

Quantum automata cannot detect biased coins, even in the limit

Ryan O’Donnell* Guy Kindler†

October 11, 2016

Abstract

Aaronson and Drucker (2011) asked whether there exists a quantum finite automaton that can distinguish fair coin tosses from biased ones by spending significantly more time in accepting states, on average, given an infinite sequence of tosses. We answer this question negatively.

1 Introduction

In a 2011 work, Aaronson and Drucker [AD11] investigated the ability of a finite automaton to distinguish, given an infinite sequence of coin tosses, whether the coins are fair or $(\frac{1}{2} \pm \epsilon)$ -biased. There are several axes of consideration discussed in [AD11], three of which we state here:

1. Whether the automaton is classical (and probabilistic), or quantum.
2. Whether $\epsilon > 0$ is “known” or not; i.e., whether the automaton can depend on ϵ .
3. The mechanism by which the automaton makes its decision. One possibility is that the automaton guesses “biased” by halting, and guesses “fair” by running forever. A laxer possibility is that the automaton always runs forever, with each of its states designated “biased” or “fair”; its final decision is based on the limiting time-average it spends in “biased” vs. “fair” states. We refer to the two mechanisms as “one-sided halting” and “limiting acceptance”.

For example, an old result of Hellman and Cover [HC70] is that even when ϵ is known and limiting acceptance is allowed, a classical automaton needs $\Omega(1/\epsilon)$ states to solve the problem. On the other hand, Aaronson and Drucker made the interesting observation that for every fixed known ϵ , there’s a quantum automaton with just 2 states that solves the problem using one-sided halting. They also showed no quantum automaton with a fixed number of states can solve the problem for every *unknown* ϵ , if the decision mechanism is one-sided halting.

Aaronson and Drucker asked whether the same negative result holds even if the automaton is allowed to use the limiting acceptance decision mechanism. Indeed, for the 48 different variations of the problem they considered, this was the only variant that remained unsolved. In 2014, Aaronson called this question one of the “Ten Most Annoying Problems in Quantum Computing” [Aar14].

In this work, we make the world of quantum computing 10% less annoying by resolving the problem in the negative. Stated informally, our main theorem is the following (a precise phrasing appears below after we give some formal definitions):

*Computer Science Dept., Carnegie Mellon Univ. Supported by NSF grant CCF-1618679. odonnell@cs.cmu.edu

†School of Computer Science and Engineering, Hebrew Univ. of Jerusalem. Supported by BSF grant 2012220. gkindler@cs.huji.ac.il

Theorem 1.1. *There is no quantum finite automaton that has the following property, simultaneously for every $\epsilon \in [-\frac{1}{2}, \frac{1}{2}] \setminus \{0\}$: Given access to an infinite sequence of coin tosses, if the coin is $(\frac{1}{2} + \epsilon)$ -biased then the automaton spends at least $2/3$ of its time guessing “biased”, and if the coin is fair then the automaton spends at least $2/3$ of its time guessing “fair”.*

Proving this theorem involves a careful understanding of the fixed points of quantum channels.

2 Classical and quantum automata

In this section we review the definitions of probabilistic and quantum finite state automata. Although we are ultimately only concerned with quantum automata, we feel it is instructive to also discuss probabilistic automata at the same time. All of our automata will have input alphabet $\Sigma = \{0, 1\}$, which may be thought of as {tails, heads}.

A classical deterministic automaton on alphabet $\Sigma = \{0, 1\}$ has some d *basic-states*,¹ an initial basic-state $i_0 \in [d]$, and transition rules $f_0, f_1 : [d] \rightarrow [d]$. Given a sequence of input symbols $w_1, w_2, w_3, \dots \in \{0, 1\}$, the automaton operates as follows: It starts in basic-state i_0 at time 0. Then, if it is in basic-state i_t at time $t \in \mathbb{N}$, it transitions to basic-state $f_{w_{t+1}}(i_t)$ at time $t + 1$. Automata also typically have their basic-states classified as “accept” or “reject”; we discuss this more later.

One can also consider classical *probabilistic* automata. These have randomized transitions, which can be encoded by a pair of $d \times d$ *stochastic matrices* S_0, S_1 . Now at any time t the automaton can be in a “probabilistic-state”, represented by a length- d probability vector π_t . (An initial probabilistic-state π_0 is also specified.) On reading symbol w_{t+1} , the automaton transitions to the probabilistic-state $\pi_{t+1} = S_{w_{t+1}}\pi_t$.

Finally, the setting for a *quantum* automaton is a d -dimensional Hilbert space \mathcal{H} (which we may think of as having an orthonormal basis of “basic-state vectors” $|1\rangle, \dots, |d\rangle$). At any time t , the automaton has a “quantum-state”, which is a density operator $\rho_t \in \mathcal{B}(\mathcal{H})$. Here $\mathcal{B}(\mathcal{H})$ denotes the set of linear operators on \mathcal{H} , and a density operator means a positive semidefinite operator of trace 1. (Probabilistic-states are the special case of quantum-states in which ρ_t is diagonal with respect to $|1\rangle, \dots, |d\rangle$.) The transition rules are now any two allowable quantum transformations Φ_0, Φ_1 ; i.e., they are *quantum channels* (*superoperators*) on $\mathcal{B}(\mathcal{H})$. Here a quantum channel means a linear map $\Phi : \mathcal{B}(\mathcal{H}) \rightarrow \mathcal{B}(\mathcal{H})$ that is completely positive and trace-preserving; an equivalent condition is that there exist (non-unique) *Kraus operators* $K_1, \dots, K_r \in \mathcal{B}(\mathcal{H})$ with $\sum_{i=1}^r K_i^\dagger K_i = \mathbb{1}$ such that $\Phi(\rho) = \sum_{i=1}^r K_i \rho K_i^\dagger$. (For more on quantum channels, see e.g. [Wol12].) Again, an initial quantum-state ρ_0 is given, and on reading symbol w_{t+1} , the automaton transitions from quantum-state ρ_t to quantum-state $\rho_{t+1} = \Phi_{w_{t+1}}(\rho_t)$.

Automata with random inputs. This paper is concerned with automata whose inputs are infinite sequences of p -biased coin tosses, $p \in [0, 1]$. More formally, we always assume the input symbols $w_1, w_2, w_3, \dots \in \{0, 1\}$ are chosen independently at random with $\Pr[w_t = 1] = p$. Because of this assumption, we can give a simplified formalization of probabilistic and quantum automata. In the case of probabilistic automata, at each time step (independently) we apply S_1 with probability p and S_0 with probability $1 - p$. It is clear that this is equivalent to simply applying the stochastic matrix $S_p := pS_1 + (1 - p)S_0$ at each time step. In other words, the probabilistic-state of a

¹There is an unfortunate terminology clash involving the word “state” — in automata theory, “states” are the basic vertices in automaton graphs, whereas in quantum mechanics a “state” usually means the “mixed quantum state” or “density operator” of a given system. Throughout we’ll refer to the former as “basic-states” and the latter as “quantum-states”.

probabilistic automaton after t time steps is simply $S_p^t \pi_0$. The setup is precisely equivalent to a Markov chain on $[d]$ with transition matrix S_p .

Similarly for quantum automata, at each time step we apply Φ_1 with probability p and Φ_0 with probability $1 - p$; this is physically equivalent to simply applying the channel $\Phi_p := p\Phi_1 + (1 - p)\Phi_0$ at each time step. (This is ultimately because being in quantum-state ρ with probability p and quantum-state ρ' with probability $1 - p$ is physically equivalent to being in quantum-state $p\rho + (1 - p)\rho'$.) Thus the quantum-state of a probabilistic automaton after t time steps is simply $\Phi_p^t(\rho_0)$; we have here the quantum analogue of a Markov chain.

Automaton acceptance probability. As discussed in Section 1, we will be considering “limiting acceptance”, the most relaxed possible notion for automaton acceptance. We first define this in the context of probabilistic automata. Here, each basic-state in $[d]$ is classified as either guessing “Fair” or “Biased”. We write $e_{\text{fair}} \in \mathbb{R}^d$ for the 0-1 indicator of the Fair states. Thus if the automaton is in probabilistic-state $\pi \in \mathbb{R}^d$, the probability it is in a Fair basic-state is $\langle e_{\text{fair}}, \pi \rangle$. We then consider, for a sequence of T coin tosses, the *average* probability with which the automaton is in a Fair basic-state:

$$f_T(p) := \frac{1}{T} \sum_{t=1}^T \langle e_{\text{fair}}, S_p^t \pi_0 \rangle = \left\langle e_{\text{fair}}, \left(\frac{1}{T} \sum_{t=1}^T S_p^t \right) \pi_0 \right\rangle.$$

Finally, we consider the limiting value of this probability:

$$f(p) := \lim_{T \rightarrow \infty} f_T(p) = \langle e_{\text{fair}}, S_p^\infty \pi_0 \rangle, \quad \text{where } S_p^\infty := \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T S_p^t.$$

Here we relied on the well-known fact that the limiting matrix S_p^∞ exists. (In fact, S_p^∞ is also a stochastic matrix, and it acts by projection onto the 1-eigenspace of S_p ; we discuss this further in Section 3.) One may then say that the probabilistic automaton “guesses Fair in the limit” if $f(p) \geq \frac{2}{3}$, and “guesses Biased in the limit” if $f(p) \leq \frac{1}{3}$. (It may be considered “indecisive” otherwise.)

The definitions for a quantum automaton are extremely similar. The automaton is assumed to come equipped with an “acceptance POVM”, $\{E_{\text{fair}}, \mathbb{1} - E_{\text{fair}}\}$. (Here $E_{\text{fair}} \in \mathcal{B}(\mathcal{H})$ is any operator satisfying $0 \leq E_{\text{fair}} \leq \mathbb{1}$, and $\mathbb{1}$ denotes the identity operator.) If the automaton is in quantum-state ρ , the probability of it measuring “Fair” is $\langle E_{\text{fair}}, \rho \rangle := \text{tr}(E_{\text{fair}}^\dagger \rho)$. We can then again define the limiting average probability of guessing “Fair” via

$$f_T(p) := \frac{1}{T} \sum_{t=1}^T \langle E_{\text{fair}}, \Phi_p^t \pi_0 \rangle = \left\langle E_{\text{fair}}, \left(\frac{1}{T} \sum_{t=1}^T \Phi_p^t \right) \pi_0 \right\rangle,$$

$$f(p) := \lim_{T \rightarrow \infty} f_T(p) = \langle E_{\text{fair}}, \Phi_p^\infty \pi_0 \rangle, \quad \text{where } \Phi_p^\infty := \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \Phi_p^t. \quad (1)$$

Again, it is known that the limiting operator Φ_p^∞ exists; this is explicitly discussed in Section 3. As before, one may say that the quantum automaton “guesses Fair in the limit” if $f(p) \geq \frac{2}{3}$, and “guesses Biased in the limit” if $f(p) \leq \frac{1}{3}$.

We may now state the main theorem of this paper:

Theorem 2.1. *In the setting of quantum automata reading p -biased bits (as described above), the function f from (1) is a continuous function of $p \in (0, 1)$.*

This theorem is a formal strengthening of Theorem 1.1, our negative result for coin distinguishing stated in Section 1. For example, it implies that if an automaton guesses “Fair” in the limit” for $p = \frac{1}{2}$, then for all sufficiently small ϵ it *cannot* guess “Biased” in the limit for $p = \frac{1}{2} \pm \epsilon$. In fact, we get the inability of quantum automata to distinguish p -biased and $(p \pm \epsilon)$ -biased coins with limiting acceptance for any fixed $p \in (0, 1)$. As noted in [AD11], this is sharp in the sense that there is a trivial 2-state deterministic classical automaton that distinguishes a 0-biased coin from any ϵ -biased coin, even with one-sided halting.

3 Outline of the proof

Here we give an outline of the proof of Theorem 2.1. At the same time, it will be instructive to outline the analogous proof in the special case of probabilistic automata. To prove that the limiting acceptance probability $f(p)$ from (1) is continuous for $p \in (0, 1)$, it is enough to prove the following:

Theorem 3.1. Φ_p^∞ is continuous for $p \in (0, 1)$.

Here for definiteness we can take the metric on channels induced by the operator norm on $\mathcal{B}(\mathcal{H})$; Theorem 2.1 then follows because matrix multiplication and inner product are continuous.

Now is a good time to review the properties of Φ_p^∞ . In general, let Φ denote any channel on $\mathcal{B}(\mathcal{H})$. Then the following are known [Wol12, Prop. 6.3] (and easy) facts: First, $\Phi^\infty := \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \Phi^t$ exists and is itself a channel. Second, as an operator Φ^∞ , acts as projection onto the fixed points $V_1(\Phi)$ of Φ . Here we are using the following notation:

Notation 3.2. For any operator A we write $V_1(A)$ for the eigenspace of A with eigenvalue 1, i.e., the *invariant* subspace for A .

As mentioned earlier, the analogous statements are true regarding S^∞ , when S is a stochastic operator. (In both the probabilistic and quantum cases, the essential point is that the operator in question has spectral radius 1.)

Returning to Theorem 3.1, certainly $\Phi_p = p\Phi_1 + (1-p)\Phi_0$ varies continuously for $p \in [0, 1]$. But what we need to prove is that the *invariant subspace* $V_1(\Phi_p)$ of Φ_p varies continuously for $p \in (0, 1)$. There is one obvious potential obstruction: the *dimension* of $V_1(\Phi_p)$ might change as p varies. (As we will see, this is actually the only obstruction.) Now in general, slightly perturbing a matrix *can* change the dimension of its 1-eigenspace. However we are not concerned with completely general perturbations: we are just considering all the convex combinations of two fixed channels Φ_0, Φ_1 . The main technical theorem in our paper will be the following:

Theorem 3.3. For any channels Φ_0, Φ_1 , the dimension $\dim V_1(\Phi_p)$ is the same for all $p \in (0, 1)$.

We will discuss the intuition for this theorem below. But first we will observe that Theorem 3.1 is an elementary linear-algebraic consequence of Theorem 3.3. This deduction of Theorem 3.1 from Theorem 3.3 is a little more familiar if we consider $\mathbb{1} - \Phi_p$ rather than Φ_p . Then Φ_p^∞ is the projection onto the kernel of $\mathbb{1} - \Phi_p$, and it is elementary that, given a continuously-parameterized family of matrices like $p \mapsto \mathbb{1} - \Phi_p$, the kernel varies continuously wherever the nullity (in this case, $\dim V_1(\Phi_p)$) is locally constant. For a simple explicit proof see, e.g., [use15].

Thus all that remains in this work is to prove Theorem 3.3. We will do this in Section 4, but first we provide some intuition and introduce a key definition, that of *combinatorially equivalent channels*.

3.1 Intuition for Theorem 3.3

All of our discussion so far applies equally to probabilistic automata defined by stochastic matrices S_0, S_1 . So let us first consider the analogue of Theorem 3.3 in this case. Here we have a family of Markov chains defined by $S_p = pS_1 + (1-p)S_0$ and we want to consider the dimension of their invariant subspaces. It is well known that the invariant subspace $V_1(S)$ of the Markov chain defined by S is spanned by a linearly independent set of invariant *probabilistic-states*. Thus $\dim V_1(S)$ is equal to the number of linearly independent (“fundamentally different”, one might say) invariant distributions.

In the study of Markov chains, it’s popular to focus on the irreducible case, in which case there is a unique invariant probability distribution; i.e., $\dim V_1(S) = 1$. However in general we must consider reducible Markov chains (the “mathematically annoying case”, as Hellman and Cover [HC70] put it). Fortunately, the theory of reducible Markov chains is well developed, and it is known that there is one linearly independent invariant distribution per every *communication class* of the Markov chain. Here the “communication classes” of the Markov chain defined by S are precisely the strongly connected components of the underlying digraph on $[d]$; i.e., the graph which has a directed edge (i, j) whenever $S_{ij} \neq 0$. Given this theory, it is easy to deduce the analogue of Theorem 3.3; the point is that for any fixed S_0, S_1 , *the underlying digraph of S_p is the same for all $p \in (0, 1)$* . Since $S_p = pS_1 + (1-p)S_0$, an edge (i, j) is present in S_p is present if and only if it is present in both S_0 and S_1 . Thus S_p has the same set (hence number) of communication classes for all $p \in (0, 1)$, as needed.

In this paper, we show there is an analogous sequence of ideas in the quantum case, using some of the recently developed theory of fixed points of quantum channels. Given a quantum channel Φ , it is known [Wol12, Cor. 6.5] that $V_1(\Phi)$ is always spanned by linearly independent *quantum-states*. The analogous notion to communication classes is that of *minimal enclosures*. Further, similar to how the communication classes of a Markov chain are determined only by the nonzero pattern of its transition matrix, the minimal enclosures of a quantum channel are determined only by its Kraus operators. We then make use of the fact that all the convex combinations Φ_p of two channels Φ_0, Φ_1 have related Kraus operators. Specifically, we introduce the following notion:

Definition 3.4. We will say that two channels Φ and $\widehat{\Phi}$ (with the same Hilbert space \mathcal{H}) are *combinatorially equivalent* if there are Kraus operators K_1, \dots, K_r for Φ and $\widehat{K}_1, \dots, \widehat{K}_r$ for $\widehat{\Phi}$ such that each K_i is proportional to some $\widehat{K}_{i'}$ and vice versa.

Given channels Φ_0, Φ_1 with Kraus operators $\{K_i^{(0)} : i \in [r_0]\}, \{K_j^{(1)} : j \in [r_1]\}$ respectively, the channel $\Phi_p = p\Phi_1 + (1-p)\Phi_0$ has Kraus operators $\{\sqrt{1-p}K_i^{(0)} : i \in [r_0]\} \cup \{\sqrt{p}K_j^{(1)} : j \in [r_1]\}$. Thus the channels Φ_p are all pairwise combinatorially equivalent for $p \in (0, 1)$ (though not necessarily for $p \in \{0, 1\}$). To show Theorem 3.3, it therefore suffices to show the following more general result:

Theorem 3.5. *Suppose Φ and $\widehat{\Phi}$ are combinatorially equivalent. Then $\dim V_1(\Phi) = \dim V_1(\widehat{\Phi})$.*

4 The last step: proof of Theorem 3.5

To prove Theorem 3.5, we use some known results concerning the decomposition of a quantum channel into irreducible components, and the structure of its invariant quantum-states. We will specifically use the key decomposition theorem appearing variously as [Wol12, Theorem 6.14], [BN12, Theorem 7], [CP15, Theorem 7.2].

Let Φ denote a quantum channel on $\mathcal{B}(\mathcal{H})$ with Kraus operators K_1, \dots, K_r . We are interested in $m = \dim V_1(\Phi)$, the dimension of the space of Φ ’s fixed points. As Φ is a quantum channel, it is

known [Wol12, Prop. 6.1] that its spectral radius is 1 and that it has at least one eigenvalue equal to 1; thus $m \geq 1$. As mentioned, it is also known [Wol12, Cor. 6.5] that $V_1(\Phi)$ is always spanned by some m linearly independent *quantum-states*.

If ρ is a quantum-state, its *support* $\text{supp}(\rho)$ is simply the range of ρ as a subspace of \mathcal{H} . The *recurrent subspace* for Φ is the subspace of \mathcal{H} defined by

$$\mathcal{R} = \text{span}\{\text{supp}(\rho) : \rho \text{ is an invariant quantum-state}\}.$$

The orthogonal complement of \mathcal{R} in \mathcal{H} is denoted \mathcal{D} ; this is the *decaying* (or *transient*) subspace. A subspace $\mathcal{V} \subseteq \mathcal{H}$ is called an *enclosure* if $\text{supp}(\rho) \subseteq \mathcal{V} \implies \text{supp}(\Phi(\rho)) \subseteq \mathcal{V}$ for all quantum-states ρ . We can relate this concept to Kraus operators via the following equivalence:

Fact 4.1. ([CP15, Proposition 4.4].) \mathcal{V} is an enclosure if and only if $K_i \mathcal{V} \subseteq \mathcal{V}$ for all Kraus operators K_i .

An enclosure \mathcal{V} is called *minimal* if it is nonzero and all enclosures $\mathcal{V}' \subseteq \mathcal{V}$ are equal to either $\{0\}$ or \mathcal{V} . It is also known [BN12, Prop. 15] that a subspace of \mathcal{H} is a minimal enclosure if and only if it is the support of an extremal invariant quantum-state, meaning one that cannot be written as a nontrivial convex combination of two distinct invariant quantum-states. One consequence is that

$$\mathcal{R} = \text{span}\{\text{supp}(\rho) : \rho \text{ is an extremal invariant quantum-state}\} = \text{span}\{\mathcal{V} : \mathcal{V} \text{ is a minimal enclosure}\}. \quad (2)$$

The theorems [Wol12, Theorem 6.14], [BN12, Theorem 7], [CP15, Theorem 7.2] characterize $V_1(\Phi)$ and the quantum-states therein in slightly different ways. To explain, we make some definitions.

Definition 4.2. (In this definition, $k, m_1, \dots, m_k, d_1, \dots, d_k$ denote positive integers.)

Given Φ , we define a *minimal enclosure decomposition* to be an orthogonal decomposition of \mathcal{H} into subspaces

$$\mathcal{H} = \mathcal{D} \oplus \bigoplus_{i=1}^k \mathcal{W}_i, \quad \text{where } \mathcal{W}_i = \bigoplus_{j=1}^{m_i} \mathcal{V}_{i,j} \quad (3)$$

in which \mathcal{D} is the decaying subspace for Φ , each $\mathcal{V}_{i,j}$ is a minimal enclosure, $\dim \mathcal{V}_{i,j} = d_i$ for all $1 \leq j \leq m_i$, and finally the following property holds: *For any minimal enclosure \mathcal{X} of Φ and any $1 \leq i \leq k$, if \mathcal{X} is not orthogonal to \mathcal{W}_i then $\mathcal{X} \subseteq \mathcal{W}_i$. (In particular, if $m_i = 1$ then \mathcal{X} must equal \mathcal{W}_i .)*

Definition 4.3. Suppose we have a minimal enclosure decomposition for Φ as above. Fix any ordered orthogonal basis for \mathcal{H} compatible with (3) (meaning the first $\dim \mathcal{D}$ elements span \mathcal{D} , the next $m_1 d_1$ elements come in m_1 groups of d_1 spanning $\mathcal{V}_{1,1}, \dots, \mathcal{V}_{1,m_1}$ respectively, etc.). Let $X \in \mathcal{B}(\mathcal{H})$, and think of X in its matrix form with respect to the ordered basis.

Then we say that X *respects the minimal enclosure decomposition* if X is block-diagonal with blocks corresponding to $\mathcal{D}, \mathcal{W}_1, \dots, \mathcal{W}_k$, and furthermore X is 0 on the \mathcal{D} -block and is of the form $A_i \otimes \rho_i$ on the \mathcal{W}_i -block for some $A_i \in \mathcal{C}^{m_i \times m_i}$ and some strictly positive density matrix $\rho_i \in \mathcal{C}^{d_i \times d_i}$. In symbols,

$$X = 0 \oplus \bigoplus_{i=1}^k A_i \otimes \rho_i.$$

(We remark that the property of respecting the minimal enclosure decomposition does not depend on the choice of the compatible orthogonal basis.)

In combination, [Wol12, Theorem 6.14], [BN12, Theorem 7] state the following:²

Theorem 4.4. *Given any channel Φ , there exists a minimal enclosure decomposition as in (3) such that $V_1(\Phi)$ consists precisely of all $X \in \mathcal{B}(\mathcal{H})$ that respect the decomposition. (An immediate consequence is that $m = \dim V_1(\Phi) = \sum_i m_i^2$.) Finally, the quantum-states that are invariant are precisely all such X with $A_i = \lambda_i \sigma_i$, where $\sigma_1, \dots, \sigma_k$ are density matrices and $\lambda_1, \dots, \lambda_k$ are nonnegative reals summing to 1.*

The statement of [CP15, Theorem 7.2] is slightly different:³

Theorem 4.5. *Given any channel Φ , at least one minimal enclosure decomposition exists. Furthermore, given any minimal enclosure decomposition*

$$\mathcal{H} = \mathcal{D} \oplus \bigoplus_{i=1}^{\widehat{k}} \widehat{\mathcal{W}}_i, \quad \text{where } \widehat{\mathcal{W}}_i = \bigoplus_{j=1}^{\widehat{m}_k} \widehat{\mathcal{V}}_{i,j},$$

every invariant quantum-state for Φ respects it. (As an immediate consequence, $m = \dim V_1(\Phi) \leq \sum_i \widehat{m}_i^2$.)

We are now able to give the proof of Theorem 3.5.

Proof of Theorem 3.5. Write $m = \dim V_1(\Phi)$ and $\widehat{m} = \dim V_1(\widehat{\Phi})$. Since Φ and $\widehat{\Phi}$ play symmetric roles, it suffices to show $\widehat{m} \leq m$. Apply Theorem 4.4 to Φ , obtaining a minimal enclosure decomposition as in (3). We have $m = \sum_{i=1}^k m_i^2$. We claim that this decomposition is also a minimal enclosure decomposition for $\widehat{\Phi}$. This will finish the proof of $\widehat{m} \leq m$, by Theorem 4.5.

To see the claim, we first observe that every enclosure \mathcal{V} for Φ is an enclosure for $\widehat{\Phi}$ (and vice versa). This follows from Fact 4.1: \mathcal{V} satisfies $K_i \mathcal{V} \subseteq \mathcal{V}$ for each Kraus operator K_i of Φ , and hence the same is true for the Kraus operators $\widehat{K}_{i'}$ of $\widehat{\Phi}$, by combinatorial equivalence of Φ and $\widehat{\Phi}$. It then follows by definition that every *minimal* enclosure for Φ is also a minimal enclosure for $\widehat{\Phi}$ (and vice versa). Finally, the claim now follows because Φ and $\widehat{\Phi}$ have the same decaying subspace (by (2)) and because Definition 4.2 of minimal enclosure decompositions depends *only* on which subspaces of \mathcal{H} are minimal enclosures. \square

References

- [Aar14] Scott Aaronson. The NEW ten most annoying questions in quantum computing, May 2014. <http://www.scottaaronson.com/blog/?p=1792>. 1
- [AD11] Scott Aaronson and Andrew Drucker. Advice coins for classical and quantum computation. In *Proceedings of the 38th Annual International Colloquium on Automata, Languages and Programming*, pages 61–72, 2011. 1, 2
- [BN12] Bernhard Baumgartner and Heide Narnhofer. The structures of state space concerning quantum dynamical semigroups. *Reviews in Mathematical Physics*, 24(2):1250001, 2012. 4, 4, 4, 2

²[Wol12] deals with the invariant subspace whereas [BN12] deals with the invariant quantum-states. The fact that the ρ_i 's are strictly positive is in [Wol12]. Finally, [BN12] does not explicitly show that the minimal enclosure decomposition satisfies the last, italicized, condition in Definition 4.2. However, it's implicit and it's easy to deduce: we know that any minimal enclosure \mathcal{X} is the support of some extremal invariant quantum-state ρ , and it's clear that if this support is not entirely within a single \mathcal{W}_i -block then ρ would not be extremal.

³The first statement of this theorem is [CP15, Proposition 7.1], except that that Proposition does not include the final, italicized, property of Definition 4.2 for those i with $m_i = 1$. However it is evident from the proof that this is an oversight; a person communication from the authors confirmed this.

- [CP15] Raffaella Carbone and Yan Pautrat. Irreducible decompositions and stationary states of quantum channels. Technical Report 1507.08404, arXiv, 2015. 4, 4.1, 4, 4, 3
- [HC70] Martin Hellman and Thomas Cover. Learning with finite memory. *Annals of Mathematical Statistics*, 41:765–782, 1970. 1, 3.1
- [use15] user1551. Continuity of the basis of the null space, March 2015. <http://math.stackexchange.com/a/1203782>. 3
- [Wol12] Michael Wolf. Quantum channels & operations: guided tour, 2012. <http://www-m5.ma.tum.de/foswiki/pub/M5/Allgemeines/MichaelWolf/QChannelLecture.pdf>. 2, 3, 3.1, 4, 4, 4, 2