

Analysis and optimization of quantum adaptive measurement protocols with the framework of preparation games

M. Weilenmann ^{1✉}, E. A. Aguilar ^{1,2✉} & M. Navascués^{1✉}

A preparation game is a task whereby a player sequentially sends a number of quantum states to a referee, who probes each of them and announces the measurement result. Many experimental tasks in quantum information, such as entanglement quantification or magic state detection, can be cast as preparation games. In this paper, we introduce general methods to design n -round preparation games, with tight bounds on the performance achievable by players with arbitrarily constrained preparation devices. We illustrate our results by devising new adaptive measurement protocols for entanglement detection and quantification. Surprisingly, we find that the standard procedure in entanglement detection, namely, estimating n times the average value of a given entanglement witness, is in general suboptimal for detecting the entanglement of a specific quantum state. On the contrary, there exist n -round experimental scenarios where detecting the entanglement of a known state optimally requires adaptive measurement schemes.

¹Institute for Quantum Optics and Quantum Information (IQOQI) Vienna Austrian Academy of Sciences, Vienna, Austria. ²AIT Austrian Institute of Technology GmbH, Vienna, Austria. ✉email: mirjam.weilenmann@oeaw.ac.at; ed.alex.aguilar@gmail.com; miguel.navascues@oeaw.ac.at

Certain tasks in quantum communication can only be conducted when all the parties involved share a quantum state with a specific property. For instance, two parties with access to a public communication channel must share an entangled quantum state in order to generate a secret key¹. If the same two parties wished to carry out a qudit teleportation experiment, then they would need to share a quantum state with an entanglement fraction beyond $1/d^2$. More generally, when only restricted quantum operations are permitted, specific types of quantum states become instrumental for completing certain information processing tasks. This is usually formalized in terms of resource theories³. Some resources, like entanglement, constitute the basis of quantum communication. Others, such as magic states, are required to carry out quantum computations⁴. Certifying and quantifying the presence of resource states with a minimum number of experiments is the holy grail of entanglement⁵ and magic state detection⁴.

Beyond the problem of characterizing resourceful states mathematically, the experimental detection and quantification of resource states is further complicated by the lack of a general theory to devise efficient measurement protocols. Such protocols would allow one to decide, at minimum experimental cost, whether a source is capable of producing resourceful states. Developing such methods is particularly important for high dimensional systems where full tomography is infeasible or in cases where the resource states to be detected are restricted to a small (convex) subset of the state space, which renders tomography excessive.

General results on the optimal discrimination between different sets of states in the asymptotic regime⁶ suggest that the optimal measurement protocol usually involves collective measurements over many copies of the states of interest, and thus would require a quantum memory for its implementation. This contrasts with the measurement scenario encountered in many experimental setups: the lack of a quantum memory often forces an experimentalist to measure each of the prepared states as soon as they arrive at the lab. In this case it is natural to consider a setting where subsequent measurements can depend on previous measurement outcomes, in which case the experimentalist is said to follow an adaptive strategy. Perhaps due to their perceived complexity, the topic of identifying optimal adaptive measurement strategies has been largely overlooked in quantum information theory.

In this paper, we propose the framework of quantum preparation games to reason about the detection and quantification of resource states in this adaptive setting. These are games wherein a player will attempt to prepare some resource, which the referee will measure and subsequently assign a score to. We prove a number of general results on preparation games, including the efficient computation of the maximum average score achievable by various types of state preparation strategies. Our results furthermore allow us to optimise over the most general measurement strategies one can follow with only a finite set of measurements, which we term Maxwell demon games. Due to limited computational resources, full optimisations over Maxwell demon games are restricted to scenarios with only $n \approx 3, 4$ rounds. For higher round numbers, say, $n \approx 20$, we propose a heuristic, based on coordinate descent, to carry these optimisations out approximately. More specifically, the outcome of the heuristic is (in general) a sub-optimal preparation game that nonetheless satisfies all the optimization constraints. In addition, we show how to devise arbitrarily complex preparation games through game composition, and yet another heuristic inspired by gradient descent. We illustrate all our techniques with examples from entanglement certification and quantification and highlight the benefit of adaptive measurement strategies in various ways. In

this regard, in contradiction to standard practice in entanglement detection, we find that the optimal n -round measurement protocol to detect the entanglement of a single, known quantum state does not consist in estimating n times the value of a given (optimised) entanglement witness. On the contrary, there exist adaptive measurement schemes that supersede any non-adaptive protocol for this task.

Results

In the following we introduce the framework of preparation games and show how to devise resource detection protocols using this framework. We illustrate our findings with applications from entanglement detection and quantification.

Quantum preparation games for resource certification and quantification. Consider the following task: a source is distributing multipartite quantum states among m separate parties who wish to quantify how entangled those states are. To this effect, the parties sequentially probe a number of m -partite states prepared by the source. Depending on the results of each experiment, they decide how to probe the next state. After a fixed number of rounds, the parties estimate the entanglement of the probed states. They naturally seek an estimate that lower bounds the actual entanglement content of the states produced during the experiment with high probability. Most importantly, if the source is unable to produce entangled states, the protocol should certify this with high probability.

Experimental scenarios whereby a source (or player) sequentially prepares quantum states that are subject to adaptive measurements (by some party or set of parties that we collectively call the referee) are quite common in quantum information. Besides entanglement detection, they are also found in magic state certification⁴, and, more generally, in the certification and quantification of any quantum state resource³. The common features of these apparently disparate quantum information processing tasks motivate the definition of quantum preparation games. See Box 1 and Fig. 1.

A preparation game G is thus fully defined by the triple (S, M, g) , where S denotes the sequence of game configuration sets $(S_k)_{k=1}^{n+1}$, and M , the set of POVMs $M \equiv \{M_{s|s}^{(k)} : s' \in S_{k+1}, s \in S_k\}_{k=1}^n$. In principle, the Hilbert space where the state prepared in round k lives could depend on k and on the current game configuration $s_k \in S_k$. For simplicity, though, we will assume that all prepared states act on the same Hilbert space, \mathcal{H} . In many practical situations, the actual, physical measurements conducted by the referee in round k will have outcomes in O , with $|O| < |S_k|$. The new game configuration $s' \in S_{k+1}$ is thus decided by the referee through some non-deterministic function of the current game configuration s and the ‘physical’ measurement outcome $o \in O$. The definition of the game POVM $\{M_{s|s}^{(k)}\}_s$ encompasses this classical processing of the physical measurement outcomes.

The expected score of a player with preparation strategy \mathcal{P} is

$$G(\mathcal{P}) \equiv \sum_{s \in S_{n+1}} p(s|\mathcal{P}, G)g(s). \quad (1)$$

In the equation, $p(s|\mathcal{P}, G)$ denotes the probability that, conditioned on the player using a preparation strategy \mathcal{P} in the game G , the final game configuration is s . For the sake of clarity, we will sometimes refer to the set of possible final configurations as \bar{S} instead of S_{n+1} .

In this paper we consider players who aim to maximise their expected score over all preparation strategies \mathcal{P} that are accessible to them, in order to convince the referee of their ability to prepare a desired resource. Intuitively, a preparation strategy is the policy

Box 1 : Quantum preparation game

Game Variables

Number of Rounds	n
Game Configuration	Unique initial configuration $S_1 = \{\emptyset\}$. In round k , there is a set of allowed configurations $S_k = \{s_1^k, s_2^k, \dots\}$. After n rounds, the game ends in $s \in S_{n+1}$.
Measurement Operators	For every game configuration $s \in S_k$, there are POVMs $\{M_{s' s}^{(k)} : s' \in S_{k+1}\}$.
Scoring Function	A (non-deterministic) function $g : S_{n+1} \rightarrow \mathbb{R}$.

The game variables, i.e., round number, possible configurations, POVMs and scoring rule, are publicly announced before the game starts.

Measurement Round Rules

1. At the beginning of round k , the current game configuration $s \in S_k$ is known to the player. The player prepares a state ρ^k according to their preparation strategy \mathcal{P} , and sends it to the referee.
2. The referee measures the quantum state ρ^k with the POVM $\{M_{s'|s}^{(k)} : s' \in S_{k+1}\}$.
3. The referee publicly announces the outcome s' of this measurement, which becomes the game configuration for the next round.

Scoring

After the n^{th} round, the player receives a score $g(s)$, where $s \in S_{n+1}$ is the final configuration.

that a player follows to decide, in each round, which quantum state to prepare. Since the player has access to the referee’s current game configuration, the player’s state preparation can depend on this. The simplest preparation strategy, however, consists in preparing independent and identically distributed (i.i.d.) copies of the same state ρ . We call such preparation schemes i.i.d. strategies and denote them as $\rho^{\otimes n}$. A natural extension of i.i.d. strategies, which we call finitely correlated strategies⁷, follows when we consider interactions with an uncontrolled environment, see Fig. 2. I.i.d. and finitely correlated strategies can be extended to scenarios where the preparation depends on the round number k . The mathematical study of these strategies is so similar to that of their round-independent counterparts, that we will not consider such extensions in this article.

Instead, we will analyze more general scenarios, where the player is limited to preparing multipartite states belonging to a specific class \mathcal{C} , e.g. separable states. In this case, given $\rho, \sigma \in \mathcal{C} \cap B(\mathcal{H})^{\otimes k}$, a player can also generate the state $p\rho + (1 - p)\sigma$ for any $p \in [0, 1]$, just by preparing ρ with probability p and σ otherwise. Thus, we can always assume $\mathcal{C} \cap B(\mathcal{H})^{\otimes k}$ to be convex for all k . The preparation strategies of such a player will be assumed fully general, e.g., the state preparation in round k can depend on k , or on the current game configuration s_k . We call such strategies \mathcal{C} -constrained.

Computing the average score of a preparation game. Even for i.i.d. strategies, a brute-force computation of the average game score would require adding up a number of terms that is exponential in the number of rounds. In the following we introduce a method to efficiently compute the average game scores for various types of player strategies.

Let $G = (S, M, g)$ be a preparation game with $M \equiv \{M_{s'|s}^{(k)} : s' \in S_{k+1}, s \in S_k\}_{k=1}^n$, and let \mathcal{C} be a set of quantum states. In principle, a \mathcal{C} -constrained player could exploit correlations between the states they prepare in different rounds to increase their average score when playing G . They could, for instance, prepare a bipartite state $\rho^{12} \in \mathcal{C}$; send part 1 to the referee in round 1 and, depending on the referee’s measurement outcome s_2 , send part 2, perhaps after acting on it with a completely positive map depending on s_2 . However, the player would be in exactly the same situation if, instead, they sent state $\rho^1 = \text{tr}_2(\rho^{12})$

in round 1 and state $\rho_{s_2}^2 \propto \text{tr}_1[(M_{s_2|\emptyset} \otimes \mathbb{I}_2)\rho^{12}]$ in round 2. There is a problem, though: the above is only a \mathcal{C} -constrained preparation strategy provided that $\rho_{s_2}^2 \in \mathcal{C}$. This motivates us to adopt the following assumption.

Assumption 1

The set of (in principle, multipartite) states \mathcal{C} is closed under arbitrary postselections with the class of measurements conducted by the referee.

This assumption holds for general measurements when \mathcal{C} is the set of fully separable quantum states or the set of states with entanglement dimension⁸ at most D (for any $D > 1$). It also holds when \mathcal{C} is the set of non-magic states and the referee is limited to conducting convex combinations of sequential Pauli measurements⁹. More generally, the assumption is satisfied when, for some convex resource theory³, \mathcal{C} is the set of resource-free states; and the measurements of the referee are resource-free. The assumption is furthermore met when the player does not have a quantum memory.

Under Assumption 1, the player’s optimal \mathcal{C} -constrained strategy consists in preparing a state $\rho_{s_k}^k \in \mathcal{C}$ in each round k , depending on k and the current game configuration s_k . Now, define $\mu_s^{(k)}$ as the maximum average score achieved by a player, conditioned on s being the configuration in round k . Then $\mu_s^{(k)}$ satisfies

$$\begin{aligned} \mu_s^{(n)} &= \max_{\rho \in \mathcal{C}} \sum_{\bar{s} \in \mathcal{S}} \text{tr}[M_{\bar{s}|s}^{(n)} \rho](g(\bar{s})), \\ \mu_s^{(k)} &= \max_{\rho \in \mathcal{C}} \sum_{s'} \text{tr}[M_{s'|s}^{(k)} \rho] \mu_{s'}^{(k+1)}. \end{aligned} \tag{2}$$

These two relations allow us to inductively compute the maximum average score achievable via \mathcal{C} -constrained strategies, $\mu_{\emptyset}^{(1)}$. Note that, if the optimizations above were carried out over a larger set of states $\mathcal{C}' \supset \mathcal{C}$, the end result would be an upper bound on the achievable maximum score. This feature will be handy when \mathcal{C} is the set of separable states, since the latter is difficult to characterize exactly^{10,11}. In either case, the computational resources to conduct the computation above scale as $O(\sum_k |S_k| |S_{k+1}|)$.

Equation (2) can also be used to compute the average score of an i.i.d. preparation strategy $\rho^{\otimes n}$. In that case, $\mathcal{C} = \{\rho\}$, and the

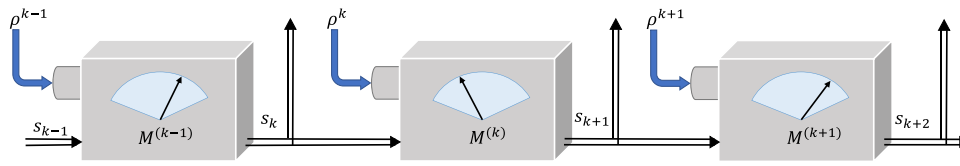


Fig. 1 Quantum preparation game from the referee’s perspective. In each round k of a preparation game, the referee (measurement box) receives a quantum state ρ^k from the player. The referee’s measurement $M^{(k)}$ will depend on the current game configuration s_k , which is determined by the measurement outcome of the previous round. In the same way, the outcome s_{k+1} of round k will determine the POVMs to be used in round $k + 1$. Recall that the player can tailor the states ρ^k to the measurements to be performed in round k , since they have access to the (public) game configuration s_k , shown with the upward arrow leaving the measurement apparatus.

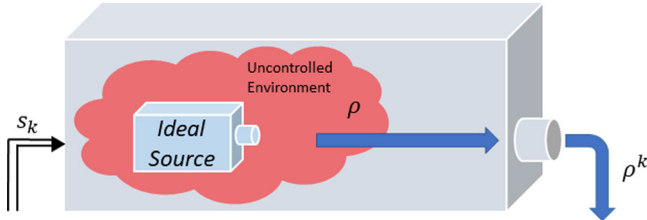


Fig. 2 Finitely correlated strategies. Suppose that a player owns a device which allows them to prepare and distribute a quantum state to the referee. Unfortunately, at each experimental preparation the player’s device interacts with an environment A . Explicitly, if the player activates their device, then the referee receives the state $\text{tr}_A[\sum_i K_i \rho K_i^\dagger]$, where ρ is the current state of the environment and $K_i : \mathcal{H}_A \rightarrow \mathcal{H}_A \otimes \mathcal{H}$ are the Kraus operators which evolve the environment and prepare the state that the referee receives. Since the same environment is interacting with each prepared state, the states that the referee receives in different rounds are likely correlated.

maximization over \mathcal{C} is trivial. Similarly, an adaptation of (2) allows us to efficiently compute the average score of finitely correlated strategies, for the details we refer to the Methods.

Optimizing preparation games. Various tasks in quantum information – including entanglement detection – have the following structure: given two sets of preparation strategies $\mathcal{S}, \mathcal{S}'$ and a score function g , we want to find a game $G = (\mathcal{S}, M, g)$ that separates these two sets, i.e., a game such that $G(\mathcal{P}) \leq \delta$, for all $\mathcal{P} \in \mathcal{S}$, and $G(\mathcal{P}) > \delta$ for all $\mathcal{P} \in \mathcal{S}'$. In some cases, we are interested to search for games where the POVMs conducted by the referee belong to a given (convex) class \mathcal{M} . This class represents the experimental limitations affecting the referee, such as space-like separation or the unavailability of a given resource.

Finding a preparation game satisfying the above constraints can be regarded as an optimization problem over the set of quantum preparation games. Consider a set \mathcal{M} of adaptive measurement protocols of the form $M \equiv \{M_{s'_k|s_k}^{(k)} : s'_k \in S_{k+1}, s_k \in S_k\}$, a selection of preparation games $\{G_M^i = (\mathcal{S}, M, g^i)\}_{i=1}^r$ and sets of preparation strategies $\{\mathcal{S}_i\}_{i=1}^r$. A general optimization over the set of quantum preparation games is a problem of the form

$$\begin{aligned} & \min_{M \in \mathcal{M}, \mathbf{v}} f(\mathbf{v}) \\ \text{s.t. } & G_M^i(\mathcal{P}) \leq v_i, \forall \mathcal{P} \in \mathcal{S}_i, i = 1, \dots, r, \\ & A \cdot \mathbf{v} \leq \mathbf{b}, \end{aligned} \tag{3}$$

where A, \mathbf{b} are a $t \times r$ matrix and a vector of length t , respectively, and $f(\mathbf{v})$ is assumed to be convex on the vector $\mathbf{v} \in \mathbb{R}^r$.

In this paper, we consider i.i.d., finitely correlated (with known or unknown environment state) and \mathcal{C} -constrained preparation strategies. The latter class also covers scenarios where a player wishes to play an i.i.d. strategy with an imperfect preparation

device. Calling ρ the ideally prepared state, one can model this contingency by assuming that, at every use, the preparation device (adversarially) produces a quantum state ρ' such that $\|\rho - \rho'\|_1 \leq \epsilon$. If, independently of the exact states prepared by the noisy or malfunctioning device, we wish the average score g^i to lie below some value v_i , then the corresponding constraint is

$$G_M^i(\mathcal{P}) \leq v_i, \forall \mathcal{P} \in \mathcal{E}_\epsilon, \tag{4}$$

where \mathcal{E}_ϵ is the set of ϵ -constrained preparation strategies, producing states in $\{\rho' : \rho' \geq 0, \text{tr}(\rho') = 1, \|\rho' - \rho\|_1 \leq \epsilon\}$.

The main technical difficulty in solving problem (3) lies in expressing conditions of the form

$$G_M(\mathcal{P}) \leq v, \forall \mathcal{P} \in \mathcal{S} \tag{5}$$

in a convex (and tractable) way. This will, in turn, depend on which type of measurement protocols we wish to optimize over. We consider families of measurement strategies M such that the matrix

$$\sum_{s_2, \dots, s_{n+1}} (g(s_{n+1})) \bigotimes_{k=1}^n (M_{s_{k+1}|s_k}^k \otimes |s_{k+1}\rangle) \tag{6}$$

depends affinely on the optimization variables of the problem. For $\mathcal{S} = \{\mathcal{P}\}$, condition (5) then amounts to enforcing an affine constraint on the optimization variables defining the referee’s measurement strategy. For finitely correlated strategies, we describe in the Methods how to phrase (5) as a convex constraint.

For \mathcal{C} -constrained strategies, the way to express (5) as a convex constraint depends more intricately on the class of measurements we aim to optimize over. Let us first consider preparation games with $n = 1$ round, where we allow the referee to conduct any $|\mathcal{S}|$ -outcome measurement from the convex set \mathcal{M} . Let \mathcal{S} represent the set of all \mathcal{C} -constrained preparation strategies, for some convex set of states \mathcal{C} . Then, condition (5) is equivalent to

$$v\mathbb{I} - \sum_{s \in \mathcal{S}} M_{s|\emptyset}^{(1)}(g(s)) \in \mathcal{C}^*. \tag{7}$$

Note that, if we replace \mathcal{C}^* in (7) by a subset thereof, relation (5) is still implied. In that case, however, there may be values of v for which relation (5) holds, but not eq. (7). As we will see later, this observation allows us to devise sound entanglement detection protocols, in spite of the fact that the dual of the set of separable states is difficult to pin down^{10,11}.

Next, we consider a particularly important family of multi-round measurement schemes, which we call Maxwell demon games. In a Maxwell demon game, the referee’s physical measurements in each round k are taken from a discrete set $\mathcal{M}(k)$. Namely, for each k , there exist sets of natural numbers A_k, X_k and fixed POVMs $\{(N_{a|x}^{(k)} : a \in A_k) : x \in X_k\} \subset B(\mathcal{H})$. The configuration space at stage k corresponds to the complete history of physical inputs x_1, \dots, x_{k-1} and outputs a_1, \dots, a_{k-1} , i.e., $s_k = (a_1, x_1, \dots, a_{k-1}, x_{k-1})$, where $s_1 = \emptyset$. Note that the cardinality of S_k grows exponentially with k . In order to decide which physical setting x_k must be measured in round k , the referee receives

advice from a Maxwell demon. The demon, who holds an arbitrarily high computational power and recalls the whole history of inputs and outputs, samples x_k from a distribution $P_k(x_k|s_k)$. The final score of the game $\gamma \in \mathcal{G}$ is also chosen by the demon, through the distribution $P(\gamma|s_{n+1})$. A Maxwell demon game is the most general preparation game that a referee can run, under the reasonable assumption that the set of experimentally available measurement settings is finite.

Let us consider

$$P(x_1, \dots, x_n, \gamma|a_0, a_1, \dots, a_n) = P(\gamma|s_{n+1}) \prod_{k=1}^n P_k(x_k|s_k), \quad (8)$$

where $a_0 = \emptyset$. Define $(y_0, \dots, y_n) \equiv (x_1, \dots, x_n, \gamma)$. As shown in¹², a collection of normalized distributions $P(y_0, \dots, y_n|a_0, \dots, a_n)$ admits a decomposition of the form (8) iff the no-signalling-to-the-past conditions

$$\sum_{y_{k+1}, \dots, y_n} P(y_0, \dots, y_n|a_0, \dots, a_n) = P(y_0, \dots, y_k|a_0, \dots, a_k) \quad \forall k \quad (9)$$

hold. For completeness, the reader can find a proof in the Methods. We can thus characterize general Maxwell demon games through finitely many linear constraints on $P(x_1, \dots, x_n, \gamma|a_0, a_1, \dots, a_n)$.

For Maxwell demon games, the matrix (6) depends linearly on the optimization variables $P(x_1, \dots, x_n, \gamma|a_0, a_1, \dots, a_n)$. Hence, we can express condition (5) as a tractable convex constraint whenever \mathcal{S} corresponds to an i.i.d. strategy, or a finitely correlated strategy with an unknown initial environment state, as described above. Enforcing (5) for \mathcal{C} -constrained strategies requires regarding the quantities $\{\mu_s^{(n)}\}_s$ in eq. (2) as optimization variables, related to P and to each other through a dualized version of the conditions (2). The reader can find a full explanation in the Methods.

Finally, we consider the set of adaptive measurement schemes with fixed POVM elements $\{M_{s|s}^{(j)} : j \neq k\}$ and variable $\{M_{s|s}^{(k)}\} \subset \mathcal{M}$, for some tractable convex set of measurements \mathcal{M} . As in the two previous cases, the matrix (6) is linear in the optimization variables $\{M_{s|s}^{(k)}\}$, so (5) can be expressed in a tractable, convex form for sets of finitely-many strategies and finitely correlated strategies with unknown environment. Similarly to the case of Maxwell demon games, enforcing (5) for \mathcal{C} -constrained strategies requires promoting $\{\mu_s^{(j)} : j \leq k\}$ to optimization variables, see the Methods for details.

Via coordinate descent, this observation allows us to conduct optimizations (3) over the set of all adaptive schemes with a fixed game configuration structure $(S_j)_{j=1}^{n+1}$. Consider, indeed, the method presented in Box 2. With this algorithm, at each iteration, the objective value $f(\mathbf{v})$ in problem (3) can either decrease or stay the same: The hope is that it returns a small enough value f^* after a moderate number L of iterations. In the Methods the reader can find a successful application of this heuristic to devise 20-round quantum preparation games.

The main drawback of this algorithm is that it is very sensitive to the initial choice of POVMs, so it generally requires several random initializations to achieve a reasonably good value of the objective function. It is therefore suitable for optimizations of $n \approx 50$ round measurement schemes. Optimizations over, say, $n = 1000$ round games risk getting stuck in a bad local minimum.

To address this issue, we provide two additional methods for the design of large- n quantum preparation games below.

Devising large-round preparation games by composition. The simplest way to construct preparation games with arbitrary round number consists in playing several preparation games, one after another. Consider thus a game where, in each round and depending on the current game configuration, the referee chooses a preparation game. Depending on the outcome, the referee changes the game configuration and plays a different preparation game with the player in the next round. We call such a game a meta-preparation game. Similarly, one can define meta-meta preparation games, where, in each round, the referee and the player engage in a meta-preparation game. This recursive construction can be repeated indefinitely.

In the Methods we show that the maximum average score of a (meta) j -game, which refers to a game at level j of the above recursive construction, can be computed inductively, through a formula akin to Eq. (2). Moreover, in the particular case that the preparation games that make up the (meta) j -game have $\{0, 1\}$ scores, one only needs to know their minimum and maximum scores to compute the (meta) j -game's maximum average score.

For simple meta-games such as play m times the $\{0, 1\}$ -scored preparation game G , count the number of wins and output 1 (0) if it is greater than or equal to (smaller than) a threshold ν , which we denote $G_\nu^{(m)}$, we find that the optimal meta-strategy for the player is to always play G optimally, thus recovering

$$p(G, \nu, m) \equiv \max_{\mathcal{P} \in \mathcal{S}} G_\nu^{(m)}(\mathcal{P}) = \sum_{k=\nu}^m \binom{m}{k} G(\mathcal{P}^*)^k (1 - G(\mathcal{P}^*))^{m-k}, \quad (10)$$

where $\mathcal{P}^* = \arg \max_{\mathcal{P} \in \mathcal{S}} G(\mathcal{P})$, from¹³. $p(G, \nu, m)$ can be interpreted as a p -value for \mathcal{C} -constrained strategies, as it measures the probability of obtaining a result at least as extreme as the observed data ν under the hypothesis that the player's strategies are constrained to belong to \mathcal{S} .

Devising large-round preparation games based on gradient descent. A more sophisticated alternative to devise many-round quantum preparation games exploits the principles behind Variational Quantum Algorithms¹⁴. These are used to optimize the parameters of a quantum circuit by following the gradient of an operator average. Similarly, we propose a gradient-based method to identify the optimal linear witness for detecting certain quantum states. Since the resulting measurement scheme is adaptive, the techniques developed so far are crucial for studying its vulnerability with respect to an adaptive preparation attack.

Box 2 : A heuristic for general optimizations over preparation games

1. Starting point: a natural number L , an optimization problem of the form (3), a sequence of sets of game configurations $S = (S_j)_{j=1}^{n+1}$, a measurement scheme $M = \{M_{s_{j+1}|s_j}^{(j)} : s_{j+1}, s_j\}$ such that $G_M^i(\mathcal{P}) \leq \nu_i, \forall \mathcal{P} \in \mathcal{S}_i, \text{ for } i = 1, \dots, r, \text{ with } A \cdot \mathbf{v} \leq \mathbf{b}$.
2. Set $l = 1$.
3. Choose an index $k \in \{1, \dots, n\}$ and, using the techniques explained in Methods (Eq. (61)), minimize the objective value of (3) over measurement schemes \tilde{M} with $\tilde{M}_{s_{j+1}|s_j}^{(j)} = M_{s_{j+1}|s_j}^{(j)}$ for all $s_j \in S_j, s_{j+1} \in S_{j+1}, j \neq k$, subject to the optimization constraints. Call f^* the objective value of the optimal measurement scheme M^* .
4. Update M and l as $M \leftarrow M^*, l \leftarrow l + 1$. If $l \geq L$, return M and f^* and stop. Otherwise, go to step 3.

Consider a set of i.i.d. preparation strategies $\mathcal{E} = \{\rho^{\otimes n} : \rho \in E\}$, and let $\{\|W(\theta)\| \leq 1 : \theta \in \mathbb{R}^m\} \subset B(\mathcal{H})$ be a parametric family of operators such that $\|\frac{\partial}{\partial \theta_x} W(\theta)\| \leq K$, for $x = 1, \dots, m$. Given a function $f : \mathbb{R}^{m+1} \rightarrow \mathbb{R}$, we wish to devise a preparation game that, ideally, assigns to each strategy $\rho^{\otimes n} \in \mathcal{E}$ an average score of

$$f(\theta_\rho, \text{tr}[W(\theta_\rho)\rho]), \tag{11}$$

with

$$\theta_\rho = \text{argmax}_\theta \text{tr}[W(\theta)\rho]. \tag{12}$$

Intuitively, $W(\theta_\rho)$ represents the optimal witness to detect some property of ρ , and both the average value of $W(\theta_\rho)$ and the value of θ_ρ hold information regarding the use of ρ as a resource.

Next, we detail a simple heuristic to devise preparation games G whose average score approximately satisfies Eq. (11). If, in addition, $f(\theta_\rho, \text{tr}[W(\theta_\rho)\rho]) \leq \delta$ for all $\rho \in \mathcal{C}$, then one would expect that $G(\mathcal{P}) \lesssim \delta$ for all \mathcal{C} -constrained strategies $\mathcal{P} \in \mathcal{S}$.

Fix the quantities $\epsilon > 0$, $\theta_0 \in \mathbb{R}^m$ and the probability distributions $\{p_k(x) : x \in \{0, 1, \dots, m\}\}$, for $k = 1, \dots, n$. For $x = 1, \dots, m$, let $\{M_a^x(\theta) : a = -1, 1\}$ be a POVM such that

$$M_1^x(\theta) - M_{-1}^x(\theta) = \frac{1}{K} \frac{\partial}{\partial \theta_x} W(\theta). \tag{13}$$

Similarly, let $\{M_{-1}^0(\theta), M_1^0(\theta)\}$ be a POVM such that

$$M_1^0(\theta) - M_{-1}^0(\theta) = W(\theta). \tag{14}$$

A gradient-based preparation game is given by the following.

1. The possible game configurations are vectors from the set $S_k = \{-(k-1), \dots, k-1\}^{m+1}$, for $k = 1, \dots, n$. Given $s_k \in S_k$, we will denote by \tilde{s}_k the vector that results when we erase the first entry of s_k .
2. At round k , the referee samples the random variable $x \in \{0, 1, \dots, m\}$ from $p_k(x)$. The referee then implements the physical POVM $M_a^x(\theta_k)$, with $\theta_k = \theta_0 + \epsilon \tilde{s}_k$, obtaining the result $a_k \in \{-1, 1\}$. The next game configuration is $s_{k+1} = s_k + a_k |x\rangle$.
3. The final score of the game is $f(\theta_n, \frac{s_n^0}{\sum_{k=1}^n p_k(0)})$.

More sophisticated variants of this game can, for instance, let ϵ depend on k , or take POVMs with more than two outcomes into account. It is worth remarking that, for fixed m , the number of possible game configurations scales with the total number of rounds n as $O(n^{m+1})$.

If the player uses an i.i.d. strategy, then the sequence of values $(\theta_k)_k$ reflects the effect of applying stochastic gradient descent¹⁵ to solve the optimization problem (12). Hence, for the i.i.d. strategy $\rho^{\otimes n}$ and $n \gg 1$, one would expect the sequence of values $(\theta_k)_k$ to converge to θ_ρ , barring local maxima. In that case, the average score of the game will be close to (11) with high probability. For moderate values of n , however, it is difficult to anticipate the average game scores for strategies in \mathcal{E} and \mathcal{S} , so that a detailed analysis with the procedure from eq. (2) becomes necessary (see the applications below for an example).

Applications: Entanglement certification as a preparation game.

A paradigmatic example of a preparation game is entanglement detection. In this game, the player is an untrusted source of quantum states, while the role of the referee is played by one or more separate parties who receive the states prepared by the source. The goal of the referee is to make sure that the source has indeed the capacity to distribute entangled states. The final score of the entanglement detection preparation game is either 1 (certified entanglement) or 0 (no entanglement certified), that is, $g : \bar{S} \rightarrow \{0, 1\}$. In this case, one can identify the final game

configuration with the game score, i.e., one can take $\bar{S} = \{0, 1\}$. The average game score is then equivalent to the probability that the referee certifies that the source can distribute entangled states.

Consider then a player who is limited to preparing separable states, i.e., a player for whom \mathcal{C} corresponds to the set of fully separable states. Call the set of preparation strategies available to such a player \mathcal{S} . Ideally, we wish to implement a preparation game such that the average game score of a player using strategies from \mathcal{S} (i.e., the probability that the referee incorrectly labels the source as entangled) is below some fixed value ϵ_I . In hypothesis testing, this quantity is known as type-I error. At the same time, if the player follows a class \mathcal{E} of preparation strategies (involving the preparation of entangled states), the probability that the referee incorrectly labels the source as separable is upper bounded by ϵ_{II} . This latter quantity is called type-II error. In summary, we wish to identify a game G such that $p(1|\mathcal{P}) \leq \epsilon_I$, for all $\mathcal{P} \in \mathcal{S}$, and $p(0|\mathcal{P}) \leq \epsilon_{II}$, for all $\mathcal{P} \in \mathcal{E}$.

In the following, we consider three types of referees, with access to the following sets of measurements:

1. Global measurements: \mathcal{M}_1 denotes the set of all bipartite POVMs.
2. 1-way Local Pauli measurements and Classical Communication (LPCC): \mathcal{M}_2 is the set of POVMs conducted by two parties, Alice and Bob, on individual subsystems, where Alice may perform a Pauli measurement first and then, depending on her inputs and outputs, Bob chooses a Pauli measurement as well. The final outcome is a function of both inputs and outcomes.
3. Local Pauli measurements: \mathcal{M}_3 contains all POVMs where Alice and Bob perform Pauli measurements x, y on their subsystems, obtaining results a, b , respectively. The overall output is $\gamma = f(a, b, x, y)$, where f is a (non-deterministic) function.

Few-round protocols for entanglement detection. We first consider entanglement detection protocols with just a single round ($n = 1$). Let $E = \{\rho_1, \dots, \rho_{r-1}\}$ be a set of $r - 1$ bipartite entangled states. Our objective is to minimise the type-II error, given a bound ϵ_I on the acceptable type-I error. To express this optimization problem as in (3), we define $\mathcal{S}_i \equiv \{\rho_i\}$, for $i = 1, \dots, r - 1$, and $\mathcal{S}_r \equiv \mathcal{S}$, the set of separable strategies. In addition, we take $f(\mathbf{v}) = v_1$ and choose A, \mathbf{b} so that $v_r = \epsilon_I, v_1 = \dots = v_{r-1}$. Finally, we consider complementary score functions $g, g' : \bar{S} \rightarrow \{0, 1\}$ and assign the scores $g^i = g$ for $i = 1, \dots, r - 1$, and $g^r = g'$. All in all, the problem to solve is

$$\begin{aligned} & \min_{(M_{s|\emptyset}^{(1)})_{s, \epsilon_{II}}} \epsilon_{II} \\ \text{s.t.} \quad & \text{tr}(M_{0|\emptyset}^{(1)} \rho_i) \leq \epsilon_{II}, \quad i = 1, \dots, r - 1, \\ & \epsilon_I \mathbb{I} - M_{1|\emptyset}^{(1)} \in \mathcal{C}^*, \\ & (M_{s|\emptyset}^{(1)})_s \in \mathcal{M}. \end{aligned} \tag{15}$$

To optimize over the dual \mathcal{C}^* of the set of separable states, as required in (15), we invoke the Doherty–Parillo–Spedalieri (DPS) hierarchy^{16,17}. As shown in the Methods, the dual of this hierarchy approximates the set of all entanglement witnesses from the inside and converges as $n \rightarrow \infty$. In the case of two qubits the DPS hierarchy already converges at the first level. Hence, the particularly simple ansatz

$$\epsilon_I \mathbb{I} - \sum_{s \in \mathcal{S}_2} M_{s|\emptyset}^{(1)}(g(s)) = V_0 + V_1^{\mathcal{T}_B}, \tag{16}$$

where $V_0, V_1 \geq 0$ and \mathcal{T}_B is the partial transpose over the second subsystem, already leads us to derive tight bounds on the possible

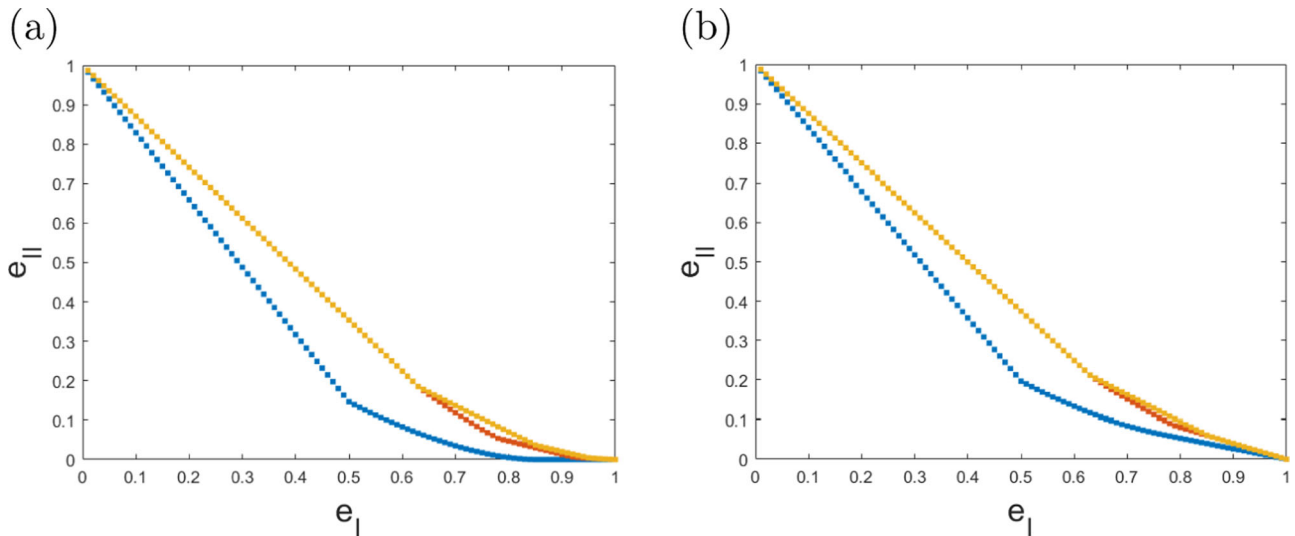


Fig. 3 One-shot entanglement certification for $|\phi\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. The referee has access to measurement strategies from the sets \mathcal{M}_1 (blue), \mathcal{M}_2 (red), \mathcal{M}_3 (yellow). We display the minimal e_{II} for fixed e_I . As each game corresponds to a hypothesis test, the most reasonable figure of merit is to quantify the type-I and type-II errors (e_I, e_{II}) a referee could achieve. These error pairs lie above the respective curves in the plots, any error-pair below is not possible with the resources at hand. Our optimisation also provides us with an explicit POVM, i.e., a measurement protocol, that achieves the optimal error pairs. **a** Entanglement detection for exact state preparation. The minimal total errors for $|\phi\rangle$ are $e_I + e_{II} = 0.6464$ with \mathcal{M}_1 , $e_I + e_{II} = 0.8152$ with \mathcal{M}_2 , and $e_I + e_{II} = 0.8153$ with \mathcal{M}_3 . For most randomly sampled states, these errors are much larger. We remark that there are also states, such as the singlet, where \mathcal{M}_2 and \mathcal{M}_3 lead to identical optimal errors. **b** Entanglement detection for noisy state preparation. To enforce that all states close to $\rho = |\phi\rangle\langle\phi|$ remain undetected with probability at most e_{II} , we need to invoke eq. (7), with $\mathcal{C} = \{\rho' : \rho' \geq 0, \text{tr}(\rho') = 1, \|\rho - \rho'\|_1 \leq \epsilon\}$. In the Methods we show how to derive the dual to this set. The plot displays the $\epsilon = 0.1$ case.

e_{II} , given e_I and the class of measurements available to the referee. For larger dimensional systems, enforcing condition (16) instead of the second constraint in (15) results in a sound but perhaps suboptimal protocol (namely, a protocol not necessarily minimizing e_{II}). Nevertheless, increasing the level of the DPS dual hierarchy generates a sequence of increasingly better (and sound) protocols whose type-II error converges to the minimum possible value asymptotically.

Eq. (15) requires us to enforce the constraint $(M_{s|\emptyset}^{(1)})_s \in \mathcal{M}$. For $\mathcal{M} = \mathcal{M}_1$, this amounts to demanding that the matrices $(M_{s|\emptyset}^{(1)})_s$ are positive semidefinite and add up to the identity. In that case, problem (15) can be cast as a semidefinite program (SDP)¹⁸.

For the cases $\mathcal{M} = \mathcal{M}_2, \mathcal{M}_3$, denote Alice and Bob’s choices of Pauli measurements by x and y , with outcomes, a, b respectively, and call $\gamma \in \{0, 1\}$ the outcome of the 1-way LPCC measurement. Then we can express Alice and Bob’s effective POVM as

$$M_{\gamma|\emptyset}^{(1)} = \sum_{a,b,x,y} P(x, y, \gamma|a, b) A_{a|x} \otimes B_{b|y}, \quad (17)$$

where the distribution $P(x, y, \gamma|a, b)$ is meant to model Alice and Bob’s classical processing of the outcomes they receive. For $\mathcal{M} = \mathcal{M}_2$, $P(x, y, \gamma|a, b)$ must satisfy the conditions

$$\sum_{\gamma} P(x, y, \gamma|a, b) = P(x) \quad \text{and} \quad \sum_{\gamma} P(x, y, \gamma|a, b) = P(x, \gamma|a), \quad (18)$$

whereas, for $\mathcal{M} = \mathcal{M}_3$, $P(x, y, \gamma|a, b)$ satisfies

$$\sum_{\gamma} P(x, y, \gamma|a, b) = P(x, y). \quad (19)$$

For $\mathcal{M} = \mathcal{M}_2, \mathcal{M}_3$, enforcing the constraint $(M_{s|\emptyset}^{(1)})_s \in \mathcal{M}$ thus requires imposing a few linear constraints on the optimization variables $P(x, y, \gamma|a, b)$. For these cases, problem (15) can therefore be cast as an SDP as well.

In Fig. 3, we compare the optimal error trade-offs for $\mathcal{M} = \mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3$ and further generalise this to scenarios, where, e.g.

due to experimental errors, the device preparing the target state ρ is actually distributing states ϵ -close to ρ in trace norm. The corresponding numerical optimisations, as well as any other convex optimization problem solved in this paper, were carried out using the semidefinite programming solver MOSEK¹⁹, in combination with the optimization packages YALMIP²⁰ or CVX²¹. We provide an example of a MATLAB implementation of these optimisations²².

We next consider the problem of finding the best strategy for $\mathcal{M} = \mathcal{M}_2, \mathcal{M}_3$ for n -round entanglement detection protocols. In this scenario, our general results for Maxwell demon games are not directly applicable. The reason is that, although both Alice and Bob are just allowed to conduct a finite set of physical measurements (namely, the three Pauli matrices), the set of effective local or LPCC measurements which they can enforce in each game round is not discrete. Nonetheless, a simple modification of the techniques developed for Maxwell demon games suffices to make the optimizations tractable. For this, we model Alice’s and Bob’s setting choices $(x_i)_i, (y_i)_i$ and final score γ of the game, depending on their respective outcomes $(a_i)_i, (b_i)_i$ through conditional distributions

$$P(x_1, y_1, x_2, y_2, \dots, x_n, y_n, \gamma|a_1, b_1, \dots, a_n, b_n). \quad (20)$$

Depending on whether the measurements in each round are taken from \mathcal{M}_2 or \mathcal{M}_3 this distribution will obey different sets of linear constraints. For the explicit reformulation of problem (3) as an SDP in this setting, we refer to the Methods.

Solving this optimization problem, we find the optimal multi-round error trade-offs for two-qubit entanglement detection in scenarios where the POVMs considered within each round are either in the set \mathcal{M}_2 (LPCC) or \mathcal{M}_3 (Local Pauli measurements), see Fig. 4.

Now let us consider the scenario from above where within each round a measurement from class \mathcal{M}_3 is applied in more detail. Does the adaptability of the choice of POVM between the rounds in a Maxwell demon game actually improve the error trade-offs?

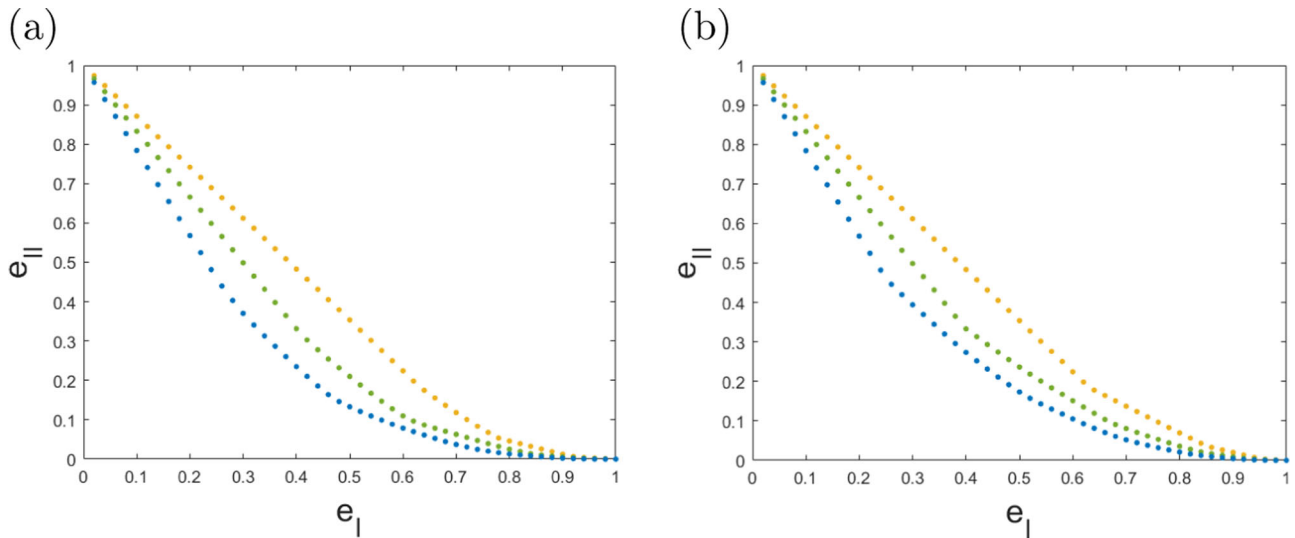


Fig. 4 Maxwell demon games played for various numbers of rounds. The referee has access to measurement strategies from the sets \mathcal{M}_2 (a) and \mathcal{M}_3 (b) within each round. The choice of the overall POVM implemented in each round will, in either case, depend on all inputs and outputs of previous rounds. The curves display the optimal trade-off between type-I and type-II error, (e_I, e_{II}) , for $n = 1$ (yellow), $n = 2$ (green) and $n = 3$ (blue) for $\mathcal{E} = \{|\phi\rangle\langle\phi|^{\otimes n}\}$.

Specifically, we aim to compare the case where the referee has to choose a POVM from \mathcal{M}_3 for each round of the game beforehand to the case where they can choose each POVM from \mathcal{M}_3 on the fly based on their previous inputs and outputs. The answer to this question is intuitively clear when we consider a set E of more than one state, since then we can conceive a strategy where in the first round we perform a measurement that allows us to get an idea which of the states in E we are likely dealing with, while in the second round we can then use the optimal witness for that state. However, more surprisingly, we find that this can also make a difference for a single state $E = \{|\psi\rangle\langle\psi|\}$. For instance, for the state $|\psi\rangle = \frac{1}{\sqrt{2}}(|0+\rangle + |1-i\rangle)$ with $| -i\rangle = \frac{1}{\sqrt{2}}(|0\rangle - i|1\rangle)$, we find that, in two-round games, the minimum value of $e_I + e_{II}$ equals 0.7979 with adaptation between rounds and 0.8006 without adaptation (see²³ or the Methods for a statistical interpretation of the quantity $e_I + e_{II}$).

This result may strike the reader as surprising: on first impulse, one would imagine that the best protocol to detect the entanglement of two preparations of a known quantum state ρ entails testing the same entanglement witness twice. A possible explanation for this counter-intuitive phenomenon is that preparations in \mathcal{E} and \mathcal{S} are somehow correlated: either both preparations correspond to ρ or both preparations correspond to a separable state. From this point of view, it is not far-fetched that an adaptive measurement strategy can exploit such correlations.

Our framework also naturally allows for the optimization over protocols with $e_{II} = 0$ and where the corresponding e_I error is being minimised, thus generalising previous work on detecting entanglement in few experimental rounds^{24,25}. Using the dual of the DPS hierarchy for full separability²⁶, we can furthermore derive upper bounds on the errors for states shared between more than two parties. Similarly, a hierarchy for detecting high-dimensional entangled states²⁷ allows us to derive protocols for the detection of high-dimensional entangled states using quantum preparation games²⁸.

Due to the exponential growth of the configuration space, optimisations over Maxwell demon adaptive measurement schemes are hard to conduct even for relatively low values of n . Devising entanglement detection protocols for $n \gg 1$ requires completely different techniques.

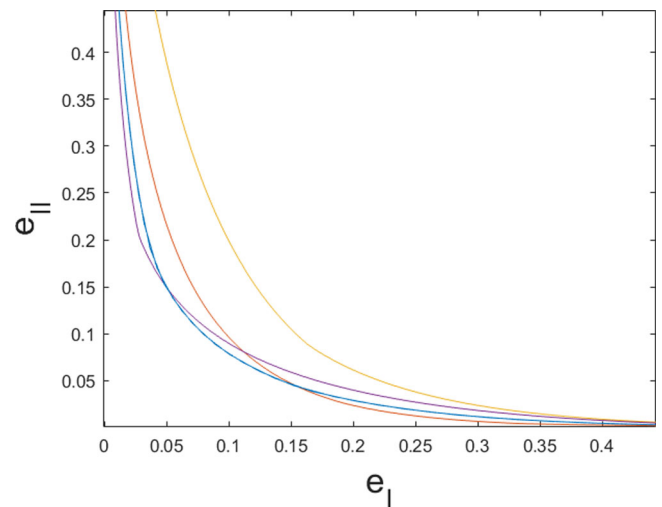


Fig. 5 Comparison of independent repetitions of 1-shot and 3-round games for $E = \{|\phi\rangle\langle\phi|\}$, in terms of the trade-off between type-I error, e_I , and type-II error, e_{II} . The games $G_{22}^{(30)}$ (yellow), $G_{25}^{(30)}$ (red) and $G_{28}^{(30)}$ (purple) are obtained through 30 independent repetitions of optimal one-shot games G restricted to measurements in \mathcal{M}_2 . These are compared to the optimal 3-round adaptive protocols G' with measurements \mathcal{M}_2 performed in each of the three rounds, independently repeated 10 times as $G_3^{(10)}$ (blue). The 1 and 3-round games G and G' were also analyzed in Fig. 4. We observe that the repetition of an adaptive protocol outperforms the others in the regime of low total error $e_I + e_{II}$.

Many-round protocols for entanglement detection. In order to devise many-round preparation games, an alternative to carrying out full optimizations is to rely on game composition. In this regard, in Fig. 5 we compare 10 independent repetitions of a 3-round adaptive strategy to 30 independent repetitions of a 1-shot protocol, based on (10). This way of composing preparation games can easily be performed with more repetitions. Indeed, for $m = 1000$ repetitions we find preparation games with errors at the order of $\approx 10^{-14}$. In the asymptotic regime, the binomial distribution of the number of 1-outcomes for a player restricted to separable strategies (see eq. (10)) can be approximated by a

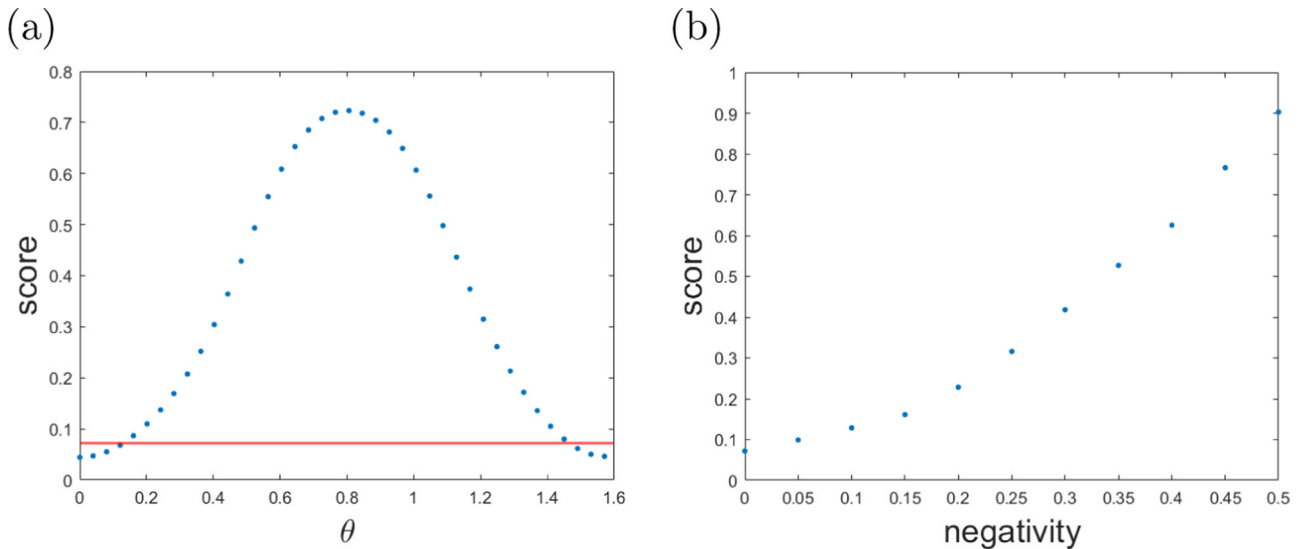


Fig. 6 Gradient descent based preparation game. The parameters are taken to be $\epsilon = 0.1$, $\lambda = 0.1$ and $\theta_0 = 0$. The probability of measuring $\{M_1^0, M_{-1}^0\}$ in round k is chosen according to $p_k(0) = \frac{1}{1+e^{-(2k-n)}}$. This captures the intuition that in the first few rounds it is more important to adjust the angle, while in later rounds the witness should be measured more often. **a** The score assigned to i.i.d. preparation strategies as a function of the parameter θ of $|\psi_\theta\rangle$ for $n = 41$ rounds for \mathcal{E} (blue) compared to the optimal separable value (red). As expected, the average game scores of the i.i.d. strategies $\{|\psi_\theta\rangle\langle\psi_\theta|^{\otimes n} : \theta\}$ mimic the shape of the curve $h(\cos^2(\theta))$ and the scores obtainable with the set of separable strategies \mathcal{S} perform significantly worse compared to the states from \mathcal{E} with angles close to $\theta = \frac{\pi}{4}$. **b** The optimal scores achievable by players capable of preparing bipartite quantum states of bounded negativity³⁰, obtained through application of eq. (2). We observe that the average score of the game constitutes a good estimator for entanglement negativity.

normal distribution. For $\epsilon_1 < \mu \equiv \frac{\nu}{m}$, this leads to a scaling as $e_1(m) \approx e^{-\frac{m(\mu-\epsilon_1)^2}{2\epsilon_1(1-\epsilon_1)}}$ (similarly for $(1 - \epsilon_{II})$ when $\mu < 1 - \epsilon_{II}$, and where the player is preparing states from \mathcal{E}).

Recall also that the heuristic presented in Box 2 is another viable option for moderate round numbers. We illustrate this in the Methods, where we use it to devise 20-round protocols for entanglement detection.

Finally, we apply gradient descent as a guiding principle to devise many-round protocols for entanglement quantification. For experimental convenience, the preparation game we develop is implementable with 1-way LOCC measurements.

We wish our protocol to be sound for i.i.d. strategies in $\mathcal{E} = \{\rho^{\otimes n} : \rho \in E\}$, with E being the set of all states

$$|\psi_\theta\rangle = \cos(\theta)|00\rangle + \sin(\theta)|11\rangle, \tag{21}$$

for $\theta \in (0, \pi/2)$. For such states, the protocol should output a reasonably good estimate of $|\psi_\theta\rangle$'s entanglement entropy $S(|\psi_\theta\rangle) = h(\cos^2(\theta))$, with $h(x) = -x \log(x) - (1-x) \log(1-x)$ the binary entropy. Importantly, if the player is limited to preparing separable states, the average score of the game should be low.

Following eq. (11), we introduce

$$W(\theta) = \frac{1}{2}[Z \otimes Z + |+\rangle\langle+| \otimes (\sin(2\theta)X + \cos(2\theta)Z)] + \frac{1}{2}[|-\rangle\langle-| \otimes (-\sin(2\theta)X + \cos(2\theta)Z)]. \tag{22}$$

This operator satisfies $\|W(\theta)\| \leq 1$ and $|\psi_\theta\rangle$ is the only eigenvector of $W(\theta)$ with eigenvalue 1. $W(\theta)$ can be estimated via 1-way LOCC with the POVM $M_{-1}^0(\theta) = \frac{1-W(\theta)}{2}$, $M_1^0(\theta) = \frac{1+W(\theta)}{2}$. Furthermore, consider

$$\frac{\partial}{\partial \theta} W = |+\rangle\langle+| \otimes (\cos(2\theta)X - \sin(2\theta)Z) - |-\rangle\langle-| \otimes (\cos(2\theta)X + \sin(2\theta)Z). \tag{23}$$

This dichotomic observable can be estimated via eq. (13) with the

1-way LOCC POVM defined by

$$M_{-1}^1(\theta) = |+\rangle\langle+| \otimes \frac{1}{2}(\mathbb{I} - \cos(2\theta)X + \sin(2\theta)Z) + |-\rangle\langle-| \otimes \frac{1}{2}(\mathbb{I} + \cos(2\theta)X + \sin(2\theta)Z), \tag{24}$$

$$M_1^1(\theta) = \mathbb{I} - M_{-1}^1(\theta),$$

which satisfies $M_1^1 - M_{-1}^1 = \frac{\partial}{\partial \theta} W$.

Let us further take $f(\theta, \nu) = h(\cos^2(\theta)) \Theta(\nu - (1 - \lambda + \lambda \delta(\theta)))$ with $0 \leq \lambda \leq 1$ and $\delta(\theta) = \max_{\rho \in \mathcal{C}} \text{tr}[W(\theta)\rho]$. This captures the following intuition: if the estimate ν of $\text{tr}[W(\theta_n)\rho]$ is below a convex combination of the maximum value achievable (namely, $\langle\psi_\theta|W(\theta_n = \theta)|\psi_\theta\rangle = 1$) and the maximum value $\delta(\theta_n)$ achievable by separable states, then the state shall be regarded as separable and thus the game score is set to zero. In Fig. 6, we illustrate how this game performs.

Discussion

We have introduced quantum preparation games as a convenient framework to analyze the certification and quantification of resources. We derived general methods to compute the (maximum) average score of arbitrary preparation games under different restrictions on the preparation devices: this allowed us to prove the soundness or security of general certification protocols. Regarding the generation of such protocols, we explained how to conduct exact (approximate) optimizations over preparation games with a low (moderate) number of rounds. In addition, we introduced two methods to devise large-round preparation games, via game composition and through gradient descent methods. These general results were applied to devise novel protocols for entanglement detection and quantification. To our knowledge, these are the first non-trivial adaptive protocols ever proposed for this task. In addition, we discovered that, against the common practice in entanglement detection, entanglement certification protocols for a known quantum state can often be improved using adaptive measurement strategies.

Even though we illustrated our general findings on quantum preparation games with examples from entanglement theory, where the need for efficient protocols is imminent, we have no doubt that our results will find application in other resource theories. With the current push towards building a quantum computer, a second use of our results that should be particularly emphasized is the certification of magic states. More generally, developing applications of our work to various resource theories, including for instance the quantification of non-locality, is an interesting direction for future work.

Another compelling line of research consists in studying the average performance of preparation games where Assumption 1 does not hold. In those games, a player can exploit the action of the referee’s measurement device to generate states outside the class allowed by their preparation device. Such games naturally arise when the player is limited to preparing resource-free states for some resource theory, but the referee is allowed to conduct resourceful measurements. An obvious motivating example of these games is the detection of magic states via general POVMs.

Finally, it would be interesting to explore an extension of preparation games where the referee is allowed to make the received states interact with a quantum system of fixed dimension in each round. This scenario perfectly models the computational power of a Noisy Intermediate-Scale Quantum (NISQ) device. In view of recent achievements in experimental quantum computing, this class of games is expected to become more and more popular in quantum information theory.

Methods

The maximum average score of finitely correlated strategies. Here we explain how to compute the maximum average score achievable in a preparation game by a player conducting a finitely correlated strategy (see Fig. 2), under the assumption that the quantum operation effected by the preparation device is known, but not the initial state of the environment.

In such preparations, the player’s device interacts with an environment A . More specifically, in each round, the referee receives a state

$$\text{tr}_A \left[\sum_i K_i \rho K_i^\dagger \right],$$

where ρ is the current state of the environment and $K_i : \mathcal{H}_A \rightarrow \mathcal{H}_A \otimes \mathcal{H}$ are the Kraus operators, which evolve the environment and prepare the state that the referee receives. Since the same environment is interacting with each prepared state, the states that the referee receives in different rounds are generally correlated.

Suppose that the referee concludes the first round of their adaptive strategy in the game configuration s . The (non-normalized) state of the environment will then be

$$\sum_{i_1, j_1} \langle I_1 | M_{s_1 \emptyset}^{(1)} | j_1 \rangle \tilde{K}_{i_1, j_1} \rho(\tilde{K}_{i_1, j_1})^\dagger, \tag{25}$$

where $\tilde{K}_{ij} = (\mathbb{I}_A \otimes \langle j |) K_i$. Iterating, we find that, if the referee observes the sequence of game configurations $\emptyset, s_2, \dots, s_n, \bar{s}$, then the final state of the environment will be

$$\sum_{i_1, j_1} \langle I_1 | M_{s_1 \emptyset}^{(1)} | j_1 \rangle \dots \langle I_n | M_{s_n s_n}^{(n)} | j_n \rangle \tilde{K}_{i_1, j_1} \dots \tilde{K}_{i_n, j_n} \rho(\tilde{K}_{i_1, j_1})^\dagger \dots (\tilde{K}_{i_n, j_n})^\dagger. \tag{26}$$

The probability to obtain such a sequence of configurations is given by the trace of the above operator. The average score of the game is thus $\text{tr}[\rho \Omega]$, where the operator Ω is defined by

$$\Omega = \sum_{s_2, \dots, s_n, \bar{s}} \sum_{i, j, l} \langle I_1 | M_{s_1 \emptyset}^{(1)} | j_1 \rangle \dots \langle I_n | M_{s_n s_n}^{(n)} | j_n \rangle (\tilde{K}_{i_1, j_1})^\dagger \dots (\tilde{K}_{i_n, j_n})^\dagger \tilde{K}_{i, l} \dots \tilde{K}_{i, j} \langle g(\bar{s}) \rangle. \tag{27}$$

Note that Ω can be expressed as the composition of a sequence of linear transformations. More concretely, consider the following recursive definition

$$\begin{aligned} \Omega_s^{(n)} &= \sum_{s \in \mathcal{S}} \sum_{i, j, l} (\tilde{K}_{i, j})^\dagger \tilde{K}_{i, l} \langle I | M_{s | s}^{(n)} | j \rangle \langle g(\bar{s}) \rangle, \\ \Omega_s^{(k)} &= \sum_{i, j, l} \sum_{s \in \mathcal{S}_{k+1}} (\tilde{K}_{i, j})^\dagger \Omega_{s, s'}^{(k+1)} \tilde{K}_{i, l} \langle I | M_{s' | s}^{(k)} | j \rangle. \end{aligned} \tag{28}$$

Then it can be verified that $\Omega = \Omega_\emptyset^{(1)}$. Calling D the Hilbert space dimension of the environment, the average score of the considered preparation game can thus be computed with $O(D^2 \sum_k |S_k| |S_{k+1}|)$ operations.

In realistic experimental situations, the player will not know the original quantum state ρ_A of the environment. In that case, we may be interested in

computing the maximum average score achievable over all allowed environment states. Let us assume that $\rho_A \in \mathcal{A}$, for some convex set \mathcal{A} . Then, the maximum average score, v , is

$$\max_{\rho_A \in \mathcal{A}} \text{tr}[\rho_A \Omega]. \tag{29}$$

In case the environment is fully unconstrained, this quantity equals the maximum eigenvalue of Ω .

This condition can be seen to be equivalent to

$$v \mathbb{I} - \Omega \in \mathcal{A}^*, \tag{30}$$

where \mathcal{A}^* denotes the dual of \mathcal{A} , i.e., $\mathcal{A}^* = \{X : \text{tr}(X\rho) \geq 0, \forall \rho \in \mathcal{A}\}$. In the particular case where the initial state of the environment is unconstrained, the condition turns into

$$v \mathbb{I} - \Omega \geq 0. \tag{31}$$

Since Ω is a linear function of the optimization variables, condition (30)– or (31)– is a convex constraint and thus we can handle it within the framework of convex optimization theory.

Enforcing \mathcal{C} -constrained preparation strategies in Maxwell-demon games. In the following we show how to turn (8) into a set of linear constraints on $\{P(y_0, \dots, y_n | a_0, \dots, a_n) : a_0, \dots, a_n\}$. We then show how to formulate the constraints (5), when \mathcal{S} is a set of \mathcal{C} -constrained strategies in terms of the variables $\{P(y_0, \dots, y_n | a_0, \dots, a_n) : a_0, \dots, a_n\}$. This allows us to treat the optimization of multi-round Maxwell demon games with convex optimization techniques.

Let us first show that (8) and (9) are equivalent. That any distribution of the form (8) satisfies (9) can be checked with a straightforward calculation. Conversely, for any set of distributions $\{P(y_0, \dots, y_n | a_0, \dots, a_n) : a_0, \dots, a_n\}$ satisfying (9), there exist distributions $P_k(x_k | s_k), P(y | s_{n+1})$ such that (8) holds¹². Indeed, one can derive the latter from $\{P(x_1, \dots, x_k | a_0, a_1, \dots, a_{k-1})\}_k$ via the relations

$$\begin{aligned} P_k(x_k | s_k) &= \frac{P(x_1, \dots, x_k | a_0, \dots, a_{k-1})}{P(x_1, \dots, x_{k-1} | a_0, \dots, a_{k-2})} \\ P(y | s_{n+1}) &= \frac{P(x_1, \dots, x_n, y | a_0, a_1, \dots, a_n)}{P(x_1, \dots, x_n | a_0, a_1, \dots, a_{n-1})}. \end{aligned} \tag{32}$$

For fixed measurements $\{N_{a|x}^{(k)} : a, x\}$, optimizations over Maxwell demon games thus reduce to optimizations over non-negative variables $P(x_1, \dots, x_n, y | a_0, a_1, \dots, a_n)$ satisfying eq. (9), positivity and normalization

$$\sum_{y_0, \dots, y_n} P(y_0, \dots, y_n | a_0, \dots, a_n) = 1 \quad \forall a_0, \dots, a_n. \tag{33}$$

We next show how to enforce the constraint (5) when \mathcal{S} corresponds to the set of \mathcal{C} -constrained preparation strategies, for some set of states \mathcal{C} . Similarly to (2), we can enforce this constraint inductively. For $k = 1, \dots, n$, let $\nu_{s_k}^{(k)}, \xi_{s_{n+1}}^{(k)}$ be optimization variables, satisfying the linear constraints

$$\xi_{s_{n+1}}^{(k)} = \sum_{y \in \mathcal{G}} y P(x_1, \dots, x_n, y | a_0, a_1, \dots, a_n), \tag{34}$$

$$\nu_{s_n}^{(n)} \mathbb{I} - \sum_{a_n, x_n} \xi_{s_{n+1}}^{(n)} N_{a_n | x_n}^{(n)} \in \mathcal{C}^*, \tag{35}$$

and

$$\nu_{s_k}^{(k)} \mathbb{I} - \sum_{a_{k+1}, x_{k+1}} \nu_{s_{k+1}}^{(k+1)} N_{a_{k+1} | x_{k+1}}^{(k+1)} \in \mathcal{C}^*. \tag{36}$$

We claim that $\nu_\emptyset^{(1)}$ is an upper bound on the maximum average score achievable by a player restricted to prepare states in \mathcal{C} . Indeed, let $\rho_{s_k}^{(k)} \in \mathcal{C}$ be the player’s preparation at stage k conditioned on the game configuration s_k . Multiply eq. (35) by $\rho_{s_n}^{(n)}$ and take the trace. Then, since eq. (35) belongs to the dual set of \mathcal{C} , we have that

$$\nu_{s_n}^{(n)} \geq \sum_{a_n, x_n} \sum_{y \in \mathcal{G}} y P(x_1, \dots, x_n, y | a_0, a_1, \dots, a_n) \text{tr} \left[N_{a_n | x_n}^{(n)} \rho_{s_n}^{(n)} \right]. \tag{37}$$

Next, we multiply both sides of the above equation by $\text{tr} [N_{a_{n-1} | x_{n-1}}^{(n-1)} \rho_{s_{n-1}}^{(n-1)}]$ and sum over the variables a_{n-1}, x_{n-1} . By eq. (36), the result will be upper bounded by $\nu_{s_{n-1}}^{(n-1)}$. Iterating this procedure, we arrive at

$$\nu_\emptyset^{(1)} \geq \sum_{a_1, \dots, a_n, x_1, \dots, x_n} \sum_{y \in \mathcal{G}} y P(x_1, \dots, x_n, y | a_0, a_1, \dots, a_n) \prod_{k=1}^n \text{tr} [N_{a_k | x_k}^{(k)} \rho_{s_k}^{(k)}]. \tag{38}$$

The right-hand side is the average score of the game. Call $\omega_{s_k}^{(k)} \in \mathcal{C}^*$ the operator expressions appearing in eqs. (35), (36). Note that, if there exist states $\rho_{s_k}^{(k)} \in \mathcal{C}$ such that $\text{tr} [\omega_{s_k}^{(k)} \rho_{s_k}^{(k)}] = 0$, i.e., if all the dual elements are tight, then the preparation strategy defined through the states $\{\rho_{s_k}^{(k)}\}$ achieves the average score $\nu_\emptyset^{(1)}$.

In sum, optimizations of the sort (3) over the set of all Maxwell demon games require optimizing over P under non-negativity and the linear constraints (9), (33). Constraints of the form (5) for $\mathcal{S} = \{P\}$ translate as extra linear constraints on P and the upper bound variable v . When \mathcal{S} corresponds to a finitely correlated strategy with unknown environment state, we can formulate condition (5) as the convex constraint (30). Finally, when \mathcal{S} corresponds to a set of \mathcal{C} -constrained

strategies, condition (5) is equivalent to enforcing constraints (34), (35) and (36) on \mathcal{P} and the slack variables $\nu, \xi_{s_{n+1}}$, with $\nu \equiv \nu_{\theta}^{(1)}$.

Computing the average score of a meta-preparation game. Our starting point is an n -round meta-game with configuration spaces $S = (S_1, S_2, \dots, S_{n+1})$, with $S_1 = \{\emptyset\}$. In each round k , the referee runs a preparation game $G_k(s_k)$. Depending on the outcome $o_k \in O_k$ of the preparation game, the referee samples the new configuration s_{k+1} from the distribution $c_k(s_{k+1}|o_k, s_k)$. The final score of the meta-game is decided via the non-deterministic function $\gamma: S_{n+1} \rightarrow \mathbb{R}$.

To find the optimal score achievable by a player using \mathcal{C} -constrained strategies, we proceed as we did for preparation games. Namely, define $\nu_s^{(k)}$ as the maximum average score achievable by a \mathcal{C} -constrained player, conditioned on the game being in configuration $s \in S_k$ at round k . Then we see that

$$\begin{aligned} \nu_s^{(n)} &= \max_{\mathcal{P} \in \mathcal{S}} \sum_{o, s'} p(o|\mathcal{P}, G_n(s)) c_n(s'|s, o) \gamma(s'), \\ \nu_s^{(k)} &= \max_{\mathcal{P} \in \mathcal{S}} \sum_{o, s'} p(o|\mathcal{P}, G_k(s)) c_k(s'|s, o) \nu_{s'}^{(k+1)}. \end{aligned} \tag{39}$$

Note that the first optimization above consists in finding the maximum average score of the preparation game $G_n(s)$ with score function $g(o) = \sum_{s'} c_n(s'|s, o) \gamma(s')$. Similarly, the second optimization corresponds to computing the maximum score of a preparation game with score function $g(o) = \sum_{s'} c_k(s'|s, o) \nu_{s'}^{(k+1)}$. Applying formula (2) iteratively, we can thus compute the average score of the meta-preparation game through $O(nn')$ operations, where n' denotes the maximum number of rounds of the considered preparation games.

Notice also that formula (39) also applies to compute the score of a meta j -preparation game, if we understand $\{G_k(s): s \in S_k, k\}$ as meta $^{j-1}$ -preparation games.

Now, consider a scenario where the games (or meta $^{j-1}$ -games) have just two possible final configurations, i.e., $o \in \{0, 1\}$. In that case,

$$\begin{aligned} \nu_{s,1}^{(n)} &= \Gamma_{s,1}^{(n)} + \max_{\mathcal{P} \in \mathcal{S}} p(0|\mathcal{P}, G_n(s)) (\Gamma_{s,0}^{(n)} - \Gamma_{s,1}^{(n)}), \\ \nu_s^{(k)} &= \Gamma_{s,1}^{(k)} + \max_{\mathcal{P} \in \mathcal{S}} p(0|\mathcal{P}, G_k(s)) (\Gamma_{s,0}^{(k)} - \Gamma_{s,1}^{(k)}), \end{aligned} \tag{40}$$

where

$$\Gamma_{s,0}^{(n)} \equiv \sum_{s'} c_n(s'|s, 0) g(s'), \quad \Gamma_{s,0}^{(k)} \equiv \sum_{s'} c_k(s'|s, 0) \nu_{s'}^{(k+1)}, \tag{41}$$

for $o = 0, 1$.

Call $p_{\max}(G)$ ($p_{\min}(G)$) the solution of the problem $\max_{\mathcal{P} \in \mathcal{S}} p(0|G, \mathcal{P})$ ($\min_{\mathcal{P} \in \mathcal{S}} p(0|G, \mathcal{P})$). Then we have that

$$\begin{aligned} \nu_s^{(n)} &= \begin{cases} \Gamma_{s,1}^{(n)} + p_{\max}(G_n(s)) (\Gamma_{s,0}^{(n)} - \Gamma_{s,1}^{(n)}), & \text{for } \Gamma_{s,0}^{(n)} > \Gamma_{s,1}^{(n)}, \\ \Gamma_{s,1}^{(n)} + p_{\min}(G_n(s)) (\Gamma_{s,0}^{(n)} - \Gamma_{s,1}^{(n)}), & \text{otherwise,} \end{cases} \\ \nu_s^{(k)} &= \begin{cases} \Gamma_{s,1}^{(k)} + p_{\max}(G_k(s)) (\Gamma_{s,0}^{(k)} - \Gamma_{s,1}^{(k)}), & \text{for } \Gamma_{s,0}^{(k)} > \Gamma_{s,1}^{(k)}, \\ \Gamma_{s,1}^{(k)} + p_{\min}(G_k(s)) (\Gamma_{s,0}^{(k)} - \Gamma_{s,1}^{(k)}), & \text{otherwise.} \end{cases} \end{aligned}$$

If the same set J of meta $^{j-1}$ -games are re-used at each round of the considered meta j -game, this formula saves us the trouble of optimizing over meta $^{j-1}$ -games for every round k and every $s \in S_k$. The complexity of computing the maximum average score is, in this case, of order $O(n) + 2|J|c$, where c is the computational cost of optimizing over a meta $^{j-1}$ -game.

Think of a meta j -game where a given meta $^{j-1}$ -game G (with scores $o \in \{0, 1\}$) is played m times, $s_k \in \{0, \dots, k-1\}$ corresponds to the number of 1's obtained, and success is declared if $s_{m+1} \geq \nu$, for some $\nu \in \{0, \dots, m\}$. Then, we have that

$$\Gamma_{s,0}^{(n)} = \Theta(s + o - \nu), \quad \Gamma_{s,0}^{(k)} = \nu_{s+o}^{(k+1)}, \tag{42}$$

where $\Theta(x) = 0$, for $x < 0$ or 1 otherwise. It is thus clear that $\Gamma_{s,0}^{(k)} \leq \Gamma_{s,1}^{(k)}$ for all k , and so the best strategy consists in always playing to maximize $p(1|G, \mathcal{P})$ in each round. In turn, this implies the binomial formula (10) derived in¹³.

Furthermore, as shown in²³, if the player uses a strategy $\mathcal{Q} \notin \mathcal{S}$ to play G , with $G(\mathcal{Q}) \geq G(\mathcal{P}^*)$, then the average value of $p(G, \nu, m)$ can be seen to satisfy

$$\sum_{\nu=0}^m p^{(m)}(\nu|\mathcal{Q}) p(G, \nu, m) \leq [1 - (G(\mathcal{Q}) - G(\mathcal{P}^*))^2]^m, \tag{43}$$

where $p^{(m)}(\nu|\mathcal{Q})$ denotes the probability of winning ν times with strategy \mathcal{Q} . This relation has important applications for hypothesis testing: if, by following the strategy \mathcal{Q} , we wish to falsify the hypothesis that the player is using a strategy in \mathcal{P} , all we need to do is play a preparation game for which $G(\mathcal{Q}) - G(\mathcal{P}^*)$ is large enough multiple times.

Optimizing over the set of separable states and its dual. In the main text, we frequently encountered convex constraints of the form

$$\nu \mathbb{I} - W \in \mathcal{C}^*, \tag{44}$$

where W is an operator and \mathcal{C} is a convex set of quantum states. Furthermore, we had to conduct several optimizations of the form

$$f^* = \max_{\rho \in \mathcal{C}} \text{tr}[W\rho]. \tag{45}$$

In the following, we will explain how to tackle these problems when \mathcal{C} corresponds to the set SEP of separable quantum states on some bipartite Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$.

In this regard, the Doherty-Parrilo-Spedalieri (DPS) hierarchy^{16,17} provides us with a converging sequence of semi-definite programming outer approximations to SEP. Consider the set E_k of $k+1$ -partite quantum states defined by

$$E_k = \{\rho_{AB_1 \dots B_k} : \Pi_k \rho_{AB_1 \dots B_k} \Pi_k = \rho_{AB_1 \dots B_k}, \rho_{AB_1 \dots B_k}^{\mathcal{T}^s} \geq 0 \forall s \in \mathcal{N} \text{ and } \text{tr}[\rho_{AB_1 \dots B_k}] = 1\}, \tag{46}$$

where Π_k is the projector onto the symmetric subspace of $\mathcal{H}_{B_1} \otimes \dots \otimes \mathcal{H}_{B_k}$; \mathcal{N} is the power set of $\{B_1, \dots, B_k\}$; and \mathcal{T}^s denotes the partial transpose over the subsystems S .

We say that the quantum state ρ_{AB} admits a Bose-symmetric PPT extension to k parts on system B iff there exists $\rho_{AB_1 \dots B_k} \in E_k$ such that $\rho_{AB} = \text{tr}_{B_2, \dots, B_k}(\rho_{AB_1 \dots B_k})$. Call SEP^k the set of all such bipartite states. Note that the condition $\rho_{AB} \in \text{SEP}^k$ can be cast as a semidefinite programming constraint.

As shown in^{16,17}, $\text{SEP}^1 \supset \text{SEP}^2 \supset \dots \supset \text{SEP}$ and $\lim_{k \rightarrow \infty} \text{SEP}^k = \text{SEP}$. Hence, for $\mathcal{C} = \text{SEP}$, we can relax optimizations over (45) by optimizing over one of the sets SEP^k instead. Since $\text{SEP}^k \supset \text{SEP}$, the solution f^k of such a semidefinite program will satisfy $f^k \geq f^*$. Moreover, $\lim_{k \rightarrow \infty} f^k = f^*$. For entanglement detection problems, the use of a relaxation of \mathcal{C} in optimizations such as (2) results in an upper bound on the maximum average game score.

To model constraints of the form (44), we similarly replace the dual of SEP by the dual of SEP^k in eq. (44), that is, as we shall show, also admits a semidefinite programming representation. Since $\text{SEP}^* \supset (\text{SEP}^k)^*$, we have that $\nu \mathbb{I} - W \in (\text{SEP}^k)^*$ implies $\nu \mathbb{I} - W \in \text{SEP}^*$. However, there might exist values of ν such that $\nu \mathbb{I} - W \in \text{SEP}^*$, but $\nu \mathbb{I} - W \notin (\text{SEP}^k)^*$. Such replacements in expressions of the form (61) will lead, as before, to an overestimation of the maximum average score of the game for the considered set of preparation strategies.

Let us thus work out a semidefinite representation for the set $(\text{SEP}^k)^*$. By duality theory¹⁸, we have that any $W \in E_k^*$ must be of the form

$$W = (C - \Pi_k C \Pi_k) + \sum_{S \in \mathcal{N}} M_S^{\mathcal{T}^s}, \tag{47}$$

for some positive semidefinite matrices $\{M_S\}_S$. Indeed, multiplying by $\rho_{AB_1, \dots, B_k} \in E_k$ and taking the trace, we find, by virtue of the defining relations (46) that the trace of ρ_{AB_1, \dots, B_k} with respect to each term in the above equation is non-negative.

Multiplying on both sides of (47) by Π_k , we arrive at the equivalent condition

$$\Pi_k W \Pi_k = \Pi_k \left(\sum_{S \in \mathcal{N}} M_S^{\mathcal{T}^s} \right) \Pi_k. \tag{48}$$

Now, let $V \in (\text{SEP}^k)^*$, and let $\rho_{AB} \in \text{SEP}^k$ with extension $\rho_{AB_1, \dots, B_k} \in E_k$. Then we have that

$$\text{tr}[V \rho_{AB}] = \text{tr}[(V \otimes \mathbb{I}_B^{\otimes k-1}) \rho_{AB_1, \dots, B_k}] \geq 0. \tag{49}$$

Since this relation must hold for all $\rho_{AB_1, \dots, B_k} \in E_k$, it follows that

$V \otimes \mathbb{I}_B^{\otimes k-1} \in E_k^*$. In conclusion, $V \in (\text{SEP}^k)^*$ iff there exist positive semidefinite matrices $\{M_S\}_S$ such that

$$\Pi_k (V \otimes \mathbb{I}_B^{\otimes k-1}) \Pi_k = \Pi_k \left(\sum_{S \in \mathcal{N}} M_S^{\mathcal{T}^s} \right) \Pi_k. \tag{50}$$

This constraint obviously admits a semidefinite programming representation.

For $\dim(\mathcal{H}_A) \dim(\mathcal{H}_B) \leq 6$, $\text{SEP}^1 = \text{SEP}$ ²⁹. In such cases, we have by eq. (50), that

$$\text{SEP}^* = \{V : V = V_0 + V_1^{\mathcal{T}^B}, V_0, V_1 \geq 0\}. \tag{51}$$

Maxwell demon games for entanglement detection. Here we provide the technical details regarding the applications of Maxwell demon games to entanglement certification. We consider honest players with i.i.d. strategies $\mathcal{E} = \{\rho^{\otimes n} : \rho \in E\}$, and we are interested in the worst-case errors $\max_{\rho \in E} e_{\Pi}(M, \rho^{\otimes n})$.

In each round k , Alice and Bob must choose the indices x_k, y_k of the measurements that they will conduct on their respective subsystems. That is, in round k Alice (Bob) will conduct the measurement $\{A_{a_k}^{(k)} : a\}$ ($\{B_{b_k}^{(k)} : b\}$), with outcome a_k (b_k). To model their (classical) decision process, we will introduce the variables $\{P(x_1, y_1, x_2, y_2, \dots, x_n, y_n, \gamma|a_1, b_1, \dots, a_n, b_n)\}$.

$P(x_1, y_1, x_2, y_2, \dots, x_n, y_n, \gamma|a_1, b_1, \dots, a_n, b_n)$ will satisfy some linear restrictions related to the no-signalling to the past condition, whose exact expression depends on how Alice and Bob conduct their measurements in each round. If, in each round, Alice and Bob make use of 1-way LOCC measurements from Alice to Bob (measurement class \mathcal{M}_2), then P will satisfy the constraints

$$\begin{aligned} \sum_{\gamma} P(x_1, y_1, \dots, x_n, y_n, \gamma|a_1, b_1, \dots, a_n, b_n) &= P(x_1, y_1, x_2, y_2, \dots, x_n, y_n|a_1, b_1, \dots, a_{n-1}, b_{n-1}, a_n), \\ \sum_{x_{k+1}, \dots, x_n, y_{k+1}, \dots, y_n} P(x_1, y_1, \dots, x_n, y_n|a_1, b_1, \dots, a_{k-1}, b_{k-1}, a_k) &= P(x_1, y_1, x_2, y_2, \dots, x_{k-1}, y_{k-1}|a_1, b_1, \dots, a_{k-2}, b_{k-2}, a_{k-1}), \\ \sum_{x_{k+1}, \dots, x_n, y_{k+1}, \dots, y_n} P(x_1, y_1, \dots, x_n, y_n|a_1, b_1, \dots, a_{k-1}, b_{k-1}, a_k) &= P(x_1, y_1, x_2, y_2, \dots, x_k|a_1, b_1, \dots, a_{k-1}, b_{k-1}). \end{aligned} \tag{52}$$

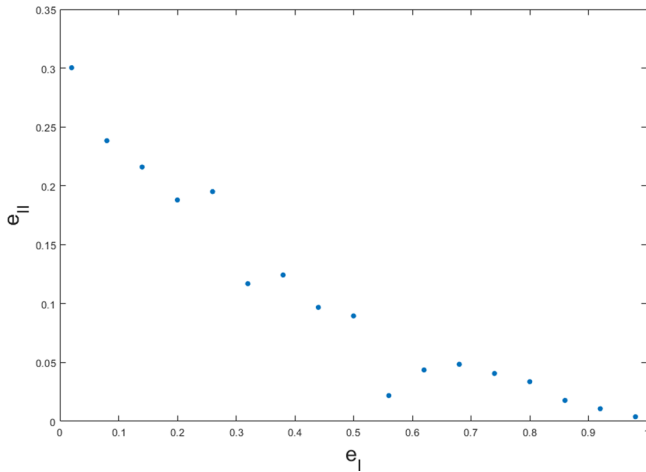


Fig. 7 Error trade-off for a finitely correlated scenario with bounded configuration space. The results are obtained with a 10-dimensional unknown environment that interacts with a maximally entangled state for $\tau = 0.1$ according to the Hamiltonian (62). There were 20 measurement rounds ($n = 20$), and in each of the first 19 rounds a 6-outcome measurement was performed, with the option of outputting 0 available as one of the outcomes of each measurement. These results were obtained through the method outlined in the main text. For each value of e_I , we plot the minimum e_{II} achieved in 10 runs (each time with a different random initialization of the measurements). Each run has been optimized until convergence was achieved. Although the type-II errors obtained are reasonably small, the curve presents large discontinuities and, in fact, is not even decreasing. Presumably, for many of the values of e_I , the initial (random) measurement scheme fed into the algorithm led the latter to a local minimum. This explains, e.g., the sudden drop of the type-II error after $e_I = 0.5$.

If, on the contrary, Alice and Bob use local measurements in each round (measurement class \mathcal{M}_3), then the constraints on P will be

$$\sum_{\gamma} P(x_1, y_1, \dots, x_n, y_n, \gamma | a_1, b_1, \dots, a_n, b_n) = P(x_1, y_1, x_2, y_2, \dots, x_n, y_n | a_1, b_1, \dots, a_{n-1}, b_{n-1}),$$

$$\sum_{a_k, \dots, a_n, \gamma_k, \dots, \gamma_n} P(x_1, y_1, \dots, x_n, y_n | a_1, b_1, \dots, a_{n-1}, b_{n-1}) = P(x_1, y_1, x_2, y_2, \dots, x_{k-1}, y_{k-1} | a_1, b_1, \dots, a_{k-2}, b_{k-2}).$$

As explained in Results, the constraints (34), (35), (36) require minor modifications, to take into account that, in each round, the set of effective measurements of Alice and Bob is not finite (although the set of local measurements of either party is). More specifically, we define $s_k = (x_1, y_1, a_1, b_1, \dots, x_{k-1}, y_{k-1}, a_{k-1}, b_{k-1})$ to be the game configuration at the beginning of round k . To enforce eq. (5) for \mathcal{C} -constrained strategies (where, in this case, \mathcal{C} denotes the set of separable quantum states), we introduce the following relations:

$$\xi_{s_{n+1}} = \sum_{\gamma \in \mathcal{G}} \gamma P(x_1, \dots, y_n, \gamma | a_0, b_0, \dots, a_n, b_n),$$

$$\nu_{s_n}^{(n)} \mathbb{I} - \sum_{a_n, b_n, x_n, y_n} \xi_{s_{n+1}} A_{a_n | x_n}^{(n)} \otimes B_{b_n | y_n}^{(n)} \in \mathcal{C}^*,$$

$$\nu_{s_k}^{(k)} \mathbb{I} - \sum_{a_k, b_k, x_k, y_k} \nu_{s_{k+1}}^{(k+1)} A_{a_{k+1} | x_{k+1}}^{(k+1)} \otimes B_{b_{k+1} | y_{k+1}}^{(k+1)} \in \mathcal{C}^*.$$

The dual of ϵ -balls of quantum states. To enforce relation (5) when \mathcal{S} corresponds to the set of strategies achievable by a player constrained to prepare states within an ϵ -ball around a quantum state ρ , we need to find the dual of the set of states

$$\mathcal{C}(\rho, \epsilon) = \{\rho' : \rho' \geq 0, \text{tr}(\rho') = 1, \|\rho' - \rho\|_1 \leq \epsilon\}.$$

For this purpose, let us consider the optimisation problem

$$\min_{\rho'} \text{tr}[M\rho']$$

s.t. $\rho' \in \mathcal{C}(\rho, \epsilon),$

which has a non-negative solution if and only if $M \in \mathcal{C}^*(\rho, \epsilon)$. This problem can be

written as

$$\min_{\rho', Z} \text{tr}[M\rho']$$

s.t. $\rho' \geq 0, \text{tr}(\rho') = 1,$

$$Z + \rho' - \rho \geq 0,$$

$$Z - \rho' + \rho \geq 0,$$

$$\text{tr}[Z] = \epsilon.$$

Now, note that the dual to this semi-definite program is

$$\max_{A, \mu, \lambda} 2 \text{tr}[A\rho] - \lambda(1 + \epsilon) - \mu$$

s.t. $A \geq 0,$

$$\lambda \mathbb{I} - A \geq 0,$$

$$(\mu + \lambda) \mathbb{I} + M - 2A \geq 0,$$

and that the two problems are strongly dual. Thus, (58) has a non-negative solution if and only if (57) does. This implies that

$$\mathcal{C}^*(\rho, \epsilon) = \{M : \exists \mu, \lambda \in \mathbb{R}, A \geq 0, \text{s.t. } \lambda \mathbb{I} - A \geq 0, (\mu + \lambda) \mathbb{I} + M - 2A \geq 0, 2 \text{tr}[A\rho] - \lambda(1 + \epsilon) - \mu \geq 0\}.$$

Round-by-round optimization of preparation games. In the following, we illustrate how to optimize the POVMs of an individual round of a preparation game. This is the main subroutine in the heuristic presented in Box 2. We then describe the application of this one-round optimisation to a particular problem: entanglement detection of finitely correlated states. We further illustrate the efficiency of our coordinate-descent-based heuristic, as presented in Box 2, in this example.

To optimise over a single game round, notice that Eq. (2) implies the conditions

$$\mu_s^{(n)} \mathbb{I} - \sum_{\bar{s} \in \mathcal{S}} (g(\bar{s})) M_{\bar{s}|s}^{(n)}, \mu_s^{(k)} \mathbb{I} - \sum_{\bar{s}} M_{\bar{s}|s}^{(k)} \mu_{\bar{s}}^{(k+1)} \in \mathcal{C}^*.$$

Optimizations over $\{M_{s_{k+1}|s_k}^{(k)} : s_k, s_{k+1}\}$ under a constraint of the form (5) can thus be achieved via the following convex optimization scheme: first, compute $\{\mu_s^{(j)} : j > k\}$ by induction via Eq. (2). Next, impose the constraints

$$\mu_{\emptyset}^{(1)} \leq \nu,$$

$$\mu_s^{(j)} \mathbb{I} - \sum_{\bar{s}} M_{\bar{s}|s}^{(j)} \mu_{\bar{s}}^{(j+1)} \in \mathcal{C}^*, \text{ for } j = 1, \dots, k.$$

$$\{M_{s_{k+1}|s_k}^{(k)} : s_{k+1}\} \subset \mathcal{M}, \text{ for } s_k \in \mathcal{S}_k.$$

Note that, in the second constraint of Eq. (61), either $M_{s'|s}^{(j)}$ or $\mu_{s'}^{(j+1)}$ is an optimization variable, but not both. This means that all the above are indeed convex constraints.

Remarkably, expressing condition (5) for i.i.d., finitely-correlated and \mathcal{C} -constrained strategies requires adding $O(\sum_{j \leq k} |\mathcal{S}_j|)$ new optimization variables, related to the original ones through a set of $O(\sum_{j \leq k} |\mathcal{S}_j|)$ constraints, all of which can be calculated with $O(\sum_j |\mathcal{S}_j| |\mathcal{S}_{j+1}|)$ operations. As long as the number of game configurations is not excessive, one can therefore carry these optimizations out for games with very large n .

We are now ready to test the practical performance of the heuristic described in Box 2. To this aim, consider the following entanglement detection scenario: an honest player is attempting to prepare the maximally entangled state $|\psi_{\frac{1}{2}}\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$, but, before being transmitted, the state interacts with the local environment ρ_A for a brief amount of time τ . Specifically, we take the environment to be a d_A -dimensional quantum system that interacts with the desired state through the Hamiltonian

$$H_I = a_A^\dagger \otimes (\mathbb{I} \otimes |0\rangle\langle 1| + |0\rangle\langle 1| \otimes \mathbb{I}) + a_A \otimes (\mathbb{I} \otimes |1\rangle\langle 0| + |1\rangle\langle 0| \otimes \mathbb{I}),$$

where a_A^\dagger and a_A are raising and lowering operators acting on the environmental system, respectively. We let the environment evolve only when it interacts with each new copy of $|\psi_{\frac{1}{2}}\rangle$. By means of global bipartite measurements \mathcal{M}_1 , we wish to detect the entanglement of the states prepared by the honest player. Our goal is thus to devise adaptive measurement protocols that detect the entanglement of a family of finitely correlated strategies of fixed interaction map, but with an unknown initial environment state.

When the player is following such a finitely correlated strategy, optimizing the k^{th} round measurements amounts to solving the following semi-definite program:

$$\min_{M^{(k)}, \{\mu^{(j)} : j \leq k\}} e_{II}$$

s.t. $\{M_{s'|s}^{(k)}\}_{s'} \subset \mathcal{M}_1, \text{ for } s \in \mathcal{S}_k$

$$\mu_{\emptyset}^{(1)} = e_I,$$

$$\mu_s^{(j)} \mathbb{I} - \sum_{\bar{s}} M_{\bar{s}|s}^{(j)} \mu_{\bar{s}}^{(j+1)} \in \mathcal{C}^*, \text{ for } j = 1, \dots, k,$$

$$\Omega(M^{(k)}) - (1 - e_{II}) \mathbb{I} \geq 0,$$

where $M^{(j)}$ ($\mu^{(j)}$) stands for $\{M_{s's}^{(j)} : s' \in S_{j+1}, s \in S_j\}$ ($\{\mu_s^{(j)} : s \in S_j\}$) and $\Omega(M^{(k)})$ is defined according to (28). The quantities $\{\mu^{(j)}; j > k\}$ do not depend on $M^{(k)}$, and hence can be computed via eq. (2) before running the optimization.

We consider a configuration space where $|S_k| = m$ for all $k = 2, 3, \dots, n$, and $S_{n+1} = \{0, 1\}$. In other words, the first $n - 1$ measurements are carried out with m -outcome POVMs, and the last measurement is dichotomic. Furthermore, in each round, we include the possibility of terminating the game early and simply outputting 0 (i.e., $0 \in S_k$). This models a scenario where the referee is convinced early that they will not be able to confidently certify the states to be entangled. Applying the coordinate-descent heuristic in Box 2 for different values of ϵ_1 , we arrive at the plot shown in Fig. 7.

Data availability

All raw data generated for the figures presented in this work are available from the corresponding author upon reasonable request.

Received: 4 December 2020; Accepted: 29 June 2021;

Published online: 27 July 2021

References

1. Yin, J. et al. Satellite-to-ground entanglement-based quantum key distribution. *Phys. Rev. Lett.* **119**, 200501 (2017).
2. Bennett, C. H. et al. Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels. *Phys. Rev. Lett.* **70**, 1895–1899 (1993).
3. Horodecki, M. & Oppenheim, J. (quantumness in the context of) resource theories. *Int. J. Mod. Phys. B* **27**, 1345019 (2012).
4. Bravyi, S. & Kitaev, A. Universal quantum computation with ideal clifford gates and noisy ancillas. *Phys. Rev. A* **71**, 022316 (2005).
5. Gühne, O. & Tóth, G. Entanglement detection. *Phys. Rep.* **474**, 1–75 (2009).
6. Mosonyi, M. & Ogawa, T. Quantum hypothesis testing and the operational interpretation of the quantum rényi relative entropies. *Commun. Math. Phys.* **334**, 1617–1648 (2015).
7. Fannes, M., Nachtergaele, B. & Werner, R. F. Finitely correlated states on quantum spin chains. *Commun. Math. Phys.* **144**, 443–490 (1992).
8. Terhal, B. M. & Horodecki, P. Schmidt number for density matrices. *Phys. Rev. A* **61**, 040301 (2000).
9. Veitch, V., Mousavian, S. A. H., Gottesman, D. & Emerson, J. The resource theory of stabilizer quantum computation. *N. J. Phys.* **16**, 013009 (2014).
10. Gurvits, L. Classical deterministic complexity of edmonds' problem and quantum entanglement, in <https://doi.org/10.1145/780542.780545> *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '03 (Association for Computing Machinery, New York, NY, USA, 2003) p. 10–19.
11. Gharibian, S. Strong np-hardness of the quantum separability problem. *Quant. Inf. Comput.* **10**, 343–360 (2010).
12. Hoffmann, J., Spee, C., Gühne, O. & Budroni, C. Structure of temporal correlations of a qubit. *N. J. Phys.* **20**, 102001 (2018).
13. Elkouss, D. & Wehner, S. (nearly) optimal p-values for all bell inequalities. *npj Quant. Inform.* **2**, 16026 (2016).
14. Bharti, K. et al. Noisy intermediate-scale quantum (nisq) algorithms, arXiv preprint arXiv:2101.08448 (2021).
15. Boyd, S., Xiao, L. & Mutapcic, A. Subgradient methods, <https://web.stanford.edu/class/ee392o/> lecture notes of EE392o, Stanford University, Autumn Quarter (2004).
16. Doherty, A. C., Parrilo, P. A. & Spedalieri, F. M. Distinguishing separable and entangled states. *Phys. Rev. Lett.* **88**, 187904 (2002).
17. Doherty, A. C., Parrilo, P. A. & Spedalieri, F. M. Complete family of separability criteria. *Phys. Rev. A* **69**, 022308 (2004).
18. Vandenberghe, L. & Boyd, S. Semidefinite programming. *SIAM Review* **38**, 49 (1996).
19. Vandenberghe, L. & Boyd, S. The mosek optimization toolbox for matlab manual. version 7.0 (revision 140). <http://docs.mosek.com/7.0/toolbox/index.html> (2020).
20. Löfberg, J. Yalmip : A toolbox for modeling and optimization in matlab, in *Proceedings of the CACSD Conference* (Taipei, Taiwan, 2004).
21. Grant, M. & Boyd, S. "CVX: Matlab software for disciplined convex programming, version 2.1", <http://cvxr.com/cvx> (2014).
22. Weilenmann, M., Aguilar, E. A. & Navascués, M. <https://github.com/MWeilenmann/Quantum-Preparation-Games> "Github repository for "quantum preparation games", (2020a).
23. Araújo, M., Hirsch, F. & Quintino, M. T. "Bell nonlocality with a single shot," (2020), <http://arxiv.org/abs/2005.13418> arXiv:2005.13418 [quant-ph].
24. Dimić, A. & Dakić, B. Single-copy entanglement detection. *npj Quantum Information* **4**, 1–8 (2018).
25. Saggio, V. et al. Experimental few-copy multipartite entanglement detection. *Nature physics* **15**, 935–940 (2019).
26. Doherty, A. C., Parrilo, P. A. & Spedalieri, F. M. Detecting multipartite entanglement. *Phys. Rev. A* **71**, 032333 (2005).
27. Weilenmann, M., Dive, B., Trillo, D., Aguilar, E. A. & Navascués, M. Entanglement detection beyond measuring fidelities. *Phys. Rev. Lett.* **124**, 200502 (2020b).
28. Hu, X.-M., et al. "Optimized detection of unfaithful high-dimensional entanglement," (2020), <http://arxiv.org/abs/arXiv:2011.02217> arXiv:2011.02217.
29. Horodecki, M., Horodecki, P. & Horodecki, R. Separability of mixed states: necessary and sufficient conditions. *Phys. Lett. A* **223**, 1–8 (1996).
30. Vidal, G. & Werner, R. F. Computable measure of entanglement. *Phys. Rev. A* **65**, 032314 (2002).

Acknowledgements

This work was supported by the Austrian Science fund (F.W.F.) stand-alone project P 30947.

Author contributions

M.W.E.A.A. and M.N. contributed to this work.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to M.W., E.A.A. or M. Nés.

Peer review information *Nature Communications* thanks the anonymous reviewer(s) for their contribution to the peer review of this work.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021, corrected publication 2022