq d_1$ and d_o Is the Output Alphabet Size. The Function β_I Is Defined in Eq. (64)

```

1: function
   COMPUTE_SUM_CAP_GRID_SEARCH( $\mathcal{N}, d_1, d_2, d_o, \epsilon$ )
2:   Set  $N = \lceil 1/\delta^2 \rceil$ , where  $\delta$  is the largest number
     satisfying  $\beta_I(\delta) \leq \epsilon/2$ 
3:   Construct the grid  $\Delta_{d_2, N}$ 
4:   return  $S^* = \max\{I^*(q) \mid q \in \Delta_{d_2, N}\}$ 
5: end function

```

As noted in Sec. IV-D.1, optimization of a β -Lipschitz-like functions using grid search can be done in polynomial time if the function β does not depend on the dimension. For the case of sum capacity computation, the function β is β_I given in Eq. (64). β_I depends on the size of the output alphabet of the MAC, and therefore, the complexity analysis is more involved. Furthermore, for any given $q \in \Delta_{d_2}$, the function $I^*(q)$ needs to be computed using convex optimization. Thus, the cost of computing I^* also needs to be included in the complexity analysis.

In Prop. 23, we show that the total cost of computing the sum capacity to a fixed precision $0 < \epsilon \leq 1$ is roughly bounded above by $\operatorname{poly}(d_1, d_o, 1/\epsilon) O(d_2^{96 \log^2(d_o)/\epsilon^2 + 2})$ (see Eq. (113) for a more precise statement). Therefore, the sum capacity can be computed to a fixed precision $\epsilon > 0$ in quasi-polynomial time. This quasi-polynomial behaviour comes from the fact that β_I depends on $\log(d_o)$. From the proof of Prop. 23, we can also infer that if one fixes the output alphabet size d_o , then the sum capacity can be computed in polynomial time.

While these results are useful from a theoretical standpoint, grid search can be slow in practice. The reason is two-fold. One, the power of the polynomial appearing in time complexity can be large even for moderately small precision and dimensions (i.e., the size of the grid becomes fairly large). Two, for each point in the grid, we need to solve a convex optimization problem for computing I^* . Together, these factors make grid search not ideal for practical implementations. We note that one can nevertheless parallelize grid search to get some improvements in the speed, though eventually even this will be too costly. With this in mind, we study another algorithm to compute the sum capacity using dense curves.

2) *Sum Capacity Computation Using Dense Curves:* In this section, we will study the application of algorithm developed in Sec. IV-D.2 for computing the sum capacity. The idea is to fill the simplex with a curve that comes within a distance of α to any given point on the simplex. Such a curve is therefore called an α -dense curve. Using such a curve, we reduce a high-dimensional optimization problem to a one-dimensional optimization problem over an interval. The one-dimensional optimization problem can be solved using the modified Piyavskii-Shubert algorithm that underlies Alg. 5.

We refer the reader to Alg. 3 which shows how to construct an α -dense curve for filling the standard simplex for $\alpha = 2(d_2 - 1)/N$, where N is a positive integer that controls the

value of α . Below, we show how to use Alg. 4 to compute the sum capacity of a two-sender MAC.

Algorithm 7 Computing the Sum Capacity of a Two-Sender MAC \mathcal{N} to a Precision $\epsilon > 0$ Using a Dense Curve. d_1, d_2 Are Input Alphabet Sizes With $d_2 \leq d_1$ and d_o Is the Output Alphabet Size. The Function β_I Is Defined in Eq. (64)

```

1: function
  COMPUTE_SUM_CAP_DENSE_CURVE( $\mathcal{N}, d_1, d_2, d_o, \epsilon$ )
2:   Compute the largest number  $\alpha > 0$  such that
    $\beta_I(\alpha) \leq \epsilon/2$ 
3:   Set  $N = \lceil 2(d_2 - 1)/\alpha \rceil$ 
4:   Construct the  $(2(d_2 - 1)/N)$ -dense curve
    $\gamma: [0, L_{\text{curve}}] \rightarrow \mathcal{D}$  as per Alg. 3
5:   Compute the maximum  $S^*$  of the function  $I^* \circ \gamma$  over
    $[0, L_{\text{curve}}]$  to a precision of  $\epsilon/2$  using Alg. 1
6:   return  $S^*$ 
7: end function

```

From Prop. 18, we can infer that the above algorithm is guaranteed to compute the sum capacity to within a precision of $\epsilon > 0$. Prop. 18 shows that the dense curve algorithm can take an exponential time in the worst case to converge. For this reason, we avoid doing a detailed complexity analysis of this algorithm for sum capacity computation, as better theoretical guarantees can be obtained using grid search. Instead, we focus on more practical gains that may be obtained using Alg. 7 over Alg. 6.

Using Alg. 7 gives another advantage. Instead of computing the sum capacity to a given precision, one can run Alg. 7 for a fixed number of time steps to get an upper bound on the sum capacity. Thus, it can also be used for bounding the sum capacity instead of computing it to a good precision.

3) *Comparing Grid Search and Dense Curve Algorithm Performance:* We compare the performance of the grid search algorithm and the dense curve algorithm for computing the sum capacity. For this purpose, we will consider a randomly constructed two-sender MAC with input alphabet sizes d_1, d_2 and output alphabet size d_o . We fix $d_1 = 10$ and $d_o = 20$ and consider the cases $d_2 = 2$ and $d_2 = 3$. Note that the smaller dimension d_2 determines the dimension for the non-convex optimization. Hence, the value of d_2 effectively decides the overall computation time. The numerical simulations are run on single core of a personal computer, and the code is implemented in Python (see Sec. E-B for more details). All the run times reported are average time taken over three repetitions of the simulation.

For $d_2 = 2$, the dense curve algorithm can be simplified to modified Piyavskii-Shubert algorithm, as noted in Sec. V-B. Thus, for $d_2 = 2$, we compare the performance of modified Piyavskii-Shubert Alg. 5 with that of grid search Alg. 6 for computing the sum capacity to a tolerance of $\epsilon = 0.15$. For the example we consider, we find that Alg. 5 takes about 0.34 s to run, whereas Alg. 6 takes about 1.51 mins to run. Thus, for this example, modified Piyavskii-Shubert algorithm is more than 250 times faster than grid search. For a tolerance of $\epsilon = 0.65$, the run time of grid search Alg. 2 improves

greatly to 1.06 s, whereas the run time of modified Piyavskii-Shubert Alg. 5 improves to 0.11 s. In this case, the modified Piyavskii-Shubert algorithm is almost 10 ten times faster than grid search. We remark that a tolerance of $\epsilon = 0.65$ is not acceptable in practice because the computed sum capacity is about 0.19 nats. We therefore use such a large value of ϵ only for benchmarking purposes.

Next, we consider the higher dimensional case of $d_2 = 3$. For this case, we compare the dense curve Alg. 7 with grid search Alg. 6. For a tolerance of $\epsilon = 0.65$, we find that the dense curve Alg. 7 takes 3.55 s whereas grid search Alg. 6 takes 3.65 min. In this case, dense curve algorithm is about 65 times faster, though as noted previously, a tolerance of $\epsilon = 0.65$ cannot be used in practice. For $\epsilon = 0.15$, the dense curve Alg. 7 takes about 1.72 mins, whereas grid search takes too long to complete. We estimate that our current implementation of grid search algorithm will take more than 3 days to complete for $\epsilon = 0.15$ on the hardware the code was executed. This shows that for the example under consideration, grid search is not a practical option.

We remark that there can be some cases where grid search performs as good as or better than the dense curve algorithm. For small dimensions, grid search scales poorly with the tolerance in comparison with dense curve and modified Piyavskii-Shubert algorithms, as evidenced from the above numerical simulations. As the dimension d_2 becomes large and d_1, d_o , and ϵ are fixed, the grid search algorithm scales polynomially with d_2 whereas dense curve algorithm scales exponentially. Despite this, the polynomial power can be so large in practice that grid search is still impractical. Thus, our numerical simulations suggest that for small dimensions and small tolerance, modified Piyavskii-Shubert algorithm (for $d_2 = 2$) and the dense curve algorithm (for $d_2 > 2$) perform better than the grid search algorithm in practice.

As noted previously, both the grid search algorithm and the dense curve algorithm (as well as modified Piyavskii-Shubert algorithm) can be parallelized to get better performance. It remains to see if these algorithms can be improved further so as to make them practical for larger values of d_2 . Finally, we show how our algorithms compare with the relaxed sum capacity.

D. Comparison With Relaxed Sum Capacity

In order to compare our algorithms for computing the sum capacity with the relaxed sum capacity, we construct a family of binary MACs parametrized by a tunable parameter $t \in [0, 1]$. These MACs have the property that for $t = 0$, the relaxed sum capacity is equal to the actual sum capacity, whereas for $t = 1$, the relaxed sum capacity is twice the actual sum capacity. By computing the sum capacity using our algorithms as well as the relaxed sum capacity for as a function of t , we can compare how well our algorithms do in relation to the convex relaxation approach.

For generating examples of such binary MACs, we construct a family of MACs that we call the noise-free subspace MAC. To that end, let \mathcal{A} and \mathcal{B} be input alphabets and let \mathcal{Z} be the output alphabet of the noise-free subspace MAC.

We assume that \mathcal{A} , \mathcal{B} and \mathcal{Z} are finite sets. The noise-free subspace corresponds to a set $\mathcal{W} \subseteq \mathcal{A} \times \mathcal{B}$. Consider the mapping $n_F: \mathcal{W} \rightarrow \mathcal{Z}$ that determines the symbol that is deterministically output by the channel when the input is in the noise-free subspace. Then, given the tuple $(\mathcal{A}, \mathcal{B}, \mathcal{Z}, \mathcal{W}, n_F)$, the MAC \mathcal{N}_F has the probability transition matrix

$$\mathcal{N}_F(z|a, b) = \begin{cases} \delta_{z, n_F(a, b)} & (a, b) \in \mathcal{W} \\ \frac{1}{|\mathcal{Z}|} & (a, b) \notin \mathcal{W} \end{cases} \quad (66)$$

When the input is in the noise-free subspace, the channel deterministically outputs the symbol selected by n_F . On the other hand, when the input is not in the noise-free subspace, the channel outputs a symbol uniformly at random. Therefore, a noise-free subspace MAC can be thought of as a generalization of nonlocal games MAC studied in Part I (see Eq. (9)).

Now, for constructing the parametrized family of binary MACs mentioned at the beginning of the section, we construct two examples of the MAC \mathcal{N}_F corresponding to the “extremities” of the parametrized family. For both of these examples, we consider the alphabets $\mathcal{A} = \{a_1, a_2\}$, $\mathcal{B} = \{b_1, b_2\}$, and $\mathcal{Z} = \{z_1, z_2\}$. The examples are labelled 0 and 1.

For the first example, we take $\mathcal{W} = \{(a_1, b_1)\}$ and $n_F(a_1, b_1) = z_1$. That is, only a single input (a_1, b_1) is transmitted noise-free. For this example, we can write the probability transition matrix as

$$\mathcal{N}_F^{(0)} = \begin{pmatrix} 1 & 0.5 & 0.5 & 0.5 \\ 0 & 0.5 & 0.5 & 0.5 \end{pmatrix}$$

where the rows correspond to $z \in \mathcal{Z}$, while the columns correspond to $(a, b) \in \mathcal{A} \times \mathcal{B}$. In App. E, we show that the sum capacity in nats is

$$S(\mathcal{N}_F^{(0)}) = h\left(\frac{4}{5}\right) - \frac{2}{5} \ln(2) \approx 0.223$$

where h is the binary entropy measured in nats. We furthermore show that the relaxed sum capacity in this case is also equal to $C(\mathcal{N}_F^{(0)}) = 0.223$ nats. Thus, for our first example, we have $C(\mathcal{N}_F^{(0)}) = S(\mathcal{N}_F^{(0)})$.

For the second example, we take $\mathcal{W} = \{(a_1, b_1), (a_2, b_2)\}$ as well as $n_F(a_1, b_1) = z_1$ and $n_F(a_2, b_2) = z_2$. That is, the inputs (a_1, b_1) and (a_2, b_2) are transmitted noise-free. The probability transition matrix in this case can be written as

$$\mathcal{N}_F^{(1)} = \begin{pmatrix} 1 & 0.5 & 0.5 & 0 \\ 0 & 0.5 & 0.5 & 1 \end{pmatrix}$$

where the rows correspond to $z \in \mathcal{Z}$, while the columns correspond to $(a, b) \in \mathcal{A} \times \mathcal{B}$. In App. E, we show that the sum capacity of the MAC $\mathcal{N}_F^{(1)}$ in nats is equal to

$$S(\mathcal{N}_F^{(1)}) = 0.5 \ln(2) \approx 0.3466.$$

On the other hand, we show that in this example the relaxed sum capacity takes the maximum possible value $C(\mathcal{N}_F^{(1)}) = \ln(2)$ nats, thus significantly overestimating the sum capacity. For this example, we have $C(\mathcal{N}_F^{(1)}) = 2 S(\mathcal{N}_F^{(1)})$.

Using these examples, we construct the parametric family of binary MACs as a convex combination of the MACs in above examples. That is,

$$\mathcal{N}_F^{(t)} = (1-t)\mathcal{N}_F^{(0)} + t\mathcal{N}_F^{(1)}$$

$$= \begin{pmatrix} 1 & 0.5 & 0.5 & 0.5(1-t) \\ 0 & 0.5 & 0.5 & 0.5(1+t) \end{pmatrix} \quad (67)$$

where $t \in [0, 1]$. Observe that for $t = 0$, we get the first example $\mathcal{N}_F^{(0)}$, whereas for $t = 1$, we get the second example $\mathcal{N}_F^{(1)}$. From the above results, we know that $C(\mathcal{N}_F^{(0)}) = S(\mathcal{N}_F^{(0)})$ whereas $C(\mathcal{N}_F^{(1)}) = 2 S(\mathcal{N}_F^{(1)})$. The MAC $\mathcal{N}_F^{(t)}$ basically interpolates between these two cases.

In Fig. 3, we plot $S(\mathcal{N}_F^{(t)})$ and $C(\mathcal{N}_F^{(t)})$ as a function of t . The sum capacity $S(\mathcal{N}_F^{(t)})$ is computed using Alg. 5 with a tolerance of $\epsilon = 0.01$, whereas relaxed sum capacity $C(\mathcal{N}_F^{(t)})$ is computed using standard techniques in convex optimization. We can see from the figure that at $t = 0$ and $t = 1$, the numerically computed values agree with the analytical results. Furthermore, we observe that $C(\mathcal{N}_F^{(t)})$ becomes a progressively worse bound on $S(\mathcal{N}_F^{(t)})$ as t ranges from 0 to 1. Thus, we demonstrate that our algorithm for computing the sum capacity does better than the relaxed sum capacity.

VI. CONCLUSION AND FUTURE DIRECTIONS OF RESEARCH

Computing the sum capacity of a multiple access channel is a nonconvex optimization problem. For MACs obtained from nonlocal games, we obtained an analytical upper bound on the sum capacity that depends only on the number of question tuples in the game and the maximum winning probability of the game when the questions are drawn uniformly at random. Our formula is an upper bound on the achievable sum rate even when the senders of the MAC can share an arbitrary set of correlations. Using this formula, we found a separation between the sum capacity and the entanglement-assisted sum rate for the 2-sender MAC obtained from the Magic Square game that is larger than the previously reported value. We also obtained separations in some other relevant scenarios using the CHSH game and multiparty parity game.

Furthermore, we studied the performance of the upper bound on the sum capacity obtained by relaxing the nonconvex problem to a convex optimization problem. By constructing the signalling game, we showed that one can obtain an arbitrarily large separation between the sum capacity and the relaxed sum capacity. With the help of numerical simulations, we argued that this separation holds even when the senders are allowed to share no-signalling correlations. These results indicate that the relaxed sum capacity can be a very poor upper bound on the sum capacity. In a recent work, Fawzi & Fermé [38] compute the no-signalling assisted sum rate for MACs, allowing feedback. In their study, they pose the question as to whether the no-signalling assisted sum rate (allowing feedback) is equal to the relaxed sum capacity. It would be interesting to see if the MAC obtained from the signalling game can be used to answer this question in the negative.

In response to the above observations, we studied algorithms to compute the sum capacity. First, we identified that the mutual information occurring in the computation of the sum capacity satisfies a Lipschitz-like property. We subsequently proposed a few algorithms one can use to optimize such functions, by appropriately modifying and generalizing existing algorithms for optimizing Lipschitz-continuous functions.

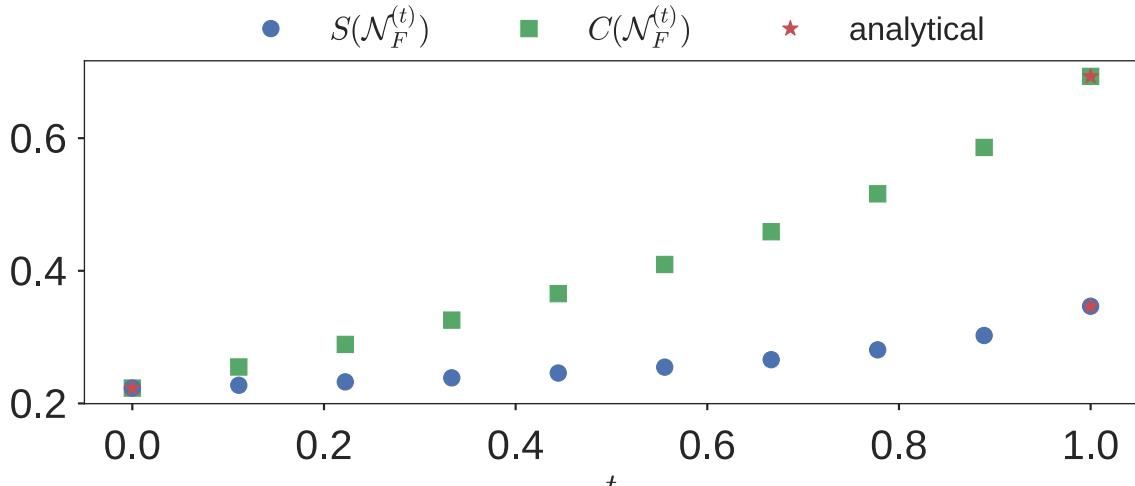


Fig. 3. Plot of the sum capacity $S(\mathcal{N}_F^{(t)})$ (in nats) and the relaxed sum capacity $C(\mathcal{N}_F^{(t)})$ (in nats) as a function of the tunable parameter t . The parameterized family of binary MACs $\mathcal{N}_F^{(t)}$ is defined in Eq. (67). The sum capacity is computed to within a tolerance of $\epsilon = 0.01$ using Alg. 5. We find that the relaxed sum capacity gives a progressively worse bound on the actual sum capacity. We also find that the analytically computed values of $S(\mathcal{N}_F^{(t)})$ and $C(\mathcal{N}_F^{(t)})$ for $t = 0$ and $t = 1$ show good agreement with the numerically computed values.

Using this, we were able to show that the sum capacity of any two-sender MAC can be computed to a fixed precision in quasi-polynomial time. Our algorithms are practically efficient for computing the sum capacity of a family of two-sender MACs that have at least one input alphabet of size two.

We remark that some other entropic quantities also satisfy a Lipschitz-like property. This is also true in the quantum setting, for example for the von Neumann entropy. Therefore, further investigation of algorithms to optimize such Lipschitz-like functions might be helpful in solving nonconvex problems in both classical and quantum information theory. In particular, it would be interesting to see if there are other practically relevant problems in information theory which would benefit from such an approach. In any case, the algorithms we present in this study for optimizing Lipschitz-like functions over the standard simplex suffer from the drawback that they are not scalable, i.e., the optimization is very costly to perform in practice as the dimension increases. Therefore, finding more practical algorithms to perform this optimization is an interesting direction for future research. In particular, the algorithm for optimization using dense curves has a scope for improvement because we only prove sub-optimal convergence guarantees.

Another avenue for performing such nonconvex optimization is to devise randomized algorithms, which might allow for faster convergence (with high probability). We remark that using randomized algorithms to optimize Lipschitz-like functions over an arbitrary compact and convex domain is not expected to give significant improvements over the deterministic algorithm we presented in this study. This is because the number of iterations needed for convergence will scale exponentially with the dimension in both the deterministic and stochastic setting [39]. It is therefore important to use information about the domain (for example, standard simplex) or impose some additional restrictions on the objective functions while designing such algorithms. It would also be interesting to see if there is a general method to determine a priori if

the convex relaxation gives a good or a bad bound on the optimum of the non-convex problem. This would allow one to use the more efficiently computable convex relaxation in relevant scenarios.

APPENDIX A BOUNDING THE CORRELATION-ASSISTED ACHIEVABLE SUM RATE OF MACS FROM NONLOCAL GAMES

Proof of Proposition 2: We solve the minimization problem $\min_{\pi \in \Delta_d} -\mathcal{I}_w(\pi)$, which is equivalent to the given maximization problem. Note that since the entries of w are either 0 or 1, we can write $\mathcal{K} = \{i \in [d] \mid w_i \neq 0\}$. Then we can write the Lagrangian for this minimization problem as

$$\begin{aligned} \mathcal{L}(\pi; \lambda, \nu) &= -\mathcal{I}_w(\pi) - \langle \lambda, \pi \rangle + \nu \left(\sum_{i=1}^d \pi_i - 1 \right) \\ &= \sum_{i=1}^d p_i \ln p_i - \sum_{i \in \mathcal{K}} \pi_i w_i \ln d + \ln d \\ &\quad - \sum_{i=1}^d \lambda_i \pi_i + \nu \left(\sum_{i=1}^d \pi_i - 1 \right), \end{aligned} \quad (68)$$

where

$$p_i = (\bar{W}\pi)_i = \pi_i w_i + \frac{1}{d} - \frac{1}{d} \sum_{j=1}^d \pi_j w_j. \quad (69)$$

The expression for the probability distribution p , which is the output of the channel \mathcal{N}_G , is obtained from Eqs. (26) and (27). The variables $\lambda_1, \dots, \lambda_d$ are the dual variables corresponding to the inequality constraint $\pi_i \geq 0$ for all $i \in [d]$, and ν is the dual variable corresponding to the equality constraint $\sum_{i=1}^d \pi_i = 1$. We also write $K = |\mathcal{K}|$.

Case 1: $0 < K < d$

We can write

$$p_i = \begin{cases} \pi_i + \frac{1}{d} - \frac{1}{d} \sum_{j \in \mathcal{K}} \pi_j & \text{for } i \in \mathcal{K} \\ \frac{1}{d} - \frac{1}{d} \sum_{j \in \mathcal{K}} \pi_j & \text{for } i \notin \mathcal{K}. \end{cases} \quad (70)$$

Since $p_i \geq 0$ for all i , we must have

$$\sum_{j \in \mathcal{K}} \pi_j \leq 1,$$

Then, we can consider two sub-cases: either (a) $\sum_{j \in \mathcal{K}} \pi_j = 1$ or (b) $\sum_{j \in \mathcal{K}} \pi_j < 1$ at the optimum.

(a) If $\sum_{j \in \mathcal{K}} \pi_j = 1$ at the optimum, then from Eq. (70), we get

$$p_i = \begin{cases} \pi_i & \text{for } i \in \mathcal{K} \\ 0 & \text{for } i \notin \mathcal{K}. \end{cases} \quad (71)$$

Furthermore, since $\sum_{j \in \mathcal{K}} \pi_j = 1$, we also have $\pi_i = 0$ for $i \notin \mathcal{K}$. It follows from Eq. (29) and Eq. (71) that $\mathcal{J}_w(\pi) = H(p)$, where p is as given in Eq. (71). Thus, p can be taken as a probability distribution on the indices \mathcal{K} , and subsequently, we obtain

$$\max_{\pi} \mathcal{J}_w(\pi) = \ln K. \quad (72)$$

(b) If $\sum_{j \in \mathcal{K}} \pi_j < 1$ at the optimum, we can infer from Eq. (70) that $p_i > 0$ for all $i \in [d]$ at the optimum. Thus, the entropy $H(p)$ is differentiable at the optimum, and consequently, we can differentiate the Lagrangian given in Eq. (69). The gradient of the Lagrangian is given as

$$\frac{\partial \mathcal{L}}{\partial \pi_j} = \begin{cases} \ln p_j - \frac{1}{d} \sum_{i=1}^d \ln p_i - \ln d - \lambda_j + \nu & \text{for } j \in \mathcal{K} \\ -\lambda_j + \nu & \text{for } j \notin \mathcal{K}. \end{cases}$$

Note that we are solving a convex optimization problem and Slater's condition holds because the constraint set is a simplex [40]. Therefore, the KKT conditions are necessary and sufficient for optimality (see Ch. 5.5 in Ref. [40]). Subsequently, the optimal distribution π satisfies [40]

(Primal feasibility) $\pi \in \Delta_d$,

(Dual feasibility) $\lambda_i \geq 0$ for all $i \in [d]$,

(Complementary slackness) $\lambda_i \pi_i = 0$ for all $i \in [d]$, and

(Stationarity) $\nabla_{\pi} \mathcal{L} = 0$.

For $j \notin \mathcal{K}$, the condition $\nabla_{\pi} \mathcal{L} = 0$ gives

$$\lambda_j = \nu. \quad (73)$$

Since $\sum_{j \in \mathcal{K}} \pi_j < 1$ by assumption and $\pi \in \Delta_d$, we must have $\pi_{j'} > 0$ for some $j' \notin \mathcal{K}$. Then, by complementary slackness, we obtain $\lambda_{j'} = 0$. Using this in Eq. (73), we get $\nu = 0$.

On the other hand, for $j \in \mathcal{K}$ the condition $\nabla_{\pi} \mathcal{L} = 0$ gives

$$\ln p_j - \frac{1}{d} \sum_{i=1}^d \ln p_i = \lambda_j + \ln d,$$

where we used the fact that $\nu = 0$. To simplify this equation further, we make the observation that

$$p_i = \begin{cases} \pi_i + \frac{p_L}{d} & \text{for } i \in \mathcal{K} \\ \frac{p_L}{d} & \text{for } i \notin \mathcal{K} \end{cases}$$

with $p_L > 0$, which follows from Eq. (25) and Eq. (70). Therefore,

$$\sum_{i=1}^d \ln p_i = \sum_{i \in \mathcal{K}} \ln p_i + (d - K) \ln(p_L/d),$$

where $d - K = |[d] \setminus \mathcal{K}|$. Thus, for $j \in \mathcal{K}$, we obtain

$$\ln p_j - \frac{1}{d} \sum_{i \in \mathcal{K}} \ln p_i = \lambda_j + \ln d + \left(1 - \frac{K}{d}\right) \ln \frac{p_L}{d}.$$

If we label the indices in \mathcal{K} as j_1, \dots, j_K , we can write the above as the following matrix equation:

$$\begin{pmatrix} 1 - \frac{1}{d} & -\frac{1}{d} & \cdots & -\frac{1}{d} \\ \vdots & \cdots & \ddots & \vdots \\ -\frac{1}{d} & -\frac{1}{d} & \cdots & 1 - \frac{1}{d} \end{pmatrix} \begin{pmatrix} \ln p_{j_1} \\ \vdots \\ \ln p_{j_K} \end{pmatrix} = \begin{pmatrix} \lambda_{j_1} + \ln d + \left(1 - \frac{K}{d}\right) \ln \frac{p_L}{d} \\ \vdots \\ \lambda_{j_K} + \ln d + \left(1 - \frac{K}{d}\right) \ln \frac{p_L}{d} \end{pmatrix} \quad (74)$$

Let $o = (1 \ \cdots \ 1)^T$ denote the vector of ones. Then, the matrix appearing in the LHS of (74) can be expressed as $\mathbb{I} - oo^T/d$. By the matrix determinant lemma (see Cor. (18.1.3) in Ref. [41]), we know that this matrix is invertible iff $1 - o^T o/d = 1 - K/d \neq 0$, which is always true by our assumption that $K < d$. Its inverse is given by the Sherman-Morrison formula [42], [43] as follows:

$$\left(\mathbb{I} - \frac{1}{d} oo^T\right)^{-1} = \mathbb{I} + \frac{1}{d - K} oo^T.$$

Using this inverse in Eq. (74), for $j \in \mathcal{K}$, we obtain

$$p_j = E_j \frac{p_L}{d},$$

where

$$E_j = \exp \left(\lambda_j + \frac{1}{d - K} \sum_{i \in \mathcal{K}} \lambda_i + \frac{d}{d - K} \ln d \right).$$

Since $\lambda_i \geq 0$ for all $i \in [d]$ due to dual feasibility, we have $E_j > 1$ for all $j \in \mathcal{K}$.

Using $p_j = \pi_j + p_L/d$ for $j \in \mathcal{K}$, we obtain

$$\pi_j = (E_j - 1) \frac{p_L}{d}. \quad (75)$$

Because $E_j > 1$, we have $\pi_j > 0$ for each $j \in \mathcal{K}$. Then, by complementary slackness, we have $\lambda_j = 0$ for all $j \in \mathcal{K}$. Therefore, we obtain

$$E_j = d^{\frac{d}{d-K}} \quad \forall j \in \mathcal{K}. \quad (76)$$

Now, we will solve for π_j for $j \in \mathcal{K}$. For this purpose, note that

$$e_j^{-1} = \frac{d}{E_j - 1}$$

is a positive number for all $j \in \mathcal{K}$. Then, using $p_L = 1 - \sum_{i \in \mathcal{K}} \pi_i$ and $\pi_j = e_j p_L$ for $j \in \mathcal{K}$ (see Eq. (75)), we can write

$$e_j^{-1} \pi_j + \sum_{i \in \mathcal{K}} \pi_i = 1,$$

which can be written in matrix form as

$$\begin{pmatrix} e_{j_1}^{-1} + 1 & 1 & \cdots & 1 \\ \vdots & \cdots & \ddots & \vdots \\ 1 & 1 & \cdots & e_{j_K}^{-1} + 1 \end{pmatrix} \begin{pmatrix} \pi_{j_1} \\ \vdots \\ \pi_{j_K} \end{pmatrix} = \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix}.$$

The matrix appearing in the LHS of the above equation can be written as $\text{diag}(e_j^{-1}) + oo^T$, which is invertible iff $1 + \sum_{j \in \mathcal{K}} e_j \neq 0$. This condition always holds because $e_j > 0$ for each $j \in \mathcal{K}$. Denoting $A = \text{diag}(e_j)$, we can use the Sherman-Morrison formula [42], [43] to write

$$\begin{pmatrix} \pi_{j_1} \\ \vdots \\ \pi_{j_K} \end{pmatrix} = Ao - \frac{o^T Ao}{1 + o^T Ao} Ao = \begin{pmatrix} \frac{e_{j_1}}{1 + \sum_{j \in \mathcal{K}} e_j} \\ \vdots \\ \frac{e_{j_K}}{1 + \sum_{j \in \mathcal{K}} e_j} \end{pmatrix}. \quad (77)$$

Then, using $e_j = (E_j - 1)/d$ for $j \in \mathcal{K}$ along with Eq. (76) and Eq. (77), we obtain

$$\pi_j = \frac{d^{\frac{d}{d-K}} - 1}{d + K(d^{\frac{d}{d-K}} - 1)} \quad \forall j \in \mathcal{K}$$

and

$$p_L = \frac{d}{d + K(d^{\frac{d}{d-K}} - 1)}.$$

Subsequently, we obtain

$$\begin{aligned} \mathcal{I}_K^* &\equiv \max_{\pi \in \Delta_d} \mathcal{I}_w(\pi) \\ &= -\frac{Kd^{\frac{d}{d-K}}}{d + K(d^{\frac{d}{d-K}} - 1)} \ln \left(\frac{d^{\frac{d}{d-K}}}{d + K(d^{\frac{d}{d-K}} - 1)} \right) \\ &\quad - \frac{(d - K)}{d + K(d^{\frac{d}{d-K}} - 1)} \ln \left(\frac{1}{d + K(d^{\frac{d}{d-K}} - 1)} \right) \\ &\quad - \frac{d}{d + K(d^{\frac{d}{d-K}} - 1)} \ln d. \end{aligned} \quad (78)$$

While the expression for \mathcal{I}_K^* looks complicated, it can be greatly simplified. Denoting

$$\delta_K = \frac{d - K}{d^{\frac{d}{d-K}}},$$

one can rearrange terms in \mathcal{I}_K^* to obtain

$$\mathcal{I}_K^* = \ln(K + \delta_K). \quad (79)$$

From Eq. (72) and Eq. (78) and the fact that $\delta_K > 0$ for $0 < K < d$, we can infer that

$$\max_{\pi \in \Delta_d} \mathcal{I}_w(\pi) = \mathcal{I}_K^*.$$

Case 2: $K = 0$ or $K = d$

$K = 0$ implies $w_i = 0$ for all $i \in [d]$. Thus, $\langle \pi, \mathbf{w} \rangle = 0$ for all $\pi \in \Delta_d$, i.e., we will always lose the game. In that case, it can be verified that we must have $\mathcal{I}_w(\pi) = 0$ for all $\pi \in \Delta_d$.

On the other hand, $K = d$ implies $w_i = 1$ for all $i \in [d]$, meaning that the strategy is perfect. In particular, we have $\langle \pi, \mathbf{w} \rangle = 1$ for all $\pi \in \Delta_d$, i.e., $p_L = 0$ for any distribution on questions. It can be verified that $\mathcal{I}_w(\pi_U) = \ln d$ for the uniform distribution $\pi_U \in \Delta_d$. \square

Proof of Proposition 3: 1) Let $\mathbf{w}^{(D)} \in \{0, 1\}^d$ denote a binary deterministic strategy that can answer $d - 1$ of the d questions correctly. Then, there is exactly one index $k \in [d]$ such that $w_k^{(D)} = 0$. We show that we can achieve $\max_{\pi} \mathcal{I}_w(\pi)$ using the deterministic strategy $\mathcal{I}_{\mathbf{w}^{(D)}}(\tilde{\pi})$ for

some appropriate distribution $\tilde{\pi} \in \Delta_d$. Note that the winning vector $\mathbf{w}^{(D)}$ may not be permitted by the game, but it still gives an upper bound on $\max_{\pi} \mathcal{I}_w(\pi)$.

Since by assumption $K^* < d$, we can find an injective function $\sigma: \mathcal{K}^* \rightarrow [d] \setminus \{k\}$. The map σ will be used to label the indices. We construct the distribution $\tilde{\pi} \in \Delta_d$ as follows:

$$\begin{aligned} \tilde{\pi}_i &= w_{\sigma^{-1}(i)} \pi_{\sigma^{-1}(i)}^*, & \text{for } i \in \sigma(\mathcal{K}^*) \\ \tilde{\pi}_i &= 0, & \text{for } i \notin \sigma(\mathcal{K}^*) \setminus \{k\} \\ \tilde{\pi}_k &= 1 - \sum_{i \in \sigma(\mathcal{K}^*)} \tilde{\pi}_i. \end{aligned} \quad (80)$$

Since σ is an injective function on \mathcal{K}^* , it is invertible when its co-domain is restricted to its range $\sigma(\mathcal{K}^*)$. Thus, $\sigma^{-1}(i)$ is well-defined for $i \in \sigma(\mathcal{K}^*)$.

Since $w_j^{(D)} = 1$ for $j \in [d] \setminus \{k\}$ and $w_k^{(D)} = 0$, we have

$$\langle \mathbf{w}^{(D)}, \tilde{\pi} \rangle = \sum_{i \in \sigma(\mathcal{K}^*)} w_{\sigma^{-1}(i)} \pi_{\sigma^{-1}(i)}^* = \sum_{j \in \mathcal{K}^*} w_j \pi_j^* = \langle \mathbf{w}, \pi^* \rangle. \quad (81)$$

That is, the winning probabilities obtained from \mathbf{w} , π^* and $\mathbf{w}^{(D)}$, $\tilde{\pi}$ are the same.

Next, we will look at the output probability of the channel in the two cases. For \mathbf{w} , π^* , the output probability is given as (see Eq. (27))

$$p_i^* = \begin{cases} w_i \pi_i^* + \frac{1}{d} - \frac{1}{d} \sum_{j \in \mathcal{K}^*} w_j \pi_j^* & \text{for } i \in \mathcal{K}^* \\ \frac{1}{d} - \frac{1}{d} \sum_{j \in \mathcal{K}^*} w_j \pi_j^* & \text{for } i \notin \mathcal{K}^*, \end{cases}$$

whereas for $\mathbf{w}^{(D)}$, $\tilde{\pi}$, the output probability is given as

$$\tilde{p}_i = \begin{cases} w_{\sigma^{-1}(i)} \pi_{\sigma^{-1}(i)}^* + \frac{1}{d} - \frac{1}{d} \sum_{j \in \mathcal{K}^*} w_j \pi_j^* & \text{for } i \in \sigma(\mathcal{K}^*) \\ \frac{1}{d} - \frac{1}{d} \sum_{j \in \mathcal{K}^*} w_j \pi_j^* & \text{for } i \notin \sigma(\mathcal{K}^*). \end{cases}$$

Here we used Eq. (81) to obtain the expression for \tilde{p} . Since $|\mathcal{K}^*| = |\sigma(\mathcal{K}^*)|$, the probability \tilde{p} is just a permutation of p^* . Furthermore, because Shannon entropy is invariant under permutation of the entries of the probability distribution, we have $H(\tilde{p}) = H(p^*)$. Then from Eq. (27), Eq. (29), and Eq. (81), we have

$$\max_{\pi \in \Delta_d} \mathcal{I}_w(\pi) = \mathcal{I}_w(\pi^*) = \mathcal{I}_{\mathbf{w}^{(D)}}(\tilde{\pi}) \leq \max_{\pi \in \Delta_d} \mathcal{I}_{\mathbf{w}^{(D)}}(\pi) = \mathcal{I}_{d-1}^*$$

where \mathcal{I}_{d-1}^* is given by Eq.(33).

2) We focus on the case where $K^* = K = d$. This implies $w_i \neq 0$ and $\pi_i^* \neq 0$ for all $i \in [d]$. Our goal is to solve the optimization problem $\max_{\pi \in \Delta_d} \mathcal{I}_w(\pi)$. Following Eq. (69), we write the Lagrangian for the problem

$$\begin{aligned} \mathcal{L}(\pi; \lambda, \nu) &= \sum_{i=1}^d p_i \ln p_i - \sum_{i=1}^d \pi_i w_i \ln d + \ln d \\ &\quad - \sum_{i=1}^d \lambda_i \pi_i + \nu \left(\sum_{i=1}^d \pi_i - 1 \right), \end{aligned} \quad (82)$$

where

$$p_i = \pi_i w_i + \frac{1}{d} - \frac{1}{d} \sum_{j=1}^d \pi_j w_j. \quad (83)$$

Since $K^* = d$, the probabilities p_i are non-zero at the optimum, and hence we can differentiate the entropy at the optimum.

Note that we are solving a convex optimization problem and Slater's condition holds [40]. Therefore, KKT conditions are necessary and sufficient for optimality [40]. From KKT conditions, it follows that $\nabla_{\pi} \mathcal{L} = 0$, which gives

$$\ln p_j - \frac{1}{d} \sum_{i=1}^d \ln p_i = \frac{\lambda_j}{w_j} - \frac{\nu}{w_j} + \ln d, \quad j \in [d].$$

Since by assumption $\pi_i \neq 0$ at the optimum for all $i \in [d]$, we have $\lambda_i = 0$ by complementary slackness. This gives

$$p_j = d \left(\prod_{i=1}^d p_i \right)^{1/d} \exp \left(-\frac{\nu}{w_j} \right) \quad (84)$$

for $j \in [d]$. By multiplying Eq. (84) for $j = 1, \dots, d$, we can infer that

$$\nu = dw_{\text{eff}} \ln d \quad (85)$$

since $p_i \neq 0$ at the optimum. Next, summing over $j \in [d]$ in Eq. (84), we obtain

$$d \left(\prod_{i=1}^d p_i \right)^{1/d} = \frac{1}{\sum_{j=1}^d \exp \left(-\frac{\nu}{w_j} \right)}. \quad (86)$$

Combining this with Eq. (84) and Eq. (85), we obtain

$$p_i = \frac{\exp \left(-\frac{dw_{\text{eff}} \ln d}{w_i} \right)}{\sum_{j=1}^d \exp \left(-\frac{dw_{\text{eff}} \ln d}{w_j} \right)}. \quad (87)$$

Therefore, at the optimum, we have a Boltzmann distribution for the probability of outcomes of the channel.

Now, instead of solving for π from Eq. (87), we find a suitable expression for the objective function $\mathcal{I}_w(\pi)$ in terms of p and w (instead of π and w). To this end, subtracting Eq. (83) corresponding to two different indices i and j and summing over the index i , we obtain

$$\sum_{i=1}^d \pi_i w_i = 1 + d(\pi_j w_j - p_j).$$

Dividing this equation by w_j and summing over j , we get

$$\begin{aligned} \sum_{j=1}^d \frac{1}{w_j} \sum_{i=1}^d \pi_i w_i &= \sum_{j=1}^d \frac{1}{w_j} + d \left(1 - \sum_{j=1}^d \frac{p_j}{w_j} \right) \\ \sum_{i=1}^d \pi_i w_i &= 1 - dw_{\text{eff}} \left(\sum_{j=1}^d \frac{p_j}{w_j} - 1 \right), \end{aligned}$$

where w_{eff} is as defined in Eq. (35). Substituting this in Eq. (29), we find that the mutual information $\mathcal{I}_w(\pi)$ can be written as a function of the outcome probability p as follows:

$$\mathcal{I}_w(\pi) = H(p) - dw_{\text{eff}} \left(\sum_{j=1}^d \frac{p_j}{w_j} - 1 \right) \ln d. \quad (88)$$

To obtain the value of $\mathcal{I}_w(\pi)$ at the optimum, we substitute Eq. (87) in Eq. (88). Denoting $\max_{\pi \in \Delta_d} \mathcal{I}_w(\pi) = \mathcal{I}^*(w)$, one can rearrange terms to obtain

$$\mathcal{I}^*(w) = dw_{\text{eff}} \ln d + \ln \left(\sum_{j=1}^d \exp \left(-\frac{dw_{\text{eff}} \ln d}{w_j} \right) \right) \quad (89)$$

$$= \ln \left(\sum_{j=1}^d \exp \left[dw_{\text{eff}} \ln d \left(1 - \frac{1}{w_j} \right) \right] \right), \quad (90)$$

where we get the last equation by noting that $dw_{\text{eff}} \ln d = \ln \exp(dw_{\text{eff}} \ln d)$. \square

Proof of Proposition 4: To obtain an upper bound on

$$\sup_{w \in \overline{\mathcal{W}}_C, w > 0} \mathcal{I}^*(w)$$

given in Eq. (38), we solve a relaxation of this optimization problem. For this purpose, note that the set over which we optimize can be written as

$$\overline{\mathcal{W}}_C > = \left\{ w \in [0, 1]^d \mid w > 0, \frac{\sum_{i=1}^d w_i}{d} \leq \omega^C(G) \right\}.$$

Let $\text{HM}(w_1, \dots, w_d)$ be the harmonic mean and $\text{AM}(w_1, \dots, w_d)$ be the arithmetic mean of w_1, \dots, w_d , respectively. Observe that

$$dw_{\text{eff}} = \text{HM}(w_1, \dots, w_d).$$

Then, since $\text{HM}(w_1, \dots, w_d) \leq \text{AM}(w_1, \dots, w_d)$ for $w_1, \dots, w_d > 0$ (with equality iff $w_1 = \dots = w_d$), we have

$$dw_{\text{eff}} \leq \frac{\sum_{i=1}^d w_i}{d} \leq \omega^C(G). \quad (91)$$

With this in mind, we define the set

$$\overline{\mathcal{W}}_{\text{eff}} = \{ w \in [0, 1]^d \mid w > 0, dw_{\text{eff}} \leq \omega^C(G) \}.$$

Then, from Eq. (91), it follows that $\overline{\mathcal{W}}_C > \subseteq \overline{\mathcal{W}}_{\text{eff}}$. Therefore, we have

$$\sup_{w \in \overline{\mathcal{W}}_C >} \mathcal{I}^*(w) \leq \sup_{w \in \overline{\mathcal{W}}_{\text{eff}}} \mathcal{I}^*(w).$$

Subsequently, we solve the optimization problem $\sup_{w \in \overline{\mathcal{W}}_{\text{eff}}} \mathcal{I}^*(w)$. To that end, we make the change of variables

$$t_i = \frac{1}{w_i}$$

for $i \in [d]$. The optimization problem then becomes

$$\begin{aligned} \sup_{\mathbf{t}} \quad & \ln \left(\sum_{i=1}^d \exp \left[\frac{d}{\sum_{j=1}^d t_j} \ln d \ (1 - t_i) \right] \right) \\ \text{s.t.} \quad & \mathbf{t} \geq 1 \\ & \frac{\sum_{i=1}^d t_i}{d} \geq \frac{1}{\omega^C(G)}. \end{aligned} \quad (92)$$

We solve Eq. (92) by splitting it into two maximizations. To that end, define

$$s = \frac{\sum_{i=1}^d t_i}{d}, \quad (93)$$

so that the objective of the optimization in Eq. (92) can be written as

$$f(\mathbf{t}, s) = f(\mathbf{t}, s) = \text{LSE} \left(-\frac{\ln d}{s} t_1, \dots, -\frac{\ln d}{s} t_d \right) + \frac{\ln d}{s}, \quad (94)$$

where $\text{LSE}(y_1, \dots, y_d) = \ln(\sum_{j=1}^d e^{y_j})$ is the log-sum-exp function. Since LSE is a convex function, we can conclude that $f(\mathbf{t}, s)$ is convex in \mathbf{t} for a fixed s . Then, Eq. (92) can be expressed as

$$\sup_{\mathbf{t} \in \overline{\mathcal{W}}_{\mathcal{C}_{\text{eff}}}^{(s)}} f(\mathbf{t}) = \sup_{s \geq (\omega^{\mathcal{C}}(G))^{-1}} \max_{\mathbf{t} \in \overline{\mathcal{W}}_{\mathcal{C}_{\text{eff}}}^{(s)}} f(\mathbf{t}, s),$$

where

$$\overline{\mathcal{W}}_{\mathcal{C}_{\text{eff}}}^{(s)} = \left\{ \mathbf{t} \in \mathbb{R}^d \mid \mathbf{t} \geq 1, \frac{\sum_{i=1}^d t_i}{d} = s \right\}$$

is a translated and scaled simplex. The extreme points of $\overline{\mathcal{W}}_{\mathcal{C}_{\text{eff}}}^{(s)}$ are given as

$$\mathbf{t}_s^{(1)} = (1, \dots, 1, 1 + (s-1)d)^T, \quad (95)$$

and its permutations are denoted by $\mathbf{t}_s^{(2)}, \dots, \mathbf{t}_s^{(d)}$. Note that the constraints in Eq. (92) imply $s \geq 1$.

Now, we are seeking to maximize the convex function $f(\mathbf{t}, s)$ over the set $\overline{\mathcal{W}}_{\mathcal{C}_{\text{eff}}}^{(s)}$ (for a fixed s). However, since f is convex and any point $\mathbf{t} \in \overline{\mathcal{W}}_{\mathcal{C}_{\text{eff}}}^{(s)}$ can be written as $\mathbf{t} = \sum_{i=1}^d \lambda_i \mathbf{t}_s^{(i)}$, we have

$$f(\mathbf{t}, s) \leq \sum_{i=1}^d \lambda_i f(\mathbf{t}_s^{(i)}, s).$$

Then, because $f(\mathbf{t}, s)$ is invariant under the permutation of the components of \mathbf{t} , we have $f(\mathbf{t}_s^{(i)}, s) = f(\mathbf{t}_s^{(j)}, s)$ for any $i, j \in [d]$. Therefore,

$$\begin{aligned} \max_{\mathbf{t} \in \overline{\mathcal{W}}_{\mathcal{C}_{\text{eff}}}^{(s)}} f(\mathbf{t}, s) &= f(\mathbf{t}_s^{(1)}, s) \\ &= \ln \left(d - 1 + \exp \left[-d \ln d \left(\frac{s-1}{s} \right) \right] \right) \end{aligned}$$

where we substituted Eq. (95) in Eq. (94) in the last equation. Then, since $f(\mathbf{t}^{(1)}, s)$ is a decreasing function of s , we can infer that

$$\begin{aligned} \sup_{s \geq (\omega^{\mathcal{C}}(G))^{-1}} \max_{\mathbf{t} \in \overline{\mathcal{W}}_{\mathcal{C}_{\text{eff}}}^{(s)}} f(\mathbf{t}, s) &= f \left(\mathbf{t}_s^{(1)}, \frac{1}{\omega^{\mathcal{C}}(G)} \right) \\ &= \ln \left(d - 1 + d^{-(1-\omega^{\mathcal{C}}(G))d} \right), \end{aligned}$$

giving the desired bound. \square

APPENDIX B

SOME NOTES ON STRATEGIES FOR NONLOCAL GAMES

We begin by characterizing the extreme points of the set of conditional distributions (over finite sets).

Proposition 13: Let \mathcal{X} and \mathcal{Y} be finite sets, and let \mathcal{C} denote the set of conditional probability distributions on \mathcal{Y} given \mathcal{X} . Then, the extreme points of \mathcal{C} correspond to conditional probabilities $p_{\mathbf{Y}|\mathbf{X}}^f(\mathbf{y}|\mathbf{x}) = \delta_{\mathbf{y}, f(\mathbf{x})}$ obtained from functions

$f: \mathcal{X} \rightarrow \mathcal{Y}$. In particular, any conditional probability distribution $p_{\mathbf{Y}|\mathbf{X}} \in \mathcal{C}$ can be written as a convex combination of $p_{\mathbf{Y}|\mathbf{X}}^f$ corresponding to functions f .

Proof: For convenience, let us denote $|\mathcal{X}| = m$ and $|\mathcal{Y}| = n$ for some $m, n \in \mathbb{N}$. We also fix a labelling of elements $\mathcal{X} = \{x_1, \dots, x_m\}$ and $\mathcal{Y} = \{y_1, \dots, y_n\}$. The set of conditional probability distributions \mathcal{C} can be thought of as the set of functions from $\mathcal{X} \rightarrow \Delta_n$. That is, for each $\mathbf{x} \in \mathcal{X}$, we have a probability distribution over \mathcal{Y} . Thus, we identify \mathcal{C} with the set $(\Delta_n)^m$.

Since Δ_n is a simplex in \mathbb{R}^n , its extreme points are $\mathbf{e}_1, \dots, \mathbf{e}_n$, where \mathbf{e}_i is the standard Euclidean basis vector for the i th coordinate. Therefore, the extreme points of $\mathcal{C} = (\Delta_n)^m$ are $(\mathbf{e}_{i_1}, \dots, \mathbf{e}_{i_m})$ for $i_1, \dots, i_m \in [n]$. We argue that such extreme points can be obtained from functions $f: \mathcal{X} \rightarrow \mathcal{Y}$.

To that end, given an extreme point $(\mathbf{e}_{i_1}, \dots, \mathbf{e}_{i_m})$ of \mathcal{C} , construct the function

$$f(\mathbf{x}_j) = \mathbf{y}_{i_j} \quad j \in [m].$$

Observe that the conditional probability distribution determined by the function f is

$$p_{\mathbf{Y}|\mathbf{X}}^f(\cdot, \mathbf{x}_j) = \delta_{\cdot, f(\mathbf{x}_j)} = \delta_{\cdot, \mathbf{y}_{i_j}} = \mathbf{e}_{i_j}$$

for $j \in [m]$. Since there are a total of $|\mathcal{Y}|^{|\mathcal{X}|}$ extreme points of \mathcal{C} , and just as many functions from \mathcal{X} to \mathcal{Y} , the above construction gives a bijective mapping between the extreme points and these functions. Since \mathcal{C} is a compact and convex set with a finite number of extreme points, it is generated as the convex hull of its extreme points by the Krein-Milman theorem [44]. \square

In the following proposition, we show that the set of no-signalling strategies is a compact and convex set. The proof of this proposition is constructive, and it can therefore be used to construct the set of no-signalling distributions numerically.

Proposition 14: Let $\mathcal{X}_1, \dots, \mathcal{X}_N$ denote the question set for N players, and let $\mathcal{Y}_1, \dots, \mathcal{Y}_N$ denote the answer set. Let \mathfrak{S}_{NS} denote the set of no-signalling strategies used by the players. That is, $p_{\mathbf{Y}|\mathbf{X}} \in \mathfrak{S}_{\text{NS}}$ iff

$$p_{Y_i|\mathbf{X}}(y_i|x_1, \dots, x_N) = p_{Y_i|\mathbf{X}_i}(y_i|x_i) \quad \forall x_k \in \mathcal{X}_k, k \in [N] \setminus \{i\} \quad (96)$$

for all $y_i \in \mathcal{Y}_i$, $x_i \in \mathcal{X}_i$, $i \in [N]$. Then, \mathfrak{S}_{NS} is a compact and convex set. Specifically, \mathfrak{S}_{NS} is a convex polytope obtained as the intersection of hyperplanes and halfspaces.

Proof: For convenience, we denote $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_N$ as the question set and $\mathcal{Y} = \mathcal{Y}_1 \times \dots \times \mathcal{Y}_N$ as the answer set. The set of all strategies $p_{\mathbf{Y}|\mathbf{X}}(y_1, \dots, y_N|x_1, \dots, x_N)$ can be written as a product of simplices $(\Delta_{|\mathcal{Y}|})^{|\mathcal{X}|}$ (see Prop. 13), and is therefore a compact set. Also note that we can write the no-signalling condition given in Eq. (96) as

$$\begin{aligned} &\sum_{\substack{y_j \in \mathcal{Y}_j \\ j \neq i}} p_{\mathbf{Y}|\mathbf{X}}(y_1, \dots, y_i, \dots, y_N|x_1, \dots, x_i, \dots, x_N) \\ &= \sum_{\substack{y_j \in \mathcal{Y}_j \\ j \neq i}} p_{\mathbf{Y}|\mathbf{X}}(y_1, \dots, y_i, \dots, y_N|x'_1, \dots, x_i, \dots, x'_N) \end{aligned}$$

$$\forall x_k, x'_k \in \mathcal{X}_k, k \in [N] \setminus \{i\} \quad (97)$$

for all $y_i \in \mathcal{Y}_i$, $x_i \in \mathcal{X}_i$, $i \in [N]$. In the above equation, the sum over $y_j \in \mathcal{Y}_j$ is a shorthand for the sum over $y_1 \in \mathcal{Y}_1, \dots, y_N \in \mathcal{Y}_N$.

We will show that the set of no-signalling strategies \mathfrak{S}_{NS} is a closed set. First, note that $\mathfrak{S}_{\text{NS}} \subseteq (\Delta_{|\mathcal{Y}|})^{|\mathcal{X}|}$, i.e., the set of no-signalling strategies is contained in the set of all strategies. Thus, we can write the elements of \mathfrak{S}_{NS} as vectors $v = (v^{(1)}, \dots, v^{(|\mathcal{X}|)})$, where $v^{(i)} \in \Delta_{|\mathcal{Y}|}$ is a probability vector. Note that $v^{(i)}$ is a $|\mathcal{Y}|$ -dimensional vector whereas v is a $|\mathcal{X}||\mathcal{Y}|$ -dimensional vector. Essentially, the vector $v^{(i)}$ denotes a probability distribution $p_{\mathbf{Y}|\mathbf{X}}(y|x^{(i)})$ over \mathcal{Y} for a fixed $x^{(i)} \in \mathcal{X}$. Using this vectorial representation of a strategy, we will write the no-signalling condition given in Eq. (97) in matrix form.

To that end, fix an ordering for the elements of \mathcal{Y} . Then, we can index the elements of $v^{(i)}$ as $v_{(y_1, \dots, y_N)}^{(i)}$ for $(y_1, \dots, y_N) \in \mathcal{Y}$ corresponding to that ordering. For each $y_i \in \mathcal{Y}_i$ ($i \in [d]$), let s_{y_i} denote the $|\mathcal{Y}|$ -dimensional vector with 1 at each index $(y'_1, \dots, y'_N) \in \mathcal{Y}$ with $y'_i = y_i$ and 0 elsewhere. For example, if $\mathcal{Y}_1 = \{a, b\}$ and $\mathcal{Y}_2 = \{c, d\}$ and we write the elements of $\mathcal{Y} = \{(a, c), (a, d), (b, c), (b, d)\}$ in that order, then, $s_a = (1, 1, 0, 0)^T$, $s_c = (1, 0, 1, 0)^T$, and so forth. Observe also that, for each $k \in [|\mathcal{X}|]$,

$$s_{y_i}^T v^{(k)} = \sum_{\substack{y_j \in \mathcal{Y}_j \\ j \neq i}} p_{\mathbf{Y}|\mathbf{X}}(y_1, \dots, y_i, \dots, y_N | x^{(k)}).$$

Similarly, fix an ordering for \mathcal{X} . For each $x_i \in \mathcal{X}_i$ with $i \in [N]$, let \mathcal{I}_{x_i} be the set that contains the indices $(x'_1, \dots, x'_N) \in \mathcal{X}$ with $x'_i = x_i$. Then, define the $|\mathcal{I}_{x_i}| \times |\mathcal{X}||\mathcal{Y}|$ matrix $S_{(x_i, y_i)}$ in block form as follows. Imagine each row of $S_{(x_i, y_i)}$ being split into $|\mathcal{X}|$ blocks of $|\mathcal{Y}|$ -dimensional vectors, i.e., $(b_1 \ \dots \ b_{|\mathcal{X}|})$ with b_i a $|\mathcal{Y}|$ -dimensional (row) vector. We label the rows of $S_{(x_i, y_i)}$ with \mathcal{I}_{x_i} . Define the row $k \in \mathcal{I}_{x_i}$ to be the block $(0_{1 \times |\mathcal{Y}|} \ \dots \ s_{y_i}^T \ \dots \ 0_{1 \times |\mathcal{Y}|})$ with $s_{y_i}^T$ in the k -th block.

Let us label the blocks $v^{(i)}$ of the vector $v = (v^{(1)}, \dots, v^{(|\mathcal{X}|)}) \in (\Delta_{|\mathcal{Y}|})^{|\mathcal{X}|}$ as $v^{(x_1, \dots, x_N)}$ for $(x_1, \dots, x_N) \in \mathcal{X}$ using the ordering of \mathcal{X} that we have fixed. Then, we have $(S_{(x_i, y_i)} v)_{(x_1, \dots, x_N)} = s_{y_i}^T v^{(x_1, \dots, x_N)}$ for $(x_1, \dots, x_N) \in \mathcal{I}_{x_i}$. That is, $S_{(x_i, y_i)} v$ is a $|\mathcal{I}_{x_i}|$ -dimensional vector with the entries

$$(S_{(x_i, y_i)} v)_{(x_1, \dots, x_N)} = \sum_{\substack{y_j \in \mathcal{Y}_j \\ j \neq i}} p_{\mathbf{Y}|\mathbf{X}}(y_1, \dots, y_i, \dots, y_N | x_1, \dots, x_i, \dots, x_N)$$

for $(x_1, \dots, x_N) \in \mathcal{I}_{x_i}$. Since the i -th component of $(x_1, \dots, x_N) \in \mathcal{I}_{x_i}$ is fixed to be x_i , the no-signalling condition given in Eq. (97) says that all the components of $S_{(x_i, y_i)} v$ are equal.

Therefore, we enforce the no-signalling condition as follows. Define a $(|\mathcal{I}_{x_i}| - 1) \times |\mathcal{I}_{x_i}|$ matrix

$$D_{x_i} = \begin{pmatrix} 1 & -1 & 0 & \dots & 0 & 0 \\ 0 & 1 & -1 & \dots & 0 & 0 \\ & & & \ddots & & \\ 0 & 0 & 0 & \dots & 1 & -1 \end{pmatrix},$$

and observe that if $r = (r_1 \ \dots \ r_{|\mathcal{I}_{x_i}|})^T$ is any $|\mathcal{I}_{x_i}|$ -dimensional vector, then $D_{x_i} r$ is the $(|\mathcal{I}_{x_i}| - 1)$ -dimensional vector $(r_1 - r_2 \ \dots \ r_{|\mathcal{I}_{x_i}| - 1} - r_{|\mathcal{I}_{x_i}|})^T$. Then, by the preceding remarks, the no-signalling condition can be written as

$$D_{x_i} S_{(x_i, y_i)} v = 0 \quad \text{for all } x_i \in \mathcal{X}_i, y_i \in \mathcal{Y}_i, i \in [N]. \quad (98)$$

Since $D_{x_i} S_{(x_i, y_i)}$ is a $(|\mathcal{I}_{x_i}| - 1) \times |\mathcal{X}||\mathcal{Y}|$ matrix and v is a $|\mathcal{X}||\mathcal{Y}|$ -dimensional vector, the equation $D_{x_i} S_{(x_i, y_i)} v = 0$ encodes $(|\mathcal{I}_{x_i}| - 1)$ hyperplanes (thinking of v as an arbitrary $|\mathcal{X}||\mathcal{Y}|$ -dimensional vector). Therefore, the set of no-signalling strategies can be written as

$$\mathfrak{S}_{\text{NS}} = \{v \in (\Delta_{|\mathcal{Y}|})^{|\mathcal{X}|} \mid D_{x_i} S_{(x_i, y_i)} v = 0 \quad \forall x_i \in \mathcal{X}_i, y_i \in \mathcal{Y}_i, i \in [N]\}. \quad (99)$$

This is the intersection of the compact set $(\Delta_{|\mathcal{Y}|})^{|\mathcal{X}|}$ with the hyperplanes defined by $D_{x_i} S_{(x_i, y_i)} v = 0$. Since hyperplanes are closed, the set \mathfrak{S}_{NS} is closed as well.

Since $(\Delta_{|\mathcal{Y}|})^{|\mathcal{X}|}$ is bounded, \mathfrak{S}_{NS} is also bounded, so that \mathfrak{S}_{NS} is compact. The convexity of \mathfrak{S}_{NS} follows from the fact that $(\Delta_{|\mathcal{Y}|})^{|\mathcal{X}|}$ and hyperplanes are convex. \square

APPENDIX C ANALYSIS OF ALGORITHMS FOR OPTIMIZATION OF LIPSCHITZ-LIKE FUNCTION

In this section, we present convergence analysis for algorithms used for Lipschitz-like optimization. We also present an algorithm for performing optimization of Lipschitz-like functions over an arbitrary compact & convex domain.

A. Optimizing Lipschitz-Like Functions Over an Interval Using Modified Piyavskii-Shubert Algorithm

We first present a convergence analysis for modified Piyavskii-Shubert algorithm that finds the global maximum of Lipschitz-like functions over a closed interval.

Proposition 15: Let $\beta: \mathbb{R}_+ \rightarrow \mathbb{R}$ be a non-negative, continuous, monotonically increasing function with $\beta(0) = 0$. Let $\mathcal{D} = [a, b]$ be a closed interval, where $a, b \in \mathbb{R}$. Let $f: \mathcal{D} \rightarrow \mathbb{R}$ be a real-valued function satisfying

$$|f(q) - f(q')| \leq \beta(|q - q'|)$$

Then, for each choice of tolerance $\epsilon > 0$, Alg. 1 terminates in a finite number of time steps. If $f(q^*)$ denotes the output of Alg. 1 corresponding to a tolerance of $\epsilon > 0$, we have $\max_{q \in \mathcal{D}} f(q) - f(q^*) \leq \epsilon$. The number of time steps required to converge with a tolerance of $\epsilon > 0$ is bounded above by $\lceil (b - a)/\delta \rceil$, where $\delta = \sup\{\delta' > 0 \mid \beta(x) < \epsilon/2 \ \forall 0 \leq x \leq \delta'\}$.

Proof: Let $i \in \mathbb{N}$ be a natural number. Given any $q^{(i)} \in [a, b]$, define the function

$$F_i(q) = f(q^{(i)}) + \beta(|q - q^{(i)}|) \quad (100)$$

for $q \in [a, b]$. Since $f(q^{(i)}) - f(q) \geq -\beta(|q - q^{(i)}|)$ by assumption, we have $f(q) \leq F_i(q)$ for all $q \in [a, b]$. Given that we initialize $q^{(0)} = a$, the maximum of $F_0(q)$ occurs at b because β is a monotonically increasing function with $\beta(0) = 0$. This justifies the assignment $q^{(1)} = b$ at the start of the algorithm.

Now, suppose that Alg. 1 terminates at the K th time step. Then, the algorithm has generated points $q^{(0)}, \dots, q^{(K)} \in [a, b]$ and a point $q^* \in [a, b]$ satisfying $F(q^*) - f(q^*) \leq \epsilon$. The function F and the point q^* are obtained as follows. The points $q^{(0)}, \dots, q^{(K)}$ are sorted at the beginning of the K th iteration, so that $a = q^{(0)} \leq q^{(1)} \leq \dots \leq q^{(K)} = b$. Then, points $\bar{q}^{(i)} \in \operatorname{argmax}_{q \in [q^{(i)}, q^{(i+1)}]} \min\{F_i(q), F_{i+1}(q)\}$ are obtained through root finding for $0 \leq i \leq K-1$. We choose $m \in \operatorname{argmax}_{0 \leq i \leq K-1} F_i(\bar{q}^{(i)})$ and set $q^* = \bar{q}^{(m)}$ and $F = F_m$.

Next, we elaborate on how $\bar{q}^{(i)} \in \operatorname{argmax}_{q \in [q^{(i)}, q^{(i+1)}]} \min\{F_i(q), F_{i+1}(q)\}$ is computed using root finding. First, consider the function $g_i(q) = F_i(q) - F_{i+1}(q)$. From the non-negativity and monotonicity of β , it follows that $F_i(q)$ is a monotonically increasing function in the interval $q \in [q^{(i)}, q^{(i+1)}]$, whereas $F_{i+1}(q)$ is a monotonically decreasing function in the interval $q \in [q^{(i)}, q^{(i+1)}]$. Therefore, $g_i(q)$ is a continuous and monotonically increasing function in the interval $q \in [q^{(i)}, q^{(i+1)}]$, where the continuity of g_i follows from that of F_i and F_{i+1} . Since $\beta(0) = 0$ and $f \leq F_i$ for all $0 \leq i \leq K$, we have $g_i(q^{(i)}) \leq 0$ and $g_i(q^{(i+1)}) \geq 0$. Therefore, the function g_i has a root $\bar{q}^{(i)}$ in the interval $[q^{(i)}, q^{(i+1)}]$, and since $g_i = F_i - F_{i+1}$, we have $F_i(\bar{q}^{(i)}) = F_{i+1}(\bar{q}^{(i)})$. From the monotonicity properties of F_i and F_{i+1} , we can infer that $\bar{q}^{(i)}$ maximizes the bounding function $\min\{F_i, F_{i+1}\}$ in the interval $[q^{(i)}, q^{(i+1)}]$, and this maximum value is equal to $F_i(\bar{q}^{(i)}) = F_{i+1}(\bar{q}^{(i)})$.

Then, because $m \in \operatorname{argmax}_{0 \leq i \leq K-1} F_i(\bar{q}^{(i)})$, $q^* = \bar{q}^{(m)}$ and $F = F_m$, we can infer that for all $q \in [q^{(i)}, q^{(i+1)}]$, we have

$$f(q) \leq \min\{F_i(q), F_{i+1}(q)\} \leq F_i(\bar{q}^{(i)}) \leq F(q^*) \leq f(q^*) + \epsilon$$

Since the above equation holds for every $0 \leq i \leq K-1$, and the intervals $[q^{(0)}, q^{(1)}], \dots, [q^{(K-1)}, q^{(K)}]$ cover $[a, b]$, we can infer that

$$\max_{q \in [a, b]} f(q) \leq f(q^*) + \epsilon$$

It remains to show that Alg. 1 terminates in a finite number of time steps. To that end, note that $\beta(x)$ is a continuous function of $x \in \mathbb{R}_+$ with $\beta(0) = 0$. Therefore, for any given $\epsilon > 0$, we can find a $\delta > 0$ such that $\beta(x) \leq \epsilon/2$ whenever $0 \leq x \leq \delta$. Let $K \in \mathbb{N}$ denote the current time step. For $0 \leq i \leq K-1$, let $\bar{q}^{(i)}$ denote a root of $F_i - F_{i+1}$ in the interval $[q^{(i)}, q^{(i+1)}]$. Let $m \in \operatorname{argmax}_{0 \leq i \leq K-1} F_i(\bar{q}^{(i)})$, $q^* = \bar{q}^{(m)}$ and $F = F_m$ as before. Since $f(q^{(m)}) - f(q^*) \leq \beta(|q^* - q^{(m)}|)$, we have $F(q^*) - f(q^*) \leq 2\beta(|q^* - q^{(m)}|)$. Therefore,

when $|q^* - q^{(m)}| \leq \delta$, we have $F(q^*) - f(q^*) \leq \epsilon$. Since $F(q^*) = F_m(q^*) = F_{m+1}(q^*)$, we can similarly infer that $F(q^*) - f(q^*) \leq \epsilon$ whenever $|q^* - q^{(m+1)}| \leq \delta$. Therefore, the algorithm terminates if either $|q^* - q^{(m)}| \leq \delta$ or $|q^* - q^{(m+1)}| \leq \delta$, where $q^* \in [q^{(m)}, q^{(m+1)}]$.

As per the procedure outlined in Alg. 1, the point q^* will join the iterates $q^{(0)}, \dots, q^{(K)}$ at the $(K+1)$ th time step. When this new iterate q^* is added, the updated intervals include $[q^{(m)}, q^*]$ and $[q^*, q^{(m+1)}]$. Since $q^* \in [q^{(m)}, q^{(m+1)}]$, we either have $|q^* - q^{(m)}| \leq |q^{(m+1)} - q^{(m)}|/2$ or $|q^* - q^{(m+1)}| \leq |q^{(m+1)} - q^{(m)}|/2$. Since there are only a finite number of intervals at each time step and the length of one of the sides of the interval where the new iterate falls is at least halved at each time step, at some large enough time step K , we will have $|q^* - q^{(m)}| \leq \delta$ or $|q^* - q^{(m+1)}| \leq \delta$. Thus, the algorithm terminates in a finite number of time steps.

The worst case scenario corresponds to the situation where $|q^{(i)} - q^{(i+1)}| = \delta$ for all $0 \leq i \leq K-1$. In this case, the algorithm terminates at the K th time step, wherein q^* falls within one of these intervals. Taking $\delta = \sup\{\delta' > 0 \mid \beta(\delta') \leq \epsilon/2\}$, there is a sequence $\delta'_n \rightarrow \delta$ with $\beta(\delta'_n) \leq \epsilon/2 \forall n$, so that by continuity of β , we have $\beta(\delta) = \lim_{n \rightarrow \infty} \beta(\delta'_n) \leq \epsilon/2$. Therefore, the number of time steps required for the algorithm to terminate is bounded above by the number $\lceil (b-a)/\delta \rceil$, where $\delta > 0$ can be taken as the largest number that satisfies $\beta(x) \leq \epsilon/2$ for $0 \leq x \leq \delta$. \square

Next, we present details of constructing and searching over the grid for Lipschitz-like optimization.

B. Lipschitz-Like Optimization Over the Standard Simplex Using Grid Search

We make the following observations about the integer grid defined in Eq. (50). For the grid $\Delta_{d,N}$ defined in Eq. (52) over the standard simplex, we have $\Delta_{d,N} = \mathcal{I}_{d,N}/N := \{n/N \mid n \in \mathcal{I}_{d,N}\}$. Therefore, the observations noted below also apply to $\Delta_{d,N}$ with appropriate modifications.

Proposition 16: Given $d, N \in \mathbb{N}_+$, let $\mathcal{I}_{d,N}$ denote the integer grid defined in Eq. (50). Then the following hold.

- 1) Any element $n \in \mathcal{I}_{d,N}$ can be written as $n = (N - \ell_{d-1}, \ell_{d-1} - \ell_{d-2}, \dots, \ell_2 - \ell_1, \ell_1)$ for some integers $0 \leq \ell_i \leq \ell_{i+1} \leq N$, $i \in [d-2]$.
- 2) The elements of $\mathcal{I}_{d,N}$ can be ordered such that any two consecutive elements are distance 2 apart in l_1 -norm. This ordering is constructive and can be implemented algorithmically.
- 3) The grid $\Delta_{d,N}$ defined in Eq. (52) is a $(2(d-1)/N)$ -net in the l_1 -norm for the standard simplex in dimension d . That is, given any $x \in \Delta_d$, there is some $z \in \Delta_{d,N}$ such that $\|x - z\|_1 \leq 2(d-1)/N$.

Proof: 1) We prove this statement by induction on the dimension. For $d = 2$, any $n \in \mathcal{I}_{d,M}$ satisfies $n_1 + n_2 = M$, and therefore, $n = (M - n_1, n_1)$ holds for all $M \in \mathbb{N}_+$. Now, assume that for any $M \in \mathbb{N}_+$, we can write each $m \in \mathcal{I}_{d-1,M}$ in dimension $d-1$ as $m = (M - s_{d-2}, \dots, s_2 - s_1, s_1)$, where $s_1, \dots, s_{d-2} \in \mathbb{N}$ satisfy $0 \leq s_i \leq s_{i+1} \leq M$ for $i \in [d-3]$. Then, given any $n \in \mathcal{I}_{d,N}$, we can write $n_d + \sum_{i=1}^{d-1} n_i = N$. Denote $\ell_{d-1} = \sum_{i=1}^{d-1} n_i$, so that $n_d = N - \ell_{d-1}$ and

$0 \leq \ell_{d-1} \leq N$. Then, since $\sum_{i=1}^{d-1} = \ell_{d-1}$, we can find some numbers $\ell_1, \dots, \ell_{d-2}$ such that $0 \leq \ell_i \leq \ell_{i+1} \leq \ell_{d-1} \leq N$ for $i \in [d-3]$ and $(n_1, \dots, n_{d-1}) = (\ell_{d-1} - \ell_{d-2}, \dots, \ell_2 - \ell_1, \ell_1)$ by assumption. Subsequently, we can write $n = (N - \ell_{d-1}, \dots, \ell_2 - \ell_1, \ell_1)$. Thus, by induction, the statement holds for all dimensions.

2) We explicitly construct an ordering to prove the assertion. The proof uses induction to obtain the desired result. Suppose that for dimension $d \in \mathbb{N}_+$ and any $N \in \mathbb{N}_+$, the elements of $\mathcal{I}_{d,N}$ can be arranged such that the first element is $(N, 0, \dots, 0)$, the last element is $(0, \dots, 0, N)$, and the l_1 -norm distance between any two consecutive elements is 2. We call this forward ordering of elements. Writing the elements of a forward ordered set gives us reverse ordering, i.e., the first element is $(0, \dots, 0, N)$, the last element is $(N, 0, \dots, 0)$, and the l_1 -norm distance between any two consecutive elements is 2.

For $d = 2$, the elements of $\mathcal{I}_{2,N}$ can be ordered as either $\{(N, 0), (N-1, 1), \dots, (1, N-1), (0, N)\}$ (forward ordering) or $\{(0, N), (1, N-1), \dots, (N-1, 1), (N, 0)\}$ (reverse ordering), so that 1-norm distance between any two consecutive elements is 2.

Assuming that this statement (induction hypothesis) holds for dimension d , we show that it also holds for $d+1$. To that end, we note using the previous result that every element of $\mathcal{I}_{d+1,N}$ can be written as $(N - s_d, \dots, s_2 - s_1, s_1)$ with $0 \leq s_i \leq s_{i+1} \leq N$ for $i \in [d-1]$. Then, we (forward) order the elements of $\mathcal{I}_{d+1,N}$ as follows. Let the first element be $(N, 0, \dots, 0)$. Choose the next sequence of elements as follows. For elements of the form $(N-1, 1 - s_{d-1}, \dots, s_2 - s_1, s_1)$ arrange the elements $(1 - s_{d-1}, \dots, s_2 - s_1, s_1)$ in forward order, which is possible by assumption. For elements of the form $(N-2, 2 - s_{d-1}, \dots, s_2 - s_1)$ arrange the elements $(2 - s_{d-1}, \dots, s_2 - s_1)$ in reverse order. Continuing this way, given elements of the form $(N - s_d, s_d - s_{d-1}, \dots, s_2 - s_1, s_1)$ for fixed s_d , arrange the elements $(s_d - s_{d-1}, \dots, s_2 - s_1, s_1)$ in forward order if s_d is odd, and in reverse order if s_d is even.

Then, for a fixed $0 \leq s_d \leq N$, if $(N - s_d, s_d - s_{d-1}, \dots, s_1')$ is the element after $(N - s_d, s_d - s_{d-1}, \dots, s_1)$ as per the above ordering, we have $\|(N - s_d, s_d - s_{d-1}, \dots, s_1') - (N - s_d, s_d - s_{d-1}, \dots, s_1)\|_1 = \|(s_d - s_{d-1}, \dots, s_1') - (s_d - s_{d-1}, \dots, s_1)\|_1 = 2$ by induction hypothesis. Next, we consider the case when s_d increases by 1. In this case, the last element of $(N - s_d, \dots, s_2 - s_1)$ is $(N - s_d, s_d, 0, \dots, 0)$ if s_d is even and it is $(N - s_d, 0, \dots, 0, s_d)$ if s_d is odd. Then, the first element of the next sequence $(N - s_d - 1, s_d + 1 - s_{d-1}, \dots, s_1)$ is $(N - s_d - 1, s_d + 1, 0, \dots, 0)$ if s_d is even and it is $(N - s_d - 1, 0, \dots, 0, s_d + 1)$ if s_d is odd. Therefore, the l_1 -norm distance between consecutive elements when s_d increases by 1 is equal to 2. Therefore, the elements of $\mathcal{I}_{d+1,N}$ can be ordered as described above. By induction, the result holds for any dimension.

3) We prove this by induction. Given dimension $d \in \mathbb{N}_+$, assume that for every $x \in \Delta_d$, there is some $z \in \Delta_{d,N}$ such that $\|x - z\|_1 \leq 2(d-1)/N$ for all $N \in \mathbb{N}_+$. For $d = 2$, we can write any $x \in \Delta_2$ as $x = (1 - x_1, x_1)$ for some $x_1 \in$

$[0, 1]$. Let $n_1 \in \operatorname{argmin}\{|x_1 - m_1/N| \mid 0 \leq m_1 \leq N\}$, so that have $|x_1 - n_1/N| \leq 1/2N$. Choosing $z = (1 - n_1/N, n_1/N)$, we obtain $\|x - z\|_1 \leq 1/N \leq 2/N$.

Now, suppose that the assumption holds for dimension d . We show that it also holds for dimension $d+1$. To that end, let $x \in \Delta_{d+1}$ be written as $x = (1 - s, x_2, \dots, x_{d+1})$, where $s = \sum_{i=2}^{d+1} x_i$. If $s = 0$, then $x = (1, 0, \dots, 0)$ and the result follows. Therefore, let $s > 0$ and consider the vector $x' = (x_2, \dots, x_{d+1})/s$, so that $x' \in \Delta_d$. Let $M = \lceil sN \rceil$, and by assumption, there is some $z' \in \Delta_{d,M}$ such that $\|x' - z'\|_1 \leq 2(d-1)/M$. Denote $z' = (n_2, \dots, n_{d+1})/M$ and define $n_1 = N - \sum_{i=2}^{d+1} n_i = N - M$. Then, $|(1 - s) - n_1/N| = |sN - M|/N \leq 1/N$ since $sN \leq M \leq sN + 1$. Then, denoting $z = (n_1, \dots, n_{d+1})/N$, we have

$$\begin{aligned} \|x - z\|_1 &= \left| (1 - s) - \frac{n_1}{N} \right| + \|sx' - z'M/N\|_1 \\ &\leq \frac{1}{N} + s \|x' - z'\|_1 + \left| s - \frac{M}{N} \right| \|z'\|_1 \\ &\leq \frac{2}{N} + 2(d-1) \frac{s}{M} \\ &\leq \frac{2d}{N} \end{aligned}$$

where we used the fact the $s/M \leq 1/N$. Thus, the statement holds by induction. \square

Prop. 16.1 helps in iterative construction of the grid, while Prop. 16.2,3 will be used in construction of dense curve. In practice, the ordering of elements given in Prop. 16.2 can be done efficiently for moderately small dimensions. We use this ordering in practice to generate the grid, and this allows for easy parallelization.

Next, based on the results of Ref. [36], we calculate the value of N that needs to be used in the grid search to converge to a precision of $\epsilon > 0$.

Proposition 17: Let $f: \mathcal{D} \rightarrow \mathbb{R}$ be a β -Lipschitz-like function satisfying Eq. (48), where $\mathcal{D} = \Delta_d$ is the standard simplex in \mathbb{R}^d . Let $N = \lceil 1/\delta^2 \rceil$, where $\delta = \sup\{\delta' > 0 \mid \beta(\delta') \leq \epsilon/2\}$. Let $\Delta_{d,N}$ be the grid defined in Eq. (52), and let $f^* = \max\{f(z) \mid z \in \Delta_{d,N}\}$. Then, we have $\max_{x \in \Delta_d} f(x) - f^* \leq \epsilon$.

Proof: This result is implied by techniques developed in Ref. [36]. For the sake of completeness, we give a proof here for the specific case of Lipschitz-like functions.

Given $f: \mathcal{D} \rightarrow \mathbb{R}$, the Bernstein polynomial approximation of f of order $N \in \mathbb{N}_+$ is defined as

$$B_N(f)(x) = \sum_{n \in \mathcal{I}_{d,N}} f\left(\frac{n}{N}\right) \frac{N!}{n!} x^n$$

for $x \in \Delta_d$, where we use the multi-index notation $n! = n_1! \cdots n_d!$, $x^n = x_1^{n_1} \cdots x_d^{n_d}$ for $n \in \mathcal{I}_{d,N}$ [36]. The sup-norm of a continuous function $f: \mathcal{D} \rightarrow \mathbb{R}$ is given by $\|f\|_\infty = \max_{x \in \mathcal{D}} |f(x)|$, where we use the fact that \mathcal{D} is compact. Next, given a continuous function $f: \mathcal{D} \rightarrow \mathbb{R}$, the modulus of continuity of f is defined as

$$\omega(f, \delta) = \max_{\substack{x, y \in \mathcal{D} \\ \|x - y\| \leq \delta}} |f(x) - f(y)|$$

with respect to a suitable norm $\|\cdot\|$ on \mathbb{R}^n [36]. If f is Lipschitz-like, then we use the norm with respect to which f satisfies the Lipschitz-like property. For a β -Lipschitz-like function f , using Eq. (48), we have

$$\omega(f, \delta) \leq \beta(\delta), \quad (101)$$

where we used the fact that β is a monotonically increasing function.

Then, given a β -Lipschitz-like function $f: \mathcal{D} \rightarrow \mathbb{R}$, using Thm. (3.2) of Ref. [36], we have

$$\|B_N(f) - f\|_\infty \leq 2\omega\left(f, \frac{1}{\sqrt{N}}\right) \leq 2\beta\left(\frac{1}{\sqrt{N}}\right)$$

where we used Eq. (101) to obtain the last inequality. In particular, this implies that $\max_{x \in \mathcal{D}} B_N(f) \geq \max_{x \in \mathcal{D}} f(x) - 2\beta(1/\sqrt{N})$. Using this result along with Lemma (3.1) of Ref. [36], we obtain

$$\max_{x \in \mathcal{D}} f(x) - 2\beta\left(\frac{1}{\sqrt{N}}\right) \leq \max_{x \in \mathcal{D}} B_N(f)(x) \leq \max_{x \in \Delta_{d,N}} f(x).$$

Therefore, choosing

$$\beta\left(\frac{1}{\sqrt{N}}\right) \leq \frac{\epsilon}{2},$$

we obtain

$$\max_{x \in \mathcal{D}} f(x) \leq \max_{x \in \Delta_{d,N}} f(x) + \epsilon$$

giving us the desired result. In other words, to compute the maximum of f to within a precision of $\epsilon > 0$, it suffices to search the grid $\Delta_{d,N}$ with $N = \lceil 1/\delta^2 \rceil$, where δ is the largest number satisfying $\beta(\delta) \leq \epsilon/2$. \square

C. Lipschitz-Like Optimization Over the Standard Simplex Using Dense Curves

We show how to optimize Lipschitz-like functions over the standard simplex using α -dense curves. Such curves get within a distance α of all points in the simplex (see Def. (10)). As we show below, by appropriately choosing α , one can perform a one-dimensional optimization to obtain the maximum of f to the desired precision.

Proposition 18: Suppose that $f: \Delta_d \rightarrow \mathbb{R}$ is a β -Lipschitz-like function satisfying Eq. (48). Then the following hold.

- 1) Given $N \in \mathbb{N}_+$, let $\gamma: [0, L_{\text{curve}}] \rightarrow \Delta_d$ be a Lipschitz-like curve constructed as per Alg. 3. Then the curve γ is $(2(d-1)/N)$ -dense in the simplex Δ_d and satisfies Eq. (54).
- 2) Alg. 4 computes the maximum of f to within a precision of $\epsilon > 0$ for any β_γ -Lipschitz-like, $(2(d-1)/N)$ -dense curve γ . Here, $N = \lceil 2(d-1)/\alpha \rceil$ with $\alpha = \sup\{\alpha' > 0 \mid \beta(\alpha') \leq \epsilon/2\}$ as noted in Alg. 4. In the worst case, the algorithm takes

$$\left\lceil \frac{2}{N\delta} \binom{N+d-1}{d-1} \right\rceil$$

time steps to converge to the maximum within a precision of $\epsilon > 0$, where $\delta = \sup\{\delta' > 0 \mid \beta(\beta_\gamma(\delta')) \leq \epsilon/2\}$. When γ is the curve generated by Alg. 3, $\alpha < 1$ and

$d \gg 1$, this amounts to $O(\alpha^{1-d}/d)$ time steps in the worst case.

Proof: 1. From Prop. 16, we know that $\Delta_{d,N}$ is a $(2(d-1)/N)$ -net. Since $\Delta_{d,N} \subseteq \text{Range}(\gamma)$, for any $x \in \Delta_d$, there is some $z \in \text{Range}(\gamma)$ such that $\|x - z\|_1 \leq 2(d-1)/N$. In other words, γ is a $(2(d-1)/N)$ -dense curve in Δ_d . Next, we show that γ is Lipschitz-like function. To that end, given $\theta, \theta' \in \mathbb{R}$, let $k = \lceil \theta N/2 \rceil$ and $k' = \lceil \theta' N/2 \rceil$. Without loss of generality, take $\theta \leq \theta'$. Given $i \in [N_{\text{grid}}]$, where N_{grid} is defined in Eq. (51), denote $x_i \in \Delta_{d,N}$ to be the i th element of $\Delta_{d,N}$ ordered as per Eq. (11). Now, if $k = k'$, then we can write $\gamma(\theta) = tx_k + (1-t)x_{k+1}$ and $\gamma(\theta') = t'x_k + (1-t')x_{k+1}$, where $t = 1 + k - \theta N/2$ and $t' = 1 + k - \theta' N/2$ (see Alg. 3). Therefore, $\|\gamma(\theta') - \gamma(\theta)\|_1 \leq |t - t'| \|x_{k+1} - x_k\|_1 = |\theta' - \theta|$, where we used the fact that $\|x_{k+1} - x_k\|_1 = 2/N$ (see Prop. 16). Thus, consider the case $k' > k$, so that

$$\begin{aligned} \|\gamma(\theta') - \gamma(\theta)\|_1 &\leq \|\gamma(\theta') - x_{k'}\|_1 \\ &\quad + \sum_{i=k+1}^{k'-1} \|x_{i+1} - x_i\|_1 + \|x_{k+1} - \gamma(\theta)\|_1 \\ &= |\theta' - \theta| \end{aligned}$$

where in the last line we used the fact that θ (and θ') is defined as the length along obtained by joining consecutive points of the grid until we reach θ (and θ' respectively). Since $\|x - y\|_1 \leq 2$ for all $x, y \in \Delta_d$, Eq. (54) follows.

2. Consider the real-valued function $g = f \circ \gamma$ defined on the interval $[0, L_{\text{curve}}]$. Here, γ is a $(2(d-1)/N)$ -dense β_γ -Lipschitz-like curve with $N = \lceil 2(d-1)/\alpha \rceil$ and $\alpha = \sup\{\alpha' > 0 \mid \beta(\alpha') \leq \epsilon/2\}$. Since $2(d-1)/N \leq \alpha$, γ is also an α -dense curve. Let $\theta^* \in \text{argmax}\{g(\theta) \mid \theta \in [0, L_{\text{curve}}]\}$, and let $\theta_\epsilon^* \in [0, L_{\text{curve}}]$ be the point output by Alg. 1. Since f is β -Lipschitz-like and γ is β_γ -Lipschitz-like, $g = f \circ \gamma$ is $\beta \circ \beta_\gamma$ -Lipschitz-like, so by Prop. 15, we know that $g(\theta^*) - g(\theta_\epsilon^*) \leq \epsilon/2$. Let $x^* \in \text{argmax}\{f(x) \mid x \in \Delta_d\}$ denote a point achieving the maximum of f . Then, since γ is an α -dense curve, there is some point $\theta_0 \in [0, L_{\text{curve}}]$ such that $\|x^* - \gamma(\theta_0)\|_1 \leq \alpha$. Noting that $g(\theta_0) \leq g(\theta^*)$, we have

$$\begin{aligned} f(x^*) - g(\theta_\epsilon^*) &\leq f(x^*) - g(\theta^*) + \frac{\epsilon}{2} \\ &\leq f(x^*) - g(\theta_0) + \frac{\epsilon}{2} \\ &\leq \beta(\|x^* - \gamma(\theta_0)\|_1) + \frac{\epsilon}{2} \\ &\leq \epsilon \end{aligned}$$

To obtain the penultimate inequality we used the Lipschitz-like property of f along with the fact that $g(\theta_0) = f(\gamma(\theta_0))$. The last inequality follows by noting that $\beta(\|x^* - \gamma(\theta_0)\|_1) \leq \beta(\alpha) \leq \epsilon/2$. The number of iterations needed to compute θ_ϵ^* in the worst-case is given by

$$N_{\text{maxiter}} = \left\lceil \frac{L_{\text{curve}}}{\delta} \right\rceil = \left\lceil \frac{2}{N\delta} \binom{N+d-1}{d-1} \right\rceil$$

where $\delta = \sup\{\delta' > 0 \mid \beta(\beta_\gamma(\delta')) \leq \epsilon/2\}$, which follows from Prop. 15. Using the fact that $(n/k)^k \leq \binom{n}{k} \leq (en/k)^k$ and $2(d-1)/\alpha \leq N \leq 2(d-1)/\alpha + 1$, the number of iterations

is bounded as

$$\frac{\alpha}{(d-1)\delta} \left(\frac{(2(d-1)/\alpha + d-1)}{d-1} \right)^{d-1} \leq N_{\maxiter} \leq \frac{\alpha}{(d-1)\delta} \left(\frac{e(2(d-1)/\alpha + d)}{d-1} \right)^{d-1}.$$

Thus, assuming $\alpha < 1$ and $d \gg 1$, we need $O(\alpha^{2-d}\delta^{-1}/d)$ time steps for convergence. For the case when γ is the curve generated by Alg. 3, we can take $\beta_\gamma(x) = x$ owing to Eq. (54), and therefore, the statement of the proposition follows. \square

D. Lipschitz-Like Optimization Over a Compact & Convex Domain

We now consider the general problem of optimizing a β -Lipschitz-like function $f: \mathcal{D} \rightarrow \mathbb{R}$ satisfying Eq. (48), where $\mathcal{D} \subseteq \mathbb{R}^d$ is a nonempty compact and convex domain. Many techniques have been developed for optimizing Lipschitz or Hölder continuous functions when \mathcal{D} is a specific set such as hypercube, simplex or is the full Euclidean space \mathbb{R}^n as noted in Sec. IV. However, when dealing with an arbitrary compact and convex domain, it is not always obvious how to encode the constraints defining the domain in the algorithm.

We circumvent this problem by looking for a function $\bar{f}: \mathbb{R}^d \rightarrow \mathbb{R}$ satisfying the following properties:

- 1) The function \bar{f} is an extension of f in the sense that \bar{f} is Lipschitz-like and $\bar{f}|_{\mathcal{D}} = f$.
- 2) The maximum of \bar{f} and f coincide. More specifically, $\sup_{x \in \mathcal{K}} \bar{f} = \max_{x \in \mathcal{D}} f$ for every set $\mathcal{D} \subseteq \mathcal{K} \subseteq \mathbb{R}^d$.
- 3) The function \bar{f} can be efficiently computed whenever f and β can be efficiently computed.

The first property ensures that we are able to use existing Lipschitz-like optimization algorithms. The second property ensures that we can compute the maximum of f by computing the maximum of \bar{f} over a convenient set \mathcal{K} containing the domain \mathcal{D} . The choice of the set \mathcal{K} will usually depend on the actual algorithm used to perform the optimization. The last property ensures that the function \bar{f} can be efficiently evaluated in practice.

We show that one can construct a function \bar{f} satisfying all the above properties for the case when β is itself a Lipschitz-like function. That is, there is some non-negative, continuous, monotonically increasing function $\kappa: \mathbb{R}_+ \rightarrow \mathbb{R}$ with $\kappa(0) = 0$ such that, for all $x, y \in \mathbb{R}_+$,

$$|\beta(x) - \beta(y)| \leq \kappa(|x - y|). \quad (102)$$

This assumption is not too restrictive for our purposes because Lipschitz continuous functions (i.e., $\beta(x) = x$), and functions β relevant to entropic quantities like the one in Eq. (64) satisfy Eq. (102) (see Prop. 20). In this case, we can define the extension of \bar{f} as follows.

Definition 19 (Lipschitz-Like Extension): Let $f: \mathcal{D} \rightarrow \mathbb{R}$ be a β -Lipschitz-like function over a compact and convex domain $\mathcal{D} \subseteq \mathbb{R}^d$. Let $\|\cdot\|$ be the norm on \mathbb{R}^d with respect to which Eq. (48) holds. Suppose that β is κ -Lipschitz-like. Then, the Lipschitz-like extension of f is the function $\bar{f}: \mathbb{R}^d \rightarrow \mathbb{R}$ defined as

$$\bar{f}(x) = f(\Pi_{\mathcal{D}}(x)) - \beta(\|x - \Pi_{\mathcal{D}}(x)\|), \quad (103)$$

where $\Pi_{\mathcal{D}} = \operatorname{argmin}_{z \in \mathcal{D}} \|z - x\|_2$ is the projection of x onto \mathcal{D} with respect to the Euclidean norm.

We remark that the choice of norm in Eq. (48), and subsequently in Eq. (103), is flexible. Moreover, because all norms on \mathbb{R}^d are equivalent, one can change the function β so as to get a Lipschitz-like property with respect to a different norm. Recall that two norms $\|\cdot\|_a$ and $\|\cdot\|_b$ on \mathbb{R}^d are said to be equivalent if there exist constants $c_1, c_2 > 0$ (possibly dimension dependent) such that $c_1 \|\cdot\|_a \leq \|\cdot\|_b \leq c_2 \|\cdot\|_a$. The caveat of using equivalence of norms to change the function β is that the modified β might end up depending on the dimension. This could negatively impact the convergence rate (for example, polynomial time convergence guarantee for optimization over the simplex using grid search might be lost if β depends on the dimension). This observation is also of relevance in Def. 19 because we define the projection with respect to the Euclidean norm, while allow an arbitrary choice of norm in Eq. (103) (also see Prop. 20).

Note that $\Pi_{\mathcal{D}}$ is well-defined since \mathcal{D} is compact and convex (see Thm. (2.5) in Ref. [44]). Since $\Pi_{\mathcal{D}}(x) = x$ for $x \in \mathcal{D}$, we can see that \bar{f} is indeed an extension of f . Roughly speaking, the extension is also Lipschitz-like because both f and β are Lipschitz-like. Furthermore, since β is a non-negative monotonically increasing function, the value of \bar{f} decreases as we move away from \mathcal{D} . Thus, the maximum of \bar{f} occurs over \mathcal{D} . Finally, so long as the projection $\Pi_{\mathcal{D}}$ can be efficiently computed, the function \bar{f} can be efficiently computed (assuming f and β can be efficiently computed).

We formalize the observations made above in the following proposition.

Proposition 20: Let $\beta: \mathbb{R}_+ \rightarrow \mathbb{R}$ be a non-negative, continuous, monotonically increasing function with $\beta(0) = 0$. Suppose that β is κ -Lipschitz-like in the sense of Eq. (102) for some appropriate κ . Let $f: \mathcal{D} \rightarrow \mathbb{R}$ be a β -Lipschitz-like function over some compact and convex domain $\mathcal{D} \subseteq \mathbb{R}^d$. Let \bar{f} be the extension of f as defined in Eq. (103). Then the following statements hold.

- 1) Define the constant $C = c_2/c_1$, where $c_1, c_2 > 0$ are obtained from equivalence of $\|\cdot\|, \|\cdot\|_2$ on \mathbb{R}^d , i.e., $c_1 \|\cdot\| \leq \|\cdot\|_2 \leq c_2 \|\cdot\|$. Then, the extension \bar{f} is a $\bar{\beta}$ -Lipschitz-like function, where

$$\bar{\beta}(x) = \beta(Cx) + \kappa((C+1)x). \quad (104)$$

- 2) Given any set $\mathcal{D} \subseteq \mathcal{K} \subseteq \mathbb{R}^d$, we have $\sup_{x \in \mathcal{K}} \bar{f}(x) = \max_{x \in \mathcal{D}} f(x)$. Moreover, any point achieving the maximum of f achieves the maximum of \bar{f} .
- 3) The modified binary entropy \bar{h} defined in Eq. (49) is an \bar{h} -Lipschitz-like function, that is,

$$|\bar{h}(x) - \bar{h}(y)| \leq \bar{h}(|x - y|)$$

for all $x, y \in \mathbb{R}_+$. Subsequently, β_I defined in Eq. (64) is a β_I -Lipschitz-like function.

Proof: 1. We first note that $\Pi_{\mathcal{D}}$ is a non-expansive mapping on \mathbb{R}^d , i.e., for all $x, y \in \mathbb{R}^d$, we have $\|\Pi_{\mathcal{D}}(x) - \Pi_{\mathcal{D}}(y)\|_2 \leq \|x - y\|_2$ [45]. By equivalence of norms, we can find a (possibly dimension-dependent) constant $C > 0$ noted in the statement of the proposition, such that

$\|\Pi_{\mathcal{D}}(x) - \Pi_{\mathcal{D}}(y)\| \leq C \|x - y\|$. From this, we can infer the following list of inequalities:

$$\begin{aligned} |\bar{f}(x) - \bar{f}(y)| &\leq |f(\Pi_{\mathcal{D}}(x)) - f(\Pi_{\mathcal{D}}(y))| \\ &\quad + |\beta(\|x - \Pi_{\mathcal{D}}(x)\|) - \beta(\|y - \Pi_{\mathcal{D}}(y)\|)| \\ &\leq \beta(\|\Pi_{\mathcal{D}}(x) - \Pi_{\mathcal{D}}(y)\|) \\ &\quad + \kappa(\|x - \Pi_{\mathcal{D}}(x)\| - \|y - \Pi_{\mathcal{D}}(y)\|) \\ &\leq \beta(C\|x - y\|) + \kappa(\|x - y\| + \|\Pi_{\mathcal{D}}(y) - \Pi_{\mathcal{D}}(x)\|) \\ &\leq \beta(C\|x - y\|) + \kappa(\|x - y\| + \|\Pi_{\mathcal{D}}(x) - \Pi_{\mathcal{D}}(y)\|) \\ &\leq \beta(C\|x - y\|) + \kappa((C + 1)\|x - y\|) \\ &= \bar{\beta}(\|x - y\|) \end{aligned}$$

To obtain the second inequality, we use the fact that f is β -Lipschitz-like and β is κ -Lipschitz-like. To obtain the third inequality, we use the reverse triangle inequality, the fact that $\Pi_{\mathcal{D}}$ is non-expansive, and that β and κ are monotonically increasing functions. The last two inequalities follow from similar observations. Finally, we note that $\bar{\beta}(x) = \beta(Cx) + \kappa((C + 1)x)$ is a non-negative, continuous, monotonically increasing function with $\bar{\beta}(0) = 0$.

2. Since β is a non-negative function, we have $\bar{f}(x) \leq f(\Pi_{\mathcal{D}}(x))$ for all $x \in \mathbb{R}^d$. Therefore, given any set $\mathcal{D} \subseteq \mathcal{K} \subseteq \mathbb{R}^d$, we have $\sup_{x \in \mathcal{K}} \bar{f}(x) \leq \sup_{x \in \mathcal{K}} f(\Pi_{\mathcal{D}}(x)) = \max_{x \in \mathcal{D}} f(x)$. The last equality follows by noting that $\Pi_{\mathcal{D}}(x) \in \mathcal{D}$ for all $x \in \mathbb{R}^d$, f is continuous and \mathcal{D} is compact. If $x^* \in \mathcal{D}$ achieves the maximum of f , then $\sup_{x \in \mathcal{K}} \bar{f}(x) = \max_{x \in \mathcal{D}} f(x) = f(x^*) = \bar{f}(x^*)$, where in the last step, we used the fact that $f = \bar{f}$ on \mathcal{D} .

3. Let \bar{h} be the modified binary entropy defined in Eq. (49). If $x \in [0, 1/2]$ then, $\bar{h}(x) = h(x)$. Furthermore, we have $h(x) = H((x, 1-x))$ for $x \in [0, 1]$, where H is the Shannon entropy. Therefore, given $x, y \in [0, 1/2]$, we have $|\bar{h}(x) - \bar{h}(y)| = |h(x) - h(y)| = |H((x, 1-x)) - H((y, 1-y))| \leq h(|x - y|) = \bar{h}(|x - y|)$, where the inequality follows from the results of Ref. [28]. Here, we used the fact that $|x - y| \leq 1/2$ when $x, y \in [0, 1/2]$. For $x, y \geq 1/2$, we have $\bar{h}(x) = \bar{h}(y) = \ln(2)$, and thus, $|\bar{h}(x) - \bar{h}(y)| \leq \bar{h}(|x - y|)$. Finally, we consider the case when $x \in [0, 1/2]$ and $y \geq 1/2$. Here, we have $|\bar{h}(x) - \bar{h}(y)| = \ln(2) - h(x) = h(1/2) - h(x) \leq h(|1/2 - x|) = \bar{h}(|1/2 - x|) \leq \bar{h}(|y - x|)$. In the last step, we used the fact that $|1/2 - x| \leq |y - x|$ and \bar{h} is monotonically increasing. Therefore, we have $|\bar{h}(x) - \bar{h}(y)| \leq \bar{h}(|x - y|)$ for all $x, y \in \mathbb{R}_+$. \square

The above extension gives some freedom in determining what algorithm to use to perform the maximization of \bar{f} , especially for the case when f is Lipschitz or Hölder continuous. One can, for example, use an unconstrained optimization method developed for optimization of Lipschitz or Hölder continuous function. Alternatively, one can embed \mathcal{D} inside a hypercube \mathcal{K} , and use an algorithm that can perform global Lipschitz optimization [37], [46], [47]. In this study, we generalize the method used in Ref. [37] using dense curves because the proof of convergence essentially follows Prop. 18. In practice, one might get faster convergence by generalizing the global optimization method using Hilbert space-filling curves studied in Ref. [46].

We use the α -dense curve for filling the hypercube $\mathcal{K} = \prod_{i=1}^d [a_i, b_i]$ used in Ref. [37], which we reproduce below for convenience. Let $d \geq 2$, $\eta > 0$, $\eta_1 = 1$ and define

$$\eta_i = \left(\frac{\eta}{\pi}\right)^{i-1} \prod_{j=2}^i \frac{1}{|a_j| + |b_j|} \quad (105)$$

for $i = 2, \dots, d$. Then, the curve $\gamma: [0, \pi/\eta_d] \rightarrow \mathcal{K}$ defined as

$$\gamma_i(\theta) = \frac{a_i - b_i}{2} \cos(\eta_i \theta) + \frac{a_i + b_i}{2} \quad (106)$$

for $i \in [d]$ is a $(\sqrt{d-1}\eta)$ -dense curve with respect to the Euclidean norm [37]. Furthermore, γ is Lipschitz continuous (with respect to the Euclidean norm) with Lipschitz constant

$$L_\gamma = \frac{1}{2} \left(\sum_{i=1}^n (|a_i| + |b_i|)^2 \eta_i^2 \right)^{1/2}. \quad (107)$$

Algorithm 8 Computing the Maximum of a β -Lipschitz-Like Function f Satisfying Eq. (48) for $\mathcal{D} = \Delta_d$, Given $\epsilon > 0$

```

1: function
    MAXIMIZE_LIPSCHITZ-LIKE_FUNC
    _COMPACT_CONVEX_DOMAIN( $d, \beta, \epsilon$ )
2:   Find  $\mathcal{K} = \prod_{i=1}^d [a_i, b_i] \subseteq \mathbb{R}^d$  such that  $\mathcal{D} \subseteq \mathcal{K}$ 
3:   Find  $c_1 > 0$  such that  $c_1 \|\cdot\| \leq \|\cdot\|_2$ 
4:   Construct the curve  $\gamma$  as per Eq. (106) for
     $\eta = c_1 \alpha / \sqrt{d-1}$ , where
     $\alpha = \sup\{\alpha' > 0 \mid \bar{\beta}(\alpha') \leq \epsilon/2\}$ 
5:   Construct the extension  $\bar{f}$  as per Eq. (103)
6:   Compute the maximum  $\bar{g}^*$  of  $\bar{g} = \bar{f} \circ \gamma$  over  $[0, \pi/\eta_d]$ 
    to a precision of  $\epsilon/2$  using Alg. 1
7:   return  $\bar{g}^*$ 
8: end function

```

Below, we show that the above algorithm will converge to the maximum within precision $\epsilon > 0$. The proof essentially adapts the ideas used to prove Prop. 18.

Proposition 21: Let $f: \mathcal{D} \rightarrow \mathbb{R}$ be a β -Lipschitz-like function with respect to the norm $\|\cdot\|$ on \mathbb{R}^d , where $\mathcal{D} \subseteq \mathbb{R}^d$ is a compact & convex set. Suppose that β is κ -Lipschitz-like. Let $\mathcal{K} \subseteq \mathbb{R}^d$ be a bounded set containing \mathcal{D} . Let γ be any α -dense, β_γ -Lipschitz-like curve from the interval $[a, b]$ into the set \mathcal{K} . Let \bar{f} be a $\bar{\beta}$ -Lipschitz-like extension of f defined in Eq. (103), where $\bar{\beta}$ is given in Eq. (104). Suppose that we compute the maximum of f using an appropriate generalization of Alg. 8 with the curve γ . Then, for $\alpha = \sup\{\alpha' > 0 \mid \bar{\beta}(\alpha') \leq \epsilon/2\}$, the algorithm computes the maximum of f to within a precision of $\epsilon > 0$. In the worst case, this algorithm takes $\lceil (b-a)/\delta \rceil$ time steps to converge, where $\delta = \sup\{\delta' > 0 \mid \bar{\beta}(\beta_\gamma(\delta')) \leq \epsilon/2\}$.

In particular, if \mathcal{K} is a hypercube containing \mathcal{D} and γ is the curve generated as per Eq. (106), a value of $\eta = c_1 \alpha / \sqrt{d-1}$ suffices to converge to the maximum of f within a precision of $\epsilon > 0$. Here, α is as defined above and $c_1 > 0$ is a (possibly dimension dependent) constant satisfying $c_1 \|\cdot\| \leq \|\cdot\|_2$. The algorithm takes $\lceil \frac{\pi}{\eta_d \delta} \rceil$ time steps to converge in the worst case, where η_d is defined in Eq. (105) and δ is as defined

above. For the case when f is a Lipschitz continuous function with Lipschitz constant L with respect to the Euclidean norm, Alg. 8 takes $O((6\pi L\sqrt{d}/\epsilon)^d)$ iterations to converge to a precision of $\epsilon > 0$. The exponential scaling with L and ϵ cannot be improved without additional assumptions on the functions or the domain.

Proof: Since $\mathcal{D} \subseteq \mathbb{R}^d$ is compact, it is bounded, and can therefore be embedded into a bounded set $\mathcal{K} \subseteq \mathbb{R}^d$. Let \bar{f} be the $\bar{\beta}$ -Lipschitz-like extension of f defined in Eq. (103), where $\bar{\beta}$ is given in Eq. (104). Let $\gamma: [a, b] \rightarrow \mathcal{K}$ be an α -dense, β_γ -Lipschitz-like curve, where $\alpha = \sup\{\alpha' > 0 \mid \bar{\beta}(\alpha') \leq \epsilon/2\}$. Consider the real-valued function $\bar{g} = \bar{f} \circ \gamma$ defined on the interval $[a, b]$. Let $\theta^* \in \operatorname{argmax}\{\bar{g}(\theta) \mid \theta \in [a, b]\}$, and let $\theta_\epsilon^* \in [a, b]$ be the point output by Alg. 1. Since \bar{f} is $\bar{\beta}$ -Lipschitz-like and γ is β_γ -Lipschitz-like, $\bar{g} = \bar{f} \circ \gamma$ is $\bar{\beta} \circ \beta_\gamma$ -Lipschitz-like. So, by Prop. 15, we know that $\bar{g}(\theta^*) - \bar{g}(\theta_\epsilon^*) \leq \epsilon/2$.

Let $x^* \in \operatorname{argmax}\{f(x) \mid x \in \Delta_d\}$ denote a point achieving the maximum of f . From Prop. 20, we know that $\bar{f}(x^*) = \sup_{x \in \mathcal{K}} \bar{f}(x) = \sup_{x \in \mathcal{D}} f(x) = f(x^*)$. Then, since γ is an α -dense curve and $\mathcal{D} \subseteq \mathcal{K}$, there is some point $\theta_0 \in [a, b]$ such that $\|x^* - \gamma(\theta_0)\|_1 \leq \alpha$. Noting that $\bar{g}(\theta_0) \leq \bar{g}(\theta^*)$, we have

$$\begin{aligned} f(x^*) - \bar{g}(\theta_\epsilon^*) &\leq \bar{f}(x^*) - \bar{g}(\theta^*) + \frac{\epsilon}{2} \\ &\leq \bar{f}(x^*) - \bar{g}(\theta_0) + \frac{\epsilon}{2} \\ &\leq \bar{\beta}(\|x^* - \gamma(\theta_0)\|_1) + \frac{\epsilon}{2} \\ &\leq \epsilon \end{aligned}$$

To obtain the penultimate inequality we used the Lipschitz-like property of \bar{f} along with the fact that $\bar{g}(\theta_0) = \bar{f}(\gamma(\theta_0))$. The last inequality follows by noting that $\bar{\beta}(\|x^* - \gamma(\theta_0)\|_1) \leq \bar{\beta}(\alpha) \leq \epsilon/2$. The number of iterations needed to compute θ_ϵ^* in the worst-case is given by

$$N_{\text{maxiter}} = \left\lceil \frac{b-a}{\delta} \right\rceil$$

where $\delta = \sup\{\delta' > 0 \mid \bar{\beta}(\beta_\gamma(\delta')) \leq \epsilon/2\}$, which follows from Prop. 15.

Now suppose that \mathcal{K} is a hypercube containing \mathcal{D} and γ is the curve given in Eq. (106). By equivalence of norms on \mathbb{R}^d , we have $c_1 \|\cdot\| \leq \|\cdot\|_2 \leq c_2 \|\cdot\|$ for some (possibly dimension dependent) constants $c_1, c_2 > 0$. Therefore, γ is $(\sqrt{d-1}\eta/c_1)$ -dense and (L_g/c_1) -Lipschitz continuous with respect to the norm $\|\cdot\|$, where L_g is given in Eq. (107). By the above results, we can choose $\eta = c_1\alpha/\sqrt{d-1}$ for convergence, where $\alpha = \{\alpha' > 0 \mid \bar{\beta}(\alpha') \leq \epsilon/2\}$. In this worst case, this algorithm takes $\lceil \pi/(\eta_d\delta) \rceil$, where η_d is defined in Eq. (105) and $\delta = \sup\{\delta' > 0 \mid \bar{\beta}(L_g\delta') \leq \epsilon/2\}$.

Consider the case when $d \geq 2$, $\mathcal{D} = [0, 1]^d$, $\mathcal{K} = \mathcal{D}$, and f is Lipschitz continuous with Lipschitz constant L (independent of the dimension) with respect to the Euclidean norm $\|\cdot\|_2$. Then, $c_1 = c_2 = 1$ as defined above, and subsequently, C defined in Prop. 20 is equal to 1. Then, since $\kappa(x) = \beta(x) = Lx$, we have $\bar{\beta}(x) = 3Lx$. Subsequently, we have $\alpha = \epsilon/6L$, and for $\epsilon < 6L$, we have $\alpha < 1$. For the curve γ , we have $\eta = \alpha/\sqrt{d-1} = \epsilon/6L\sqrt{d-1}$, $\eta_d = (\eta/\pi)^{d-1}$, $\delta = \epsilon/6LL\gamma$,

and $L\gamma \leq \sqrt{d}/2$ assuming that $\alpha < 1$. Thus, the Alg. 8 takes

$$\left\lceil \frac{6\pi^d LL\gamma}{\eta^{d-1}\epsilon} \right\rceil \leq \left\lceil \frac{(6\pi L\sqrt{d})^d}{\epsilon^d} \right\rceil$$

iterations to converge. It is known that any algorithm (in the sense of a black box model) needs at least $O((L/2\epsilon)^d)$ iterations to compute the optimum of a L -Lipschitz function over the unit hypercube to a precision of $\epsilon > 0$ (see Ref. [27] for details). Thus, one cannot, in general improve the $1/\epsilon^d$ scaling for Alg. 8 without additional assumptions on \mathcal{D} or the class of functions that are optimized by the algorithm. \square

The above result shows that the exponential scaling with the dimension cannot be improved without additional assumptions on the functions or the domain. It is likely the case that the \sqrt{d} factor is sub-optimal and comes from the choice of the curve γ in Eq. (106), so one might be able to improve that factor by using better constructions for the curve or by resorting to a different approach altogether.

We remark that the above algorithm will perform worse than Alg. 2 and Alg. 4 in practice when \mathcal{D} is the standard simplex. This is because we essentially use no information about the domain except for the projection $\Pi_{\mathcal{D}}$ in constructing the extension \bar{f} (see Eq. (103)). Indeed, when using grid search specially designed for $\mathcal{D} = \Delta_d$, we can get polynomial scaling with the dimension for a fixed precision (see Prop. 17). Thus, it might be preferable to resort to methods designed specifically for the particular domain and function class of interest, as opposed to general algorithms like Alg. 8. That said, an advantage of Alg. 8 is that the maximum is computed by finding successively better approximations. Therefore, one can decide to terminate the computation after some fixed number of iterations in order to obtain an upper bound on the maximum.

APPENDIX D ANALYSIS OF ALGORITHMS FOR COMPUTING THE SUM CAPACITY FOR TWO-SENDER MACS

Proof of Prop. 12: For any probability distributions $p_1^Z, p_2^Z \in \Delta_{d_o}$, we have $|H(p_1^Z) - H(p_2^Z)| \leq \theta \log(d_o - 1) + h(\theta)$, where $\theta = \|p_1^Z - p_2^Z\|_1/2$ and h is the binary entropy function [28]. In particular, if we define \bar{h} as in Eq. (49), then we have

$$\begin{aligned} |H(p_1^Z) - H(p_2^Z)| &\leq \frac{1}{2} \log(d_o - 1) \|p_1^Z - p_2^Z\|_1 \\ &\quad + \bar{h}\left(\frac{1}{2} \|p_1^Z - p_2^Z\|_1\right) \end{aligned} \quad (108)$$

because $h(x) \leq \bar{h}(x)$ for all $x \in [0, 1]$.

Now, note that using the same ideas employed in obtaining $I(p, q)$ given in Eq. (60), we can interchange the roles of $p \in \Delta_{d_1}$ and $q \in \Delta_{d_2}$ to obtain

$$I(p, q) = H(A_p q) - \langle b_p, q \rangle$$

where

$$A_p(z, b_2) = \sum_{b_1 \in \mathcal{B}_1} N(z|b_1, b_2)p(b_1)$$

and

$$b_p(b_2) = - \sum_{b_1 \in \mathcal{B}_1} p(b_1) \sum_{z \in \mathcal{Z}} N(z|b_1, b_2) \log(N(z|b_1, b_2)).$$

Then, for any $q, q' \in \Delta_{d_2}$ and a fixed $p \in \Delta_{d_1}$, we have

$$\begin{aligned} & |I(p, q) - I(p, q')| \\ & \leq |H(A_p q) - H(A_p q')| + |\langle b_p, q - q' \rangle| \\ & \leq \frac{1}{2} \log(d_o - 1) \|A_p q - A_p q'\|_1 \\ & \quad + \bar{h} \left(\frac{1}{2} \|A_p q - A_p q'\|_1 \right) + \|b_p\|_\infty \|q - q'\|_1 \end{aligned} \quad (109)$$

where in the second step, we used Eq. (108) and Hölder's inequality.

Now, we note that $\|A_p q - A_p q'\|_1 \leq \|A_p\|_{1 \rightarrow 1} \|q - q'\|_1$, where $\|A_p\|_{1 \rightarrow 1}$ is the induced matrix norm with respect to l_1 -norm on the domain and the co-domain of A_p . This norm is equal to the maximum (absolute value) column sum of A_p . Since A_p is a left stochastic matrix, all of its entries are non-negative and all of its columns sum to 1, and therefore, $\|A_p\|_{1 \rightarrow 1} = 1$. Similarly, we have $0 \leq b_p(b_2) \leq H_{\mathcal{N}}^{\max}$ for all $b_2 \in \mathcal{B}_2$ (see Eq. (65) for definition of $H_{\mathcal{N}}^{\max}$). Therefore, we have $\|b_p\|_\infty \leq H_{\mathcal{N}}^{\max}$. Then, since \bar{h} a monotonically increasing function, from Eq. (109), we obtain

$$\begin{aligned} |I(p, q) - I(p, q')| & \leq \frac{1}{2} \log(d_o - 1) \|q - q'\|_1 + \bar{h} \left(\frac{1}{2} \|q - q'\|_1 \right) \\ & \quad + H_{\mathcal{N}}^{\max} \|q - q'\|_1 \\ & = \beta_I (\|q - q'\|_1) \end{aligned} \quad (110)$$

Since Eq. (110) holds for every $p \in \Delta_{d_1}$, we have for fixed $q, q' \in \Delta_{d_2}$

$$\max_{p \in \Delta_{d_1}} I(p, q) \leq \max_{p \in \Delta_{d_1}} I(p, q') + \beta_I (\|q - q'\|_1)$$

and

$$\max_{p \in \Delta_{d_1}} I(p, q') \leq \max_{p \in \Delta_{d_1}} I(p, q) + \beta_I (\|q - q'\|_1)$$

Since $I^*(q) = \max_{p \in \Delta_{d_1}} I(p, q)$, we obtain $|I^*(q) - I^*(q')| \leq \beta_I (\|q - q'\|_1)$. \square

Proposition 22: Let \mathcal{N} be any two-sender MAC with input alphabets of size d_1, d_2 , and output alphabet of size d_o . Suppose that that $d_1, d_o \geq 2$ and $d_2 = 2$. Then, for $0 < \epsilon \leq 3 \log(d_o)$, the number of iterations required by the while loop in Alg. 5 to compute the sum capacity of the MAC \mathcal{N} to a precision $\epsilon > 0$ is bounded above by

$$\left\lceil \frac{2(\mu^2 + 4)}{(\epsilon\mu + 4) - \sqrt{16 + 8\epsilon\mu - 4\epsilon^2}} \right\rceil \quad (111)$$

where $\mu = 3 \log(d_o)$. In particular, for a precision $0 < \epsilon \leq 3$ chosen independent of d_o , the number of iterations is bounded above by $O(\log(d_o)/\epsilon)$ for $\log(d_o) \gg 1$.

The total cost of computing the sum capacity to a precision $\epsilon > 0$, including the costs of performing intermediate steps such as computing I^* , sorting and root-finding, is at most polynomial in d_1, d_o and $1/\epsilon$.

Proof: To compute the total number of iterations required to compute the sum capacity to a precision $\epsilon > 0$, we break the analysis into the following steps.

- 1) Number of iterations required for the while loop to converge in Alg. 5.
- 2) Cost of computing $I^*(s)$ in Alg. 5 to a precision $\epsilon_I > 0$.
- 3) Cost of sorting the points $s^{(0)}, \dots, s^{(k)}$ in Alg. 5.
- 4) Cost for finding a root of $g_i(s)$ in Alg. 5 to a precision $\epsilon_r > 0$.

From these, we can compute the total cost of converging to a precision $\epsilon > 0$.

1. We begin by computing the number of iterations required for Alg. 5 to converge to a precision $\epsilon > 0$. To that end, note that for obtaining Alg. 5 from Alg. 1, we used the fact that $\|q_s - q_{s'}\|_1 = 2|s - s'|$, where given any $s \in [0, 1]$, we define $q_s = (s, 1 - s)$. Therefore, we have an additional factor of 2 in $\beta_I(2|s - s'|)$ in Alg. 5. Therefore, to make Alg. 5 the same as Alg. 1, we define $\beta(x) = \beta_I(2x)$. Then, from Prop. 15, we know that the number of iterations required to converge to the maximum within an error of $\epsilon > 0$ is bounded above by $\lceil 1/\delta \rceil$, where we used the fact that $\mathcal{D} = [0, 1]$ for the optimization. The number $\delta > 0$ is chosen such that $\beta(x) \leq \epsilon/2$ whenever $0 \leq x \leq \delta$.

In the following analysis, we measure entropy in bits. To proceed, we note that the binary entropy h satisfies the inequality $h(x) \leq 2\sqrt{x(1-x)}$ [48]. Further, for any probability transition matrix \mathcal{N} , the quantity $H_{\mathcal{N}}^{\max}$ defined in Eq. (65) is bounded above by $\log(d_o)$, where d_o is the size of the output alphabet. For $x \leq 1/2$, we have $h(x) = \bar{h}(x)$, where \bar{h} is the modified binary entropy function defined in Eq. (49). Then, from Eq. (64) and the definition $\beta(x) = \beta_I(2x)$, it follows that

$$\beta(x) < \mu x + 2\sqrt{x(1-x)}$$

for $x \leq 1/2$, where $\mu = 3 \log(d_o)$. Then, solving the inequality $\mu x + 2\sqrt{x(1-x)} \leq \epsilon/2$, we obtain

$$x \leq \frac{(\epsilon\mu + 4) - \sqrt{16 + 8\epsilon\mu - 4\epsilon^2}}{2(\mu^2 + 4)}$$

which holds whenever $\epsilon \leq \mu$. Therefore, we choose

$$\delta = \frac{(\epsilon\mu + 4) - \sqrt{16 + 8\epsilon\mu - 4\epsilon^2}}{2(\mu^2 + 4)}$$

which is a positive number for $0 < \epsilon \leq \mu$. Then, from Prop. 15, we know that the number of iterations to converge to a precision $\epsilon > 0$ is bounded above by

$$K_{\text{PS}} = \left\lceil \frac{2(\mu^2 + 4)}{(\epsilon\mu + 4) - \sqrt{16 + 8\epsilon\mu - 4\epsilon^2}} \right\rceil$$

where the subscript PS stands for Piyavskii-Shubert.

2. Next, we calculate the cost of computing $I^*(s)$ to a precision $\epsilon_I > 0$. This is a non-trivial cost because is obtained by solving a convex optimization problem. Specifically, $I^*(s) = \max_{p \in \Delta_{d_1}} H(A_{q_s} p) - \langle b_{q_s}, p \rangle$, where $q_s = (s, 1 - s)$. This cost depends on the algorithm one uses to solve the optimization problem. For example, if one uses an interior point method based on the log-barrier, then one needs

at most $O(d_1^3\sqrt{d_1+d_o}\log((d_1+d_o)/\epsilon_I))$ flops to converge to the optimum within a precision of $\epsilon_I > 0$ (see Ch. (11.5) in Ref. [40]). We denote the cost of computing I^* as K_I .

3. If we use quicksort to sort the numbers $s^{(0)}, \dots, s^{(k)}$, then in the worst case, one needs $O(k^2)$ operations to perform this sorting.

4. If we compute the root to a precision $\epsilon_r > 0$ using bisection, then one needs at most $O(\log(1/\epsilon_r))$ iterations to find this root.

The total cost of sorting (including all K_{PS} iterations) is bounded above by $O(K_{\text{PS}}^3)$. The number of times the root needs to be computed is bounded above by $O(K_{\text{PS}}^2)$ and the number of times the function I^* is calculated in the algorithm is $O(K_{\text{PS}}^2 \log(1/\epsilon_r))$. Since we can only compute I^* to a precision ϵ_I and find the root to a precision ϵ_r , we need to choose these small enough so that the total error is below ϵ . In particular, if we choose $\epsilon_I = \epsilon_r$ such that $O(K_{\text{PS}}^2 \log(1/\epsilon_I))\epsilon_I = \epsilon/2$, then using the fact that $-x \ln(x) \geq x - x^2$ for $x > 0$, we can infer that it is sufficient to choose $\epsilon_I = O(\epsilon/K_{\text{PS}}^2)$. Then running the while loop in Alg. 5 to a precision $\epsilon/2$ so that the total error is within a tolerance of ϵ , we can infer that the total number of iterations needed is bounded above by

$$O(K_I K_{\text{PS}}^2 \log(K_{\text{PS}}^2/\epsilon))$$

where we used the fact that K_I is at least as large as K_{PS} . Therefore, the total cost to converge to the optimum is at most polynomial in d_1 , d_o , and $1/\epsilon$. \square

Next, we obtain an upper bound on the number of iterations needed to compute the sum capacity of an arbitrary two-sender MAC to a fixed precision. The algorithm we use to analyze this is grid search explained in Sec. IV-D.1, and many of the calculations follow the ideas given in Prop. 22.

Proposition 23: Let \mathcal{N} be any two-sender MAC with input alphabets of size d_1 , d_2 , and output alphabet of size d_o . Suppose that that $d_1, d_2, d_o \geq 2$ with $d_1 \geq d_2$. Suppose that all the entropies are measured in bits and denote $\mu = 3 \log(d_o)$. Then, for a fixed $0 < \epsilon \leq 3$, the number of iterations required by the grid search in Alg. 5 to compute the sum capacity of the MAC \mathcal{N} to a precision $\epsilon > 0$ is bounded above by

$$\left(\frac{e\epsilon^2}{\mu^2}d_2 + e\right)^{\frac{4\mu^2}{\epsilon^2}+1} \quad (112)$$

when $\mu \geq \max\{(8/\epsilon) + 2\sqrt{(16/\epsilon^2) - 1}, \epsilon/2\}$.

The total cost of computing the sum capacity to a precision $\epsilon > 0$, including the cost of computing I^* , is at most

$$\text{poly}\left(d_1, d_o, \frac{1}{\epsilon}\right) \left(\frac{e\epsilon^2}{36} \frac{d_2}{(\log(d_o))^2} + e\right)^{\frac{96(\log(d_o))^2}{\epsilon^2}+2} \quad (113)$$

for $0 < \epsilon \leq 3$ and $\mu \geq \max\{(16/\epsilon) + 2\sqrt{(64/\epsilon^2) - 1}, \epsilon/4\}$.

Proof: The sum capacity of a two sender MAC can be expressed as $S(\mathcal{N}) = \max_{q \in \Delta_{d_2}} I^*(q)$, where I^* is defined in Eq. (62). Note that I^* is β_I -Lipschitz-like, where β_I is defined in Eq. (64). By Prop. 17, we know that the number of iterations needed for the grid search algorithm to compute

the sum capacity is equal to

$$K_{\text{GS}} = \binom{N + d_2 - 1}{d_2 - 1},$$

where $N = \lceil 1/\delta^2 \rceil$ and $\delta = \sup\{\delta' > 0 \mid \beta_I(\delta') \leq \epsilon/2\}$. In particular, the number $\delta > 0$ can be chosen such that $\beta_I(x) \leq \epsilon/2$ whenever $0 \leq x \leq \delta$ to obtain an upper bound on the number of iterations.

As noted in the proof of Prop. 22, the binary entropy h (measured in bits) satisfies the inequality $h(x) \leq 2\sqrt{x(1-x)}$ [48]. Further, the quantity H_N^{\max} defined in Eq. (65) is bounded above by $\log(d_o)$. For $x \leq 1/2$, we have $h(x) = \bar{h}(x)$, where \bar{h} is the modified binary entropy function defined in Eq. (49). Then, from Eq. (64), it follows that

$$\beta_I(x) < \frac{\mu}{2}x + \sqrt{x(2-x)}$$

for $x \leq 1/2$, where $\mu = 3 \log(d_o)$. Then, solving the inequality $(\mu/2)x + \sqrt{x(2-x)} \leq \epsilon/2$, we obtain

$$x \leq \frac{(\epsilon\mu + 4) - 2\sqrt{4 + 2\epsilon\mu - \epsilon^2}}{\mu^2 + 4}$$

which holds whenever $\epsilon \leq \mu$. Therefore, we choose

$$\delta = \frac{(\epsilon\mu + 4) - 2\sqrt{4 + 2\epsilon\mu - \epsilon^2}}{\mu^2 + 4}$$

which is a positive number for $0 < \epsilon \leq \mu$. For a fixed $0 < \epsilon \leq \min\{16, \mu\}$ and $\mu \geq \max\{(8/\epsilon) + 2\sqrt{(16/\epsilon^2) - 1}, \epsilon/2\}$, it can be verified that

$$\frac{\mu}{\epsilon} \leq \frac{1}{\delta} \leq \frac{2\mu}{\epsilon}.$$

Further, since $\binom{n}{k} \leq (en/k)^k$, we can write $\binom{N+d_2-1}{d_2-1} = \binom{N+d_2-1}{N} \leq (e(N+d_2-1)/N)^N$. Then, noting that $1/\delta^2 \leq N \leq 1/\delta^2 + 1$, the number of iterations needed for grid search to converge to a precision $0 < \epsilon \leq \min\{\mu, 16\}$ is bounded above by

$$K_{\text{GS}} \leq \left(\frac{e\epsilon^2}{\mu^2}d_2 + e\right)^{\frac{4\mu^2}{\epsilon^2}+1}$$

when $\mu \geq \max\{(8/\epsilon) + 2\sqrt{(16/\epsilon^2) - 1}, \epsilon/2\}$. In order to avoid specifying simultaneous conditions on both ϵ and μ , we note that $\mu = 3 \log(d_o) \geq 3$ for $d_o \geq 2$. Therefore, we can simply assume $0 < \epsilon \leq 3$ in the above equations.

Next, as noted in the proof of Prop. 22, the cost of computing I^* to a precision of $\epsilon_I > 0$ is at most $O(d_1^3\sqrt{d_1+d_o}\log((d_1+d_o)/\epsilon_I))$. Choosing $\epsilon_I = \epsilon/2K_{\text{GS}}$, we can ensure that after K_{GS} calls to I^* , the error is at most $\epsilon/2$. Then, solving the grid search to a precision of $\epsilon/2$, we can infer that the total cost of computing the sum capacity is at most

$$\begin{aligned} & O\left(d_1^3\sqrt{d_1+d_o}\log\left(\frac{2(d_1+d_o)K_{\text{GS}}}{\epsilon}\right)\right) \left(\frac{e\epsilon^2}{4\mu^2}d_2 + e\right)^{\frac{16\mu^2}{\epsilon^2}+1} \\ & \leq O\left(d_1^3\sqrt{d_1+d_o}\log\left(\frac{2(d_1+d_o)}{\epsilon}\right)\right) \left(\frac{e\epsilon^2}{4\mu^2}d_2 + e\right)^{\frac{32\mu^2}{\epsilon^2}+2} \end{aligned}$$

when $0 < \epsilon \leq 6$ and $\mu \geq \max\{(16/\epsilon) + 2\sqrt{(64/\epsilon^2) - 1}, \epsilon/4\}$. To obtain the last inequality, we used the fact that $\log(K_{\text{GS}}) \leq K_{\text{GS}}$ for $K_{\text{GS}} > 0$. \square

APPENDIX E NOISE-FREE SUBSPACE MAC

We compute the sum capacity and the relaxed sum capacity for the examples constructed using the noise-free subspace MAC defined in Eq. (66). For both the examples, we consider the input alphabets $\mathcal{A} = \{a_1, a_2\}$, $\mathcal{B} = \{b_1, b_2\}$ and the output alphabet $\mathcal{Z} = \{z_1, z_2\}$.

A. Example 1

For the first example, the MAC $\mathcal{N}_F^{(0)}$ has the probability transition matrix

$$\mathcal{N}_F^{(0)} = \begin{pmatrix} 1 & 0.5 & 0.5 & 0.5 \\ 0 & 0.5 & 0.5 & 0.5 \end{pmatrix}.$$

Let us denote the input probability distribution of sender A as $(p, 1-p)$ corresponding to the symbols (a_1, a_2) . Similarly, denote the input probability distribution of sender B as $(q, 1-q)$ corresponding to the symbols (b_1, b_2) . When the input to the MAC is a product distribution, the output probability distribution determined by the channel is

$$p^Z(z_1) = \frac{1+pq}{2} \text{ and } p^Z(z_2) = \frac{1-pq}{2}.$$

The mutual information $I(A, B; Z)$ between the senders and the receiver is given by

$$I(A, B; Z) = -\frac{1+pq}{2} \ln\left(\frac{1+pq}{2}\right) - \frac{1-pq}{2} \ln\left(\frac{1-pq}{2}\right) - (1-pq)\log(2).$$

We wish to compute the sum capacity

$$S(\mathcal{N}_F^{(0)}) = \max_{0 \leq p, q \leq 1} I(A, B; Z).$$

For $(p, q) \in \{(0, 0), (0, 1), (1, 0), (1, 1)\}$, we have $I(A, B; Z) = 0$, and hence we can focus on the interior of the domain. For maximization over p for a fixed q , by complementary slackness (Lagrangian not written here), we can simply set the derivative of $I(A, B; Z)$ with respect to p equal to 0. The derivative is given as

$$\frac{\partial I}{\partial p} = \frac{q}{2} \ln\left(4 \frac{(1-pq)}{(1+pq)}\right).$$

Setting this equal to zero gives $pq = 3/5$. In this case, no maximization over q is necessary, and subsequently, we find that $S(\mathcal{N}_F^{(0)}) = h(4/5) - (2/5)\ln(2)$, where h is the binary entropy function.

For computing $C(\mathcal{N}_F^{(0)})$, we maximize over all probability distributions over the inputs. For an arbitrary input probability distribution $p(a, b)$, we can write the mutual information between the inputs and the output as

$$I(A, B; Z) = -\frac{1+p(a_1, b_1)}{2} \ln\left(\frac{1+p(a_1, b_1)}{2}\right) - \frac{1-p(a_1, b_1)}{2} \ln\left(\frac{1-p(a_1, b_1)}{2}\right) - (1-p(a_1, b_1))\log(2).$$

In this case, we again find that the maximum is attained at $p(a_1, b_1) = 3/5$, and subsequently, the maximum value of the mutual information is equal to $h(4/5) - (2/5)\ln(2)$ as before. Thus, the capacity corresponding to the relaxed sum region matches the actual sum capacity.

B. Example 2

For the second example, the probability transition matrix is given as

$$\mathcal{N}_F^{(1)} = \begin{pmatrix} 1 & 0.5 & 0.5 & 0 \\ 0 & 0.5 & 0.5 & 1 \end{pmatrix}.$$

As before, denote the input probability distribution of sender A as $(p, 1-p)$ and that of sender B as $(q, 1-q)$. Then, given an input probability distribution to the MAC $\mathcal{N}_F^{(1)}$, the output probability distribution is given by

$$p(z_1) = \frac{p+q}{2} \text{ and } p(z_2) = 1 - \frac{(p+q)}{2}.$$

Then, the mutual information $I(A, B; Z)$ between the senders and the receiver can be written as

$$I(A, B; Z) = -\frac{p+q}{2} \ln\left(\frac{p+q}{2}\right) - \left(1 - \frac{(p+q)}{2}\right) \ln\left(1 - \frac{(p+q)}{2}\right) - (p+q-2pq)\ln(2).$$

We wish to compute the sum capacity $S(\mathcal{N}_F^{(1)}) = \max_{0 \leq p, q \leq 1} I(A, B; Z)$. The edge cases $(p, q = 0, 1)$ will be handled separately. First, we perform the maximization over p for each fixed q ,

$$I^*(q) = \max_{0 \leq p \leq 1} I(A, B; Z).$$

We look for maxima in the interior $(0, 1)$, which can be obtained through $\frac{\partial}{\partial p} I(A, B; Z) = 0$. Note that the output probability is $(1, 0)$ or $(0, 1)$ only when $(p, q) \in \{(0, 0), (1, 1)\}$, and therefore, the derivative of $I(A, B; Z)$ is well-defined in the interior. This derivative is given as

$$\frac{\partial I}{\partial p} = \frac{1}{2} \ln\left(\frac{2-(p+q)}{p+q}\right) - (1-2q)\ln(2).$$

Setting the derivative to zero, we find that $p+q = 2/(\kappa_q + 1)$, where $\kappa_q = 2^{2-4q}$. Therefore, the function $I^*(q)$ can be written as

$$I^*(q) = h\left(\frac{1}{\kappa_q + 1}\right) - \left(\frac{2}{\kappa_q + 1} - 2\left(\frac{2}{\kappa_q + 1} - q\right)q\right)\ln(2),$$

where h is the binary entropy. Now, to compute the sum capacity, we wish to perform the maximization $S(\mathcal{N}_F^{(1)}) = \max_{0 \leq q \leq 1} I^*(q)$. Since $I^*(q)$ is a continuous function of q , the maximum is either attained at the interior or at the boundaries. The maxima in the interior can be found via $\frac{\partial}{\partial q} I^*(q) = 0$:

$$\ln\left(\frac{1}{\kappa_q}\right) = \ln(2) \left[(4q-2) - \frac{1}{\ln(2)} \left(\frac{\kappa_q + 1}{\kappa_q} \right) (q(\kappa_q + 1) - 1) \right]$$

$$\Rightarrow q = \frac{1}{\kappa_q + 1}.$$

It can be verified that $q = 1/2$ is a solution to the above equation, corresponding to which we have $p = 1/2$. Furthermore, this solution is unique. Thus, in the interior of the domain, the maximum occurs at $p = 1/2$, $q = 1/2$. Correspondingly, the maximum value of $I(A, B; Z)$ in the interior is equal to $0.5 \ln(2)$.

Now, we check the edge cases. For $p = 0$, we need to maximize $g(q) = I(p = 0, q) = h(q/2) - q \ln(2)$ over q . At $q = 0, 1$, we have $g(0) = g(1) = 0$. Then, the maximum in the interior can be obtained by setting the derivative with respect to q to zero. This gives $q = 2/5$, and since $g(2/5) \approx 0.223 < 0.5 \ln(2)$, this is not the global maximum. On the other hand, for $p = 1$, we need to maximize $g(q) = I(p = 1, q) = h((1+q)/2) - (1-q) \ln(2)$. At $q = 0, 1$, we have $g(0) = g(1) = 0$. The maximum in the interior can be obtained by setting the derivative of $g(q)$ with respect to q equal to zero. This gives $q = 3/5$, and correspondingly, we have $g(3/5) \approx 0.223 < 0.5 \ln(2)$. Since $I(p, q)$ is symmetric under interchange of p and q , the same results will be obtained when beginning with $q = 0$ or $q = 1$ and then maximizing over p . Therefore, we find that the sum capacity is equal to $S(\mathcal{N}_F^{(1)}) = 0.5 \ln(2)$.

Now, we compute the relaxed sum capacity. For this, note that output probability distribution of the MAC $\mathcal{N}_F^{(1)}$ when using an arbitrary input probability distribution $p(a, b)$ is given as

$$p(z_1) = \frac{1}{2} + \frac{p(a_1, b_1) - p(a_2, b_2)}{2}$$

and

$$p(z_2) = \frac{1}{2} - \frac{(p(a_1, b_1) - p(a_2, b_2))}{2}.$$

Correspondingly, the mutual information between the senders and the receivers is given by

$$I(A, B; Z) = h\left(\frac{1}{2} + \frac{p(a_1, b_1) - p(a_2, b_2)}{2}\right) - (1 - p(a_1, b_1) - p(a_2, b_2)) \ln(2).$$

It can be verified that for $p(a_1, b_1) = p(a_2, b_2) = 1/2$ we have $I(A, B; Z) = \ln(2)$, which is the maximum possible value for the mutual information. Therefore, $C(\mathcal{N}_F^{(1)}) = \ln(2)$.

ACKNOWLEDGMENT

The authors would like to thank Mohammad Alhejji for helpful discussions and comments on the manuscript. Akshay Seshadri would like to thank Stephen Becker for helpful discussions and for suggesting to look into Lipschitz grid search algorithms for nonconvex optimization.

CODE AVAILABILITY

An open source implementation of the codes used in this study can be found at <https://github.com/akshaysehshadri/sum-capacity-computation>.

This repository contains code to perform the following tasks: 1) Maximize Lipschitz-like functions over the interval

(Alg. 1) and standard simplex using grid search (Alg. 2) and dense curves (Alg. 4), 2) Compute the sum capacity of two-sender MACs by implementing the aforementioned algorithms, and 3) Find the maximum winning probability of an N -player non-local game using no-signalling strategies and a corresponding optimal NS strategy.

REFERENCES

- [1] A. Seshadri, F. Leditzky, V. Siddhu, and G. Smith, "On the separation of correlation-assisted sum capacities of multiple access channels," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2022, pp. 2756–2761.
- [2] C. E. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.*, vol. 27, no. 3, pp. 379–423, 1948.
- [3] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Hoboken, NJ, USA: Wiley-Interscience, 2006.
- [4] M. Chiang and S. Boyd, "Geometric programming duals of channel capacity and rate distortion," *IEEE Trans. Inf. Theory*, vol. 50, no. 2, pp. 245–258, Feb. 2004. [Online]. Available: https://web.stanford.edu/boyd/papers/pdf/gp Chan_cap.pdf
- [5] S. Arimoto, "An algorithm for computing the capacity of arbitrary discrete memoryless channels," *IEEE Trans. Inf. Theory*, vol. IT-18, no. 1, pp. 14–20, Jan. 1972.
- [6] R. E. Blahut, "Computation of channel capacity and rate-distortion functions," *IEEE Trans. Inf. Theory*, vol. IT-18, no. 4, pp. 460–473, Jul. 1972.
- [7] R. Ahlswede, "The capacity region of a channel with two senders and two receivers," *Ann. Probab.*, vol. 2, no. 5, pp. 805–814, Oct. 1974.
- [8] H. Liao, "Multiple access channels," Ph.D. thesis, Dept. Elect. Eng., Univ. Hawaii, Honolulu, HI, USA, 1972.
- [9] E. Calvo, D. P. Palomar, J. R. Fonollosa, and J. Vidal, "On the computation of the capacity region of the discrete MAC," *IEEE Trans. Commun.*, vol. 58, no. 12, pp. 3512–3525, Dec. 2010.
- [10] J. Böhler and G. Wunder, "A note on capacity computation for the discrete multiple access channel," *IEEE Trans. Inf. Theory*, vol. 57, no. 4, pp. 1906–1910, Apr. 2011.
- [11] F. Leditzky, M. A. Alhejji, J. Levin, and G. Smith, "Playing games with multiple access channels," *Nature Commun.*, vol. 11, no. 1, pp. 1–5, Mar. 2020.
- [12] Y. Quek and P. W. Shor, "Quantum and superquantum enhancements to two-sender, two-receiver channels," *Phys. Rev. A, Gen. Phys.*, vol. 95, no. 5, May 2017, Art. no. 052329.
- [13] J. Nötzel, "Entanglement-enabled communication," *IEEE J. Sel. Areas Inf. Theory*, vol. 1, no. 2, pp. 401–415, Aug. 2020.
- [14] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge, U.K.: Cambridge Univ. Press, 2000.
- [15] G. Brassard, A. Broadbent, and A. Tapp, "Quantum pseudo-telepathy," *Found. Phys.*, vol. 35, no. 11, pp. 1877–1907, 2005.
- [16] C. Palazuelos and T. Vidick, "Survey on nonlocal games and operator space theory," *J. Math. Phys.*, vol. 57, no. 1, Jan. 2016, Art. no. 015220.
- [17] N. D. Mermin, "Hidden variables and the two theorems of John Bell," *Rev. Modern Phys.*, vol. 65, no. 3, pp. 803–815, Jul. 1993.
- [18] V. Russo, "Extended nonlocal games," Ph.D. dissertation, D. R. Cheriton School Comput. Sci., Univ. Waterloo, Waterloo, ON, Canada, 2017.
- [19] N. D. Mermin, "Simple unified form for the major no-hidden-variables theorems," *Phys. Rev. Lett.*, vol. 65, p. 3373, Jul. 1990.
- [20] A. Peres, "Incompatible results of quantum measurements," *Phys. Lett. A*, vol. 151, nos. 3–4, pp. 107–108, Dec. 1990.
- [21] P. K. Aravind, "A simple demonstration of Bell's theorem involving two observers and no probabilities or inequalities," 2002, *arXiv:quant-ph/0206070*.
- [22] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, "Proposed experiment to test local hidden-variable theories," *Phys. Rev. Lett.*, vol. 23, p. 880, Oct. 1969.
- [23] R. Cleve, P. Hoyer, B. Toner, and J. Watrous, "Consequences and limits of nonlocal strategies," in *Proc. 19th IEEE Annu. Conf. Comput. Complex.*, Jan. 2004, pp. 236–249.
- [24] S. Popescu and D. Rohrlich, "Quantum nonlocality as an axiom," *Found. Phys.*, vol. 24, no. 3, pp. 379–385, 1994.
- [25] G. Brassard, A. Broadbent, and A. Tapp, "Multi-party pseudo-telepathy," in *Proc. Workshop Algorithms Data Struct.* Cham, Switzerland: Springer, 2003, pp. 1–11.

- [26] B. Toner, "Monogamy of non-local quantum correlations," *Proc. Roy. Soc. A, Math., Phys. Eng. Sci.*, vol. 465, no. 2101, pp. 59–69, Jan. 2009.
- [27] E. de Klerk, "The complexity of optimizing over a simplex, hypercube or sphere: A short survey," *Central Eur. J. Oper. Res.*, vol. 16, no. 2, pp. 111–125, Jun. 2008.
- [28] Z. Zhang, "Estimating mutual information via Kolmogorov distance," *IEEE Trans. Inf. Theory*, vol. 53, no. 9, pp. 3280–3282, Sep. 2007.
- [29] K. M. R. Audenaert, "A sharp continuity estimate for the von Neumann entropy," *J. Phys. A, Math. Theoretical*, vol. 40, no. 28, p. 8127, 2007.
- [30] S. A. Piyavskii, "An algorithm for finding the absolute extremum of a function," *USSR Comput. Math. Math. Phys.*, vol. 12, no. 4, pp. 57–67, Jan. 1972.
- [31] B. O. Shubert, "A sequential method seeking the global maximum of a function," *SIAM J. Numer. Anal.*, vol. 9, no. 3, pp. 379–388, 1972.
- [32] D. Lera and Y. D. Sergeyev, "Global minimization algorithms for Hölder functions," *BIT Numer. Math.*, vol. 42, no. 1, pp. 119–133, 2002.
- [33] R. J. Vanderbei, "Extension of Piyavskii's algorithm to continuous global optimization," *J. Global Optim.*, vol. 14, no. 2, pp. 205–216, 1999.
- [34] R. H. Mladineo, "An algorithm for finding the global maximum of a multimodal, multivariate function," *Math. Program.*, vol. 34, no. 2, pp. 188–200, Mar. 1986.
- [35] M. Rahal and A. Ziadi, "A new extension of Piyavskii's method to Hölder functions of several variables," *Appl. Math. Comput.*, vol. 197, no. 2, pp. 478–488, 2008.
- [36] E. De Klerk, D. Den Hertog, and G. Elabwabi, "On the complexity of optimization over the standard simplex," *Eur. J. Oper. Res.*, vol. 191, no. 3, pp. 773–785, Dec. 2008.
- [37] A. Ziadi, Y. Cherrault, and G. Mora, "Global optimization: A new variant of the Alienor method," *Comput. Math. Appl.*, vol. 41, nos. 1–2, pp. 63–71, Jan. 2001.
- [38] O. Fawzi and P. Ferme, "Beating the sum-rate capacity of the binary adder channel with non-signaling correlations," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2022, pp. 2750–2755.
- [39] C. Malherbe and N. Vayatis, "Global optimization of Lipschitz functions," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 2314–2323.
- [40] S. Boyd, S. P. Boyd, and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [41] D. A. Harville, *Matrix Algebra From a Statistician's Perspective*. New York, NY, USA: Springer, 1997.
- [42] J. Sherman and W. J. Morrison, "Adjustment of an inverse matrix corresponding to a change in one element of a given matrix," *Ann. Math. Statist.*, vol. 21, no. 1, pp. 124–127, 1950.
- [43] M. S. Bartlett, "An inverse matrix adjustment arising in discriminant analysis," *Ann. Math. Statist.*, vol. 22, no. 1, pp. 107–111, Mar. 1951.
- [44] J. B. Conway, *A Course in Functional Analysis* (Graduate Texts in Mathematics), vol. 96. New York: Springer, 2007.
- [45] V. Balestro, H. Martini, and R. Teixeira, "Convex analysis in normed spaces and metric projections onto convex bodies," 2019, *arXiv:1908.08742*.
- [46] D. Lera and Y. D. Sergeyev, "Lipschitz and Hölder global optimization using space-filling curves," *Appl. Numer. Math.*, vol. 60, nos. 1–2, pp. 115–129, 2010.
- [47] Y. D. Sergeyev, R. G. Strongin, and D. Lera, *Introduction to Global Optimization Exploiting Space-Filling Curves*. Cham, Switzerland: Springer, 2013.
- [48] F. Topsøe, "Bounds for entropy and divergence for distributions over a two-element set," *J. Inequalities Pure Appl. Math.*, vol. 2, no. 2, p. 300, 2001.

Akshay Seshadri received the B.Tech. and M.Tech. degrees in aerospace engineering from the Indian Institute of Technology Madras. He is currently pursuing the Ph.D. degree in physics with the University of Colorado Boulder. His research interests include quantum information theory, probability and statistics, convex and non-convex optimization, and dynamical systems theory.

Felix Leditzky received the Ph.D. degree from the University of Cambridge in 2017. From 2016 to 2019, he was a Post-Doctoral Research Associate at the JILA, University of Colorado Boulder. From 2019 to 2020, he was a Joint Post-Doctoral Fellow at the Institute for Quantum Computing, University of Waterloo, and the Perimeter Institute of Theoretical Physics, Waterloo, ON, Canada. He is currently an Assistant Professor in the Department of Mathematics at the University of Illinois Urbana–Champaign, and a member of the Illinois Quantum Information Science and Technology Center. His research interests include mathematical and information-theoretic aspects of quantum information theory.

Vikesh Siddhu received the Ph.D. degree in physics from Carnegie Mellon University (CMU), USA, in 2020. He held post-doctoral positions at the CMU's Tepper School of Business, JILA–University of Colorado/NIST, and the IBM Quantum where he works presently. His research interests include quantum information theory, (non) convex optimization, and related subjects.

Graeme Smith received the Ph.D. degree in theoretical physics from Caltech in 2006. He was a Research Associate in computer science with the University of Bristol from 2006 to 2007. He was a Research Staff Member from 2010 to 2016 and a Post-Doctoral Researcher from 2007 to 2010 at IBM Thomas J. Watson Research Center. He is currently an Associate Professor of physics at the University of Colorado Boulder and a JILA Fellow. His research interests include computation and communication in noisy settings, quantum Shannon theory, coding theory, cryptography, quantum estimation and detection, optical communications, and the physics of information. He was awarded an NSF CAREER Award in 2016.