

Constrained and Vanishing Expressivity of Quantum Fourier Models

Hela Mhiri^{1,2}, Léo Monbroussou^{1,3}, Mario Herrero-González^{1,4}, Slimane Thabet^{1,5}, Elham Kashefi^{1,4}, and Jonas Landman^{4,6}

¹Sorbonne Université, CNRS, LIP6, 75005 Paris, France

²ENSTA Paris, Institut Polytechnique de Paris, France

³CEMIS, Direction Technique, Naval Group, 83190 Ollioules, France

⁴Quantum Software Lab, School of Informatics, University of Edinburgh, United Kingdom

⁵PASQAL SAS, 7 avenue Léonard de Vinci, 91300 Massy, France

⁶QC Ware, Palo Alto, USA and Paris, France

In this work, we highlight an unforeseen behavior of the expressivity of Parameterized Quantum Circuits (PQCs) for machine learning. A large class of these models, seen as Fourier series whose frequencies are derived from the encoding gates, were thought to have their Fourier coefficients mostly determined by the trainable gates. Here, we demonstrate a new correlation between the Fourier coefficients of the quantum model and its encoding gates. In addition, we display a phenomenon of vanishing expressivity in certain settings, where some Fourier coefficients vanish exponentially as the number of qubits grows. These two behaviors imply novel forms of constraints which limit the expressivity of PQCs, and therefore imply a new inductive bias for quantum models. The key concept in this work is the notion of a frequency redundancy in the Fourier series spectrum, which determines its importance. Those theoretical behaviors are observed in numerical simulations.

1 Introduction

Quantum Machine Learning (QML) is an important field of study as an application of quantum computing [1]. While many fault-tolerant algorithms propose significant advantages over their classical counterparts, many technological milestones must be reached

Hela Mhiri[✉]: hela.mhiri@outlook.fr

for implementation. Variational quantum algorithms [2] are promising candidates for near term QML methods. Indeed, a popular approach consists of training Parametrized Quantum Circuits (PQCs) as neural networks. To do so, the classical data must be embedded in the Hilbert space through encoding gates, whereas trainable gates that depend on internal parameters will be optimized during the hybrid training procedure.

Many studies have been conducted to understand the potential and limitations of quantum models. Multiple works focus on the trainability of such models and highlight the issues of local minima and barren plateaus [3, 4, 5, 6, 7, 8].

On the other hand, another fundamental question concerns the expressivity of these models, namely which hypothesis class the quantum model is exploring. From the seminal paper [9], we know that if one considers an encoding scheme where the classical input is encoded as the time evolution of some Hamiltonian, the quantum model generated by the PQC can be described as a Fourier series in the classical input. The spectrum is determined by the encoding layers while the Fourier coefficients are mainly controlled by the trainable layers (See Section 2 for more details).

In this work, we highlight a new connection between the Fourier coefficients and the encoding gates as illustrated in Fig.1. This connection is made through the new concept of frequency *redundancy*, which captures the number of times a frequency appears in the spectrum (see Definition 1 for details).




Figure 1: Parameterized quantum models with classical input x and parameter vector θ , a) can be seen as Fourier series in the classical input with frequencies $\omega \in \Omega$ and associated Fourier coefficients $c_{\omega}(\theta)$. b) illustrates the relation between the frequencies redundancies $|R(\omega)|$, i.e. the number of times a frequency appears in the spectrum, and their Fourier coefficients variance $\text{Var}[c_{\omega}(\theta)]$. This connection constrains the expressivity of quantum Fourier models.

The concept of frequency redundancy emerged from previous works [10, 11] attempting to find a classical approximation to any quantum Fourier model. Indeed, one could use the same frequencies in the spectrum, or a sample from them, to train a classical Fourier series that is guaranteed to have the same or better performance than the quantum model on generic machine learning tasks. If this question remains open in general [12, 13], it was observed in [10] that some frequencies were more important to include in the Fourier series, as their Fourier coefficients had always a greater contribution.

In this paper, we formally establish that some Fourier coefficients are less constrained than others depending on the choice of the encoding Hamiltonians. Using the notion of redundancy, we show that the encoding strategy leads to an *inductive bias* in the quantum model, as the Fourier coefficients with high redundancies will have a greater impact. To prove the connection between a frequency redundancy and its importance in the Fourier model, we show that the variance of any Fourier coefficient is highly dependent on the redundancy of its frequency. We further use this result to establish the *vanishing expressivity* phenomenon, whereby some or all Fourier coefficients can suffer from exponentially vanishing variance as the number of qubits increases. This vanishing expressivity phenomenon can be interpreted as the effective model expressivity. Namely, it implies that while using an exponentially big spectrum may seem at first as a trivial strategy to increase the model's expressivity, one should examine more carefully the *effective* expressivity of the model. Moreover,

we establish another constraint on the 2-norm of the Fourier coefficients vector that holds for any quantum Fourier Model.

The manuscript is structured as follows. In Section 2, we introduce the framework we consider in the study of quantum models from PQCs. We also define relevant mathematical tools and concepts required for our study. In Section 3, we present our main results showing the inductive bias in the quantum model that arises from spectrum properties. In Section 4, we discuss in more details the impact of the proven Fourier model constraints on the model behavior. Finally, we present numerical simulations in Section 5 supporting our theoretical results and we conclude in Section 6.

Related Work

In [9], it has been shown that a large class of parameterized quantum models with classical data encoding can be expressed as Fourier series. Moreover, by considering reuploading schemes, the spectrum size can be efficiently increased leading to quantum models that serve as universal function approximators [14, 9]. Hence, this framework arises as a powerful tool to further assess the performance and potential of such quantum models. Specifically, much focus has been accorded to the study of the quantum spectrum and its characteristics [9, 15, 16] by exploring different choices of the encoding Hamiltonians. Besides, in [17], the author proposed to learn the frequencies of the quantum model that fit best the learning task at hand. On the other hand, a variety of works used this Fourier representation to propose dequantization

schemes of quantum models. In [10], an efficient construction of a classical model based solely on the circuit description was proposed. The classical model was expected to have similar performances to the quantum model on the training data. However, it has been pointed out in [12, 13] that both models may not necessarily converge to the same solution, leading to different generalization performances. This statement has been further examined in [18] and in [19] where the authors show that there exist quantum models that cannot be dequantized by any efficient method. Thus, it is now highly important to understand the inductive bias of quantum Fourier models that differentiate them from their classical “dequantizers”.

In this work, we focus on studying the bias in the quantum model Fourier coefficients and how it is linked to the structure of the quantum circuit (i.e. encoding strategy, observable, etc.). In this line of work, [15] related the phenomenon of benign overfitting in quantum interpolating Fourier models to the frequency distribution and state preparation showing that the encoding strategy provides an inductive bias that impacts directly the generalization performance. Moreover, [20] provided encoding-dependent generalization bounds for such models. Nonetheless, these works provide generic results about quantum Fourier models performance without taking into account the parameterized Fourier coefficients structure in finer detail. In a more recent work [21], the authors explored encoding dependent concentration of the Fourier coefficients in the quantum reuploading scheme under the Haar measure assumption proving a phenomenon of vanishing expressivity for high frequencies in the limit of an infinite depth circuit. In addition, [22] explored exponential concentration sources in quantum Reservoir computing by using the Fourier representation. In this work, we further explore the degree to which the Fourier coefficients concentrate around their mean under different assumptions on the trainable unitaries of the circuit and show that the encoding strategy constraints the Fourier coefficients’ variance for arbitrary trainable unitaries, limiting their theoretical expressivity.

2 Framework

In this Section, we present the framework used throughout this work. We first describe the considered circuit structure and then recall how to define the Fourier representation of the associated quantum model in Section 2.1. Moreover, we introduce the notion of frequency redundancy and how it could be tuned through the choice of the encoding strategy (i.e. encoding Hamiltonians). Finally, in Section 2.2, we discuss key metrics for quantifying the expressivity of quantum models.

2.1 Quantum Fourier Model

We consider a standard supervised learning task, where a parameterized function f , called *model*, must be optimized to match targets in a finite dataset. We define *quantum models* on n qubits as the family of parameterized functions $f : \mathcal{X} \times \Theta \rightarrow \mathbb{R}$ obtained by measuring the expectation value of some hermitian observable O , such that

$$f(x, \theta) = \langle 0 | U(x, \theta)^\dagger O U(x, \theta) | 0 \rangle, \quad (1)$$

where $U(x, \theta)$ is a 2^n -dimensional unitary, $\theta \in \Theta$ is the vector of trainable parameters and $x = (x_1, \dots, x_D) \in \mathcal{X} \subset \mathbb{R}^D$ is the classical data vector with D components.

We consider a circuit unitary composed of alternating *encoding* and *trainable* layers as depicted in Fig.1 of the form

$$U(x, \theta) = W^{L+1}(\theta) \left[\prod_{l=1}^L S^l(x) W^l(\theta) \right], \quad (2)$$

where L is the total number of circuit layers (i.e. a circuit layer is made of an encoding layer and a trainable layer), $W^l(\theta)$ s are formed by trainable gates depending on the parameter vector θ , which is optimized during training. On the other hand, $S^l(x)$ s only depend on input data values and are fixed during the training. This scheme where the classical data is encoded using L alternated layers was originally proposed in [14] and the associated model of the form in Eq.(1) is denoted by quantum reuploading model.

In the remainder of this work, we adopt the *Hamiltonian encoding* strategy where the classical input components are encoded as the time evolution of some Hamiltonians $H_l^{(k)}$ such that the encoding unitaries are of the form

$S^l(x) = \prod_{k=1}^D e^{-ix_k H_l^{(k)}}$. From the seminal work [9], we know that if one considers the Hamiltonian encoding strategy, then the quantum model generated by the circuit described in Eq.(2)xxx can be written as a truncated Fourier Series with a spectrum Ω depending on the eigenvalues of the encoding Hamiltonians. The associated Fourier coefficients depend mainly on the trainable unitaries. Under these assumptions, we call the obtained model a quantum Fourier model (QFM), which is expressed as follows

$$f(x, \theta) = \sum_{\omega \in \Omega} c_{\omega}(\theta) e^{i\omega^T x}. \quad (3)$$

The above equation tends to imply that $c_{\omega}(\theta)$ is solely determined by the parameterized unitaries $W^l(\theta)$ s. However, in this work, we show that the dependence of the Fourier coefficients on the encoding gates is more subtle. First, we begin by highlighting the relation between the frequencies and the encoding Hamiltonian's eigenvalues. To do so, we expand here Eq.(1) in the case of one-dimensional input vectors ($D = 1$). The generalization of the Fourier expansion to higher dimensional input is detailed in Appendix A.1.1. We denote by $d = 2^n$ the dimension of the Hilbert space (with n the number of qubits) and we assume without loss of generality¹ that $S^l(x) = \text{diag}(\lambda_1^l, \dots, \lambda_d^l)$. We also drop the explicit dependence on θ in $W^l(\theta)$ for ease of notation and obtain

$$c_{J,J'} = \langle 0 | \prod_{l=1}^L W^{l\dagger} |j'_l\rangle\langle j_l| \tilde{O} \prod_{l=L}^1 |j'_l\rangle\langle j_l| W^l |0\rangle, \quad (4)$$

$$f(x, \theta) = \sum_{J,J' \in \llbracket 1, d \rrbracket^L} c_{J,J'} e^{-ix \left(\sum_{l=1}^L \lambda_{j_l}^l - \lambda_{j'_l}^l \right)},$$

where we introduce the shorthand $\tilde{O} = W^{L+1\dagger} O W^{L+1}$. We note that $J = (j_1, \dots, j_L)$ is a multi-index where each component $j_l \in \llbracket 1, d \rrbracket$ refers to the choice of the j^{th} eigenvalue of the Hamiltonian H_l (J maps to a path in the tree from Fig.2).

From Eq.(4), we see that the spectrum Ω can be constructed from the eigenvalues of the

¹One can simply consider that $S^l(x) = P D P^{-1}$ and inject P in the expression of W^l and P^{-1} in W^{l+1} .

encoding Hamiltonians in each layer as follows

$$\Omega = \left\{ \sum_{j_l \in J} \lambda_{j_l}^l - \sum_{j'_l \in J'} \lambda_{j'_l}^l \mid (J, J') \in \llbracket 1, d \rrbracket^{2L} \right\}. \quad (5)$$

We note here that the spectrum Ω contains redundant frequencies by construction but in the remainder of this work we consider that Ω denotes the set of distinct frequencies.

As shown in Fig.2, the choice of two *paths* (J, J') in the quantum spectrum tree leads to the generation of a frequency ω by computing the difference between the sum of eigenvalues over the pair of paths. One can easily notice that several pairs of paths could lead to the same frequency. This can happen if an eigenvalue is degenerate, or if several paths of the tree end at the same leaf value (sum of eigenvalues over a path), or eventually if several pairs have the same difference value. The number of those paths evolves with the choice of the different encoding Hamiltonians, the degeneracy of their eigenvalues, and the number L of circuit layers.

By grouping paths (J, J') in Eq.(4) leading to a certain frequency, we identify the Fourier coefficient expression, i.e. $c_{\omega} = \sum_{J,J' \in R(\omega)} c_{J,J'}$. We also formally define the *Frequency Generator* $R(\omega)$ as the set of all paths leading to the generation of the frequency ω and denote the cardinality of this set by the frequency *redundancy* $|R(\omega)|$. As we will demonstrate in Section 3, the redundancy of a frequency will have a crucial role in characterizing the expressivity of QFMs.

Definition 1 (Frequency Generator). *Consider an L -layer quantum Fourier model as described in Eq.(1-5) with one-dimensional input vectors ($D = 1$). For a given frequency ω , we define its generator $R(\omega)$ as the set of eigenvalue indices leading to the generation of ω .*

$$R(\omega) = \left\{ (J, J') \in \llbracket 1, d \rrbracket^{2L} \mid \sum_{j_l \in J} \lambda_{j_l}^l - \sum_{j'_l \in J'} \lambda_{j'_l}^l = \omega \right\} \quad (6)$$

We call the redundancy of a frequency ω the size of its Generator, i.e. $|R(\omega)|$.

Since $\sum_{\omega \in \Omega} |R(\omega)| = 2^{2nL} = d^{2L}$ by construction, the normalized redundancies $\{|\tilde{R}(\omega)|\}_{\omega \in \Omega} := \left\{ \frac{|R(\omega)|}{d^{2L}} \right\}_{\omega \in \Omega}$ define a natural weighted probability distribution over the spectrum Ω . Therefore, by considering different




Figure 2: **Quantum spectrum tree.** The frequencies of a quantum Fourier model are derived from the eigenvalues of the encoding Hamiltonians. Each path in the quantum spectrum tree represents a different choice of eigenvalues $\lambda_{j_\ell}^\ell$ from Hamiltonians H_ℓ s. Some edges are duplicated, expressing that some eigenvalues are degenerate for some Hamiltonians (not at scale). Each frequency ω in the model comes from the difference of two paths in the tree, as shown in the example in red. For a quantum model acting on $n = 2$ qubits with $L = 3$ circuit layers and encoding Hamiltonians H_1, H_2 and H_3 , we give in red a pair of paths in the tree $J = (4, 3, 3), J' = (4, 1, 1)$ generating the frequency $\omega = (\lambda_4^1 + \lambda_3^2 + \lambda_3^3) - (\lambda_4^1 + \lambda_1^2 + \lambda_1^3)$.

encoding Hamiltonians, one can obtain different probability distributions over the spectrum that will impact the behavior of the associated quantum model.

For example, Let's consider the standard case of Pauli encoding [9], where single qubit rotation gates are used to encode the classical input $x \in \mathbb{R}$ as the rotation angle. In this case, the encoding Hamiltonian in each layer is a Pauli string. If the Pauli strings do not contain the identity, then the obtained spectrum is simply $\Omega = \llbracket -nL, nL \rrbracket$. Moreover, it has been shown that the spectrum distribution defined by the redundancies follows a standard Gaussian distribution [15]. Hence, this encoding strategy gives rise to a linear size spectrum (linear in n and L) and concentrates the redundancies in the lower values as detailed in Appendix A.2.

On the contrary, the exponential encoding strategy introduced in [16], which uses scaled Pauli rotations (See Appendix A.2), leads to an exponential size spectrum of consecutive integer frequencies. Specifically, the obtained spectrum is $\Omega = \llbracket -\frac{3^{nL}-1}{2}, \frac{3^{nL}-1}{2} \rrbracket$ and some frequencies (not necessarily high frequencies) have redundancies that do not scale exponentially in n and L . However, if one wants to obtain a fully non-degenerate spectrum (except for the null frequency), then a single circuit layer made of a non-local encoding Hamiltonian must be used as mentioned in [16] and explained in Appendix A.2. This is the case for the *Golomb* encoding

introduced in [15] where the size of the spectrum is exponentially large ($|\Omega| = 2^{\binom{d}{2}} + 1$) and all non zero coefficients have a redundancy of one. A non-degenerate spectrum can also be obtained using a global Hamiltonian with distinct real eigenvalues.

As illustrated in Fig.3, one could choose a particular set of Hamiltonians to design a quantum model with a specific spectrum distribution. In this work, we will show that this choice does not only impact the spectrum of the quantum Fourier model but also the Fourier coefficients concentration. In addition, previous work [10] showed the possibility of classically approximating such PQC for Machine Learning tasks. While having an exponential size spectrum may be a leeway to avoid this classical approximation, we will show in Section 3 that models with a large spectrum tend to have more constrained Fourier coefficients, hence limiting their expressivity and making their classical approximation potentially more efficient.

2.2 Expressivity measure of quantum Fourier models

In this section, we discuss relevant expressivity measures of quantum models of the form in Eq.(1) using a parameterized circuit of the form in Eq.(2) equipped with Hamiltonian encoding. Namely, in the context of supervised learning, the model expressivity is tightly related to the class of functions (in the input data) it can generate.




Figure 3: **Spectrum distributions.** We compare two Hamiltonian encoding strategies leading to very different spectrums (x-axis) and distributions (normalized height of the bars). On the left, we present an illustration of a highly degenerate spectrum (e.g. Pauli encoding spectrum distribution), and on the right an illustration of a weakly degenerate spectrum (e.g. Golomb encoding spectrum distribution).

Consequently, it follows naturally that the model dependence on classical input parameters and trainable parameters is conceptually different. As shown in Eq.(3), the choice of the encoding strategy dictates the general model structure expressed as a truncated Fourier series in the classical input $x \in \mathcal{X}$. Hence, the model expressivity boils down to identifying the spectrum Ω and characterizing the attainable Fourier coefficients values as the trainable parameters $\theta \in \Theta$ vary. This definition of model expressivity can invoke unitary expressivity measures such as the distance of the trainable unitaries distribution (induced by the trainable parameter distribution) to a 2-design on the unitary group [23, 24]. Nevertheless, properly characterizing the expressivity of the quantum model requires further analysis of how the use of *expressive* trainable unitaries $W^l(\theta)$ s translates into a measure of “expressivity” of Fourier coefficients for a fixed spectrum Ω . To do so, we first recall the definition of a unitary approximate 2-design and then define its plausible consequences on Fourier models expressivity. Specifically, we introduce the notion of model exponential concentration and vanishing expressivity in Definition 2 and Definition 3 respectively.

By restricting our attention to the family of parametrized unitaries of the form $\{W^l(\theta)\}_{\theta \in \Theta}$ associated to each trainable unitary in the circuit expression in Eq.(2), this family of unitaries is said to form an ε -approximate 2-design if its distribution over the unitary group induced by the uniform distribution over the parameter space Θ matches the Haar measure up to second moment for some error ε . Formally, defining the 2nd moment superoperator of the distribution

over $\{W(\theta)\}_{\theta \in \Theta}$ as

$$M_{W(\Theta)}^{(2)} = \int_{\text{Unif}(\Theta)} dW(\theta) W(\theta)^{\otimes 2} \otimes (W(\theta)^*)^{\otimes 2}, \quad (7)$$

the ε deviation from a 2-design is given by some norm² of the superoperator

$$\mathcal{A}_{W(\Theta)}^{(2)} = M_{W_H}^{(2)} - M_{W(\Theta)}^{(2)}, \quad (8)$$

where $M_{W_H}^{(2)}$ and $M_{W(\Theta)}^{(2)}$ are the second moments of the Haar distribution and the parametrization induced unitary distribution respectively. An exact 2-design is then achieved when $\|\mathcal{A}_{W(\Theta)}^{(2)}\| = 0$ and the family of parameterized unitaries is considered to be *maximally* expressive. However, such an assumption on the trainable unitaries leads to the model *exponential concentration* phenomena where the variance of the model vanishes exponentially in the number of qubits [4, 7, 23]. For completeness, we formally define the model exponential concentration as follows.

Definition 2 (Model Exponential Concentration). *Consider a quantum model $f(x, \theta)$ of the form in Eq. (1) and assume that the trainable parameters are sampled uniformly over Θ . The model is said to exhibit an exponential concentration phenomenon for the parameter hyperspace Θ when*

$$\text{Var}_{\theta \sim \Theta}[f(x, \theta)] = \mathcal{O}\left(\frac{1}{b^n}\right), \forall x \in \mathcal{X} \quad (9)$$

for some constant $b > 1$.

The above definition of model concentration is defined independently of its Fourier decomposition. Nevertheless, we can obtain

²Different norms can be used to characterize the ε -distance to a 2-design. They are indeed equivalent up to exponential factors in the number of qubits [25].

a similar phenomenon at the level of the model’s Fourier coefficients. Specifically, one can study how the variance of each Fourier coefficient (with respect to the uniform distribution over the trainable parameter hyperspace) decays in the number of qubits. Notably, we say that a QFM suffers from *vanishing expressivity* if some or all of its Fourier coefficients are exponentially concentrated around their mean.

Definition 3 (Vanishing Expressivity). *Consider a quantum Fourier model of the form in Eq. (1) with the Fourier decomposition in Eq. (3) and a spectrum Ω . Let us also assume that the trainable parameters are sampled uniformly. The Fourier model is said to suffer from vanishing expressivity when some Fourier coefficients have an exponentially vanishing variance, i.e.*

$$\exists \Omega_{\text{vanish}} \subset \Omega \mid \forall \omega \in \Omega_{\text{vanish}} \quad \text{Var}_{\theta}[c_{\omega}(\theta)] = \mathcal{O}\left(\frac{1}{b^n}\right) \quad (10)$$

for some constant $b > 1$.

Although the above definition of vanishing expressivity resembles the model concentration phenomenon introduced in Definition 2, they do indeed differ on an operational level. Specifically, the model exponential concentration implies that an exponential number of shots is required to resolve exponentially small changes in the model for a fixed input $x \in \mathcal{X}$. This makes optimization inefficient and non-scalable when considering random initialization [26]. However, the *vanishing expressivity* phenomenon introduced in Definition 3 does not directly impact the model’s trainability guarantees. This is because, in practice, we do not measure individual Fourier coefficients or use them to compute gradients. Nevertheless, one may still wonder if the trained model will be able to reach an initially exponentially vanishing Fourier coefficient. If such a correlation exists between initially vanishing Fourier coefficients and their reachability in the final trained model, it does justify the *vanishing expressivity* connotation. While we do not provide analytical evidence of this correlation, we do support our intuition with numerical simulations, showing that vanishing coefficients are hard to train in Section 4.2.

3 Main Results

In this section, we present our main theorems and corollaries on expressivity constraints in quantum Fourier models. Specifically, we study the concentration of Fourier coefficients by computing their variance under different assumptions about the trainable unitaries distribution. We then show that the variance is always constrained by the frequency redundancy and that some Fourier coefficients may exhibit an exponential concentration phenomenon, leading to *vanishing expressivity*.

To do so, we start by considering the global 2-design hypothesis on the trainable unitaries and provide an exact expression of the Fourier coefficients variance in Theorem 1 for a single-layer model and in Theorem 6 for a reuploading model. Secondly, we relax the global 2-design assumption and provide an upper bound on Fourier coefficients variance under the ε -approximate 2-design hypothesis in Theorem 7. Finally, we consider the brick-wise circuit architecture with local 2-design blocks and give an upper bound on the variance in this setting. This circuit architecture falls within the ε -approximate 2-design assumption, but with more structure that allows to take into account the locality of the observable.

We note that in the remainder of the manuscript, we focus on one-dimensional input vectors ($D = 1$). However, our results can be easily extended to the high-dimensional setting if we assume that the encoding unitaries within a single layer commute, as detailed in Appendix A.1.1. Beyond this assumption, the analysis becomes more intricate and we do not cover it in this work.

3.1 Trainable layers as global 2-design

As described in Section 2.2, in the unitary space, the expressivity of trainable unitaries is often characterized by how uniformly they explore the unitary group and a parametrized unitary is said to be maximally expressive if its distribution approximates the Haar measure. However, it has been shown in [4, 7], that the quantum model and its gradient exhibit an exponential concentration phenomena under the 2-design assumption, resulting in an unexpressive model in practice.

Here, we explore the implications of considering maximally expressive trainable unitaries (i.e. each of the trainable layers forms an exact 2-design) on the Fourier coefficients variance.

First, we present an exact expression of the Fourier coefficients variance for a QFM with a single circuit layer ($L = 1$) in Theorem 1. We then extend the result to a reuploading model with $L \geq 1$ in Theorem 6, Appendix B.4.

Theorem 1 (Fourier coefficients variance with 2-design trainable unitaries, Informal). *Consider a quantum model of the form in Eq.(1) and a parametrized circuit of the form in Eq.(2) with $L = 1$ layers and fixed encoding Hamiltonians resulting in a spectrum Ω . We assume that each of the trainable layers $W^l(\theta)$, $l \in \{1, 2\}$ form independently a 2-design. The expectation and variance of each Fourier coefficient $c_\omega(\theta)$ for the frequencies $\omega \in \Omega$ appearing in the model Fourier decomposition in Eq.(3) are given by*

$$\begin{aligned} \mathbb{E}_\theta[c_\omega(\theta)] &= \frac{\text{Tr}(O)}{d} \delta_\omega^0, \\ \text{Var}_\theta[c_\omega(\theta)] &\in \Theta \left(\alpha \frac{|\tilde{R}(\omega)|}{d} - \frac{\alpha}{d^2} \delta_\omega^0 \right). \end{aligned} \quad (11)$$

Here we recall that $d = 2^n$ and introduce the normalized frequency redundancy $|\tilde{R}(\omega)| := |R(\omega)|/d^2$ and the constant $\alpha := (d\|O\|_2^2 - \text{Tr}(O)^2)/d^2$ which depends on the observable O .

Theorem 1 establishes that, under the 2-design assumption on the model's trainable unitaries, the variance of a Fourier coefficient depends linearly on its (normalized) frequency redundancy up to some prefactor. Here we recall that the normalized frequencies sum up to one. Thus, Theorem 1 implies that while frequencies with high redundancies exhibit a relatively large variance, the ones with low redundancies are way more constrained. In other words, the distribution of the Fourier coefficients is dictated by the redundancies and hence by the encoding Hamiltonians. In Section 5, we illustrate this dependence through the numerical study of two models with different encoding Hamiltonians corresponding to spiked and flat frequency distributions (i.e. defined by normalized redundancies over the model's spectrum). Moreover, we generalize the result of Theorem 1 to the setting of reuploading models

with $L \geq 1$ alternating layers in Theorem 6. Similarly, we prove that a Fourier coefficient variance is linear in its frequency redundancy. The formal version of Theorem 1, Theorem 6 and their proof are presented in Appendix B.1.

Apart from the linear dependence of the Fourier coefficient variance on their respective redundancies, one can also deduce the decay rate of the Fourier coefficients variance from their exact expressions given in Theorem 1 and Theorem 6. This is captured via Corollary 1.

Corollary 1 (Fourier coefficients variance decay with 2-design trainable unitaries, Informal). *Consider a quantum model of the form in Eq.(1) and a parametrized circuit of the form in Eq.(2) with $L \geq 1$ layers and fixed encoding Hamiltonians resulting in a spectrum Ω . We assume that each of the trainable layers $W^l(\theta)$ form independently a 2-design. The variance of each Fourier coefficient $c_\omega(\theta)$ for the frequencies $\omega \in \Omega$ appearing in the model Fourier decomposition in Eq.(3) is upper bounded by*

$$\text{Var}_\theta[c_\omega(\theta)] \in \mathcal{O} \left(\alpha \frac{|\tilde{R}(\omega)|}{d} \right), \quad (12)$$

where we recall that $d = 2^n$ and α is a constant given by $\alpha := (d\|O\|_2^2 - \text{Tr}(O)^2)/d^2$.

From Corollary 1, one can see that, under reasonable assumptions on the observable norm, all Fourier coefficients variance decay exponentially in the number of qubits. Specifically, the prefactor α can be bounded by a constant for any observable satisfying $\|O\|_2^2 \in \mathcal{O}(d)$. Additionally, the normalized redundancy $|\tilde{R}(\omega)|$ is, by definition, bounded by one. This implies that irrespective of the frequency redundancy and thus of the encoding strategy, all coefficients concentrate exponentially toward their mean value. This result can be viewed as an exponential concentration statement of each Fourier coefficient in a reuploading model, aligning with results in [4] about the exponential concentration of the model under the 2-design assumption.

While it is useful to show the connection between the Fourier coefficients and the spectrum redundancies, considering global 2-design is a strong assumption leading to the model's exponential concentration. In practice, it is improbable that one will use trainable layers

forming 2-design for learning purposes. Thus we propose in the following to relax this hypothesis by first considering trainable layers as forming approximate 2-design and then as made of local 2-design blocks and that for models with a single layer ($L = 1$).

3.2 Trainable layers as global ε -approximate 2-design

Let us now consider the broader setting where the trainable unitaries $W^l(\theta)$ form each an ε -approximate 2-design. By moving away from highly expressive trainable unitaries, one might question if it is feasible to break free from the constraining redundancy dependence of the Fourier coefficients variance established in Theorem 1, or if such dependency is an inductive bias of the quantum model which still hold even when the trainable unitaries are not maximally expressive (i.e. do not form exact 2-designs).

To do so, we further build on results from [23] about the model concentration for approximate 2-design unitaries and explore encoding dependent concentration for single Fourier components, giving a finer interpretation of the model's expressivity through the Fourier lens.

In the following theorem, we provide an upper bound on the Fourier coefficients variance for a single-layer circuit formed by arbitrary trainable layers.

Theorem 2 (Fourier coefficients variance decay with approximate 2-design trainable unitaries, Informal). *Consider a quantum model of the form in Eq.(1) and a parametrized circuit of the form in Eq.(2) with $L = 1$ layers and fixed encoding Hamiltonians resulting in a spectrum Ω . We assume that each of the trainable layer $W^l(\theta)$, $l \in \{1, 2\}$ form independently an ε -approximate 2-design. The variance of each Fourier coefficient $c_\omega(\theta)$ for the frequencies $\omega \in \Omega$ appearing in the model Fourier decomposition in Eq.(3) is upper bounded as*

$$\text{Var}[c_\omega] \in \mathcal{O}(Q_\varepsilon(|\tilde{R}(\omega)|)) , \quad (13)$$

where Q_ε is a polynomial of degree at most 2 in the normalized frequency redundancy $|\tilde{R}(\omega)|$ defined for different ε measures as

$$Q_{\varepsilon_\diamond} = \|O\|_2^2 |\tilde{R}(\omega)|_{\varepsilon_\diamond} + \|O\|_1^2 \varepsilon_\diamond^2 , \quad (14)$$

$$Q_{\varepsilon_\infty} = \frac{\|O\|_2^2}{d} \sqrt{|\tilde{R}(\omega)|_{\varepsilon_\infty}} + d^2 \|O\|_2^2 |\tilde{R}(\omega)|_{\varepsilon_\infty}^2 , \quad (15)$$

$$Q_{\varepsilon_M} = \|O\|_2^2 |\tilde{R}(\omega)|_{\varepsilon_M} + d^2 \|O\|_2^2 |\tilde{R}(\omega)|_{\varepsilon_M}^2 . \quad (16)$$

Here, we use the shorthand $\varepsilon_\diamond := \|\mathcal{A}^{(2)}\|_\diamond$ for the diamond norm defined in Eq.(B38), $\varepsilon_\infty := \|\mathcal{A}^{(2)}\|_\infty$ for the spectral norm and $\varepsilon_M := d^2 \max_{i,j} |\mathcal{A}^{(2)}|_{i,j}$. We also recall that $\mathcal{A}^{(2)}$ is a superoperator defined in Eq.(8).

Theorem 7 shows that the variance of a Fourier coefficient in the approximate 2-design setting is constrained by the combined action of the normalized frequency redundancy $|\tilde{R}(\omega)|$ and the ε -distance of the trainable unitaries to a 2-design. Specifically, for a fixed choice of the trainable unitaries distribution and thus for a fixed ε value, the degree to which each Fourier coefficient concentrates around its mean is constrained by its corresponding normalized frequency redundancy. Therefore, we prove that the vanishing expressivity phenomenon, whereby some Fourier coefficients exhibit exponentially decaying variance, may still hold beyond the 2-design assumption. We note that the bounds in Eqs.(14)-(16) correspond to different norms used to quantify the distance from a 2-design. These ε -distance definitions are equivalent up to some prefactors [25] and we include all of them because one bound may be tighter than the other depending on the interplay between the observable norm³, the frequency redundancy and the ε -distance scalings. Precisely, while ε_∞ saturates at 1 and $\varepsilon_\diamond$ at 2, the monomial-based epsilon ε_M can take up values up to d^2 . Additionally, we recall that the normalized frequencies $|\tilde{R}(\omega)|$ take up values within $[1/d^2, 1]$. Hence, the bounds can be used to prove the vanishing expressivity phenomenon introduced in Definition 3 for frequencies with relatively low redundancies. Namely, a frequency with a normalized redundancy that counterbalance the observable norm will exhibit an exponential decay

³We consider Schatten p -norms defined as $\|O\|_p := (\text{Tr}[(\sqrt{O^\dagger O})^p])^{1/p}$.

on average over trainable unitaries forming an approximate to 2-design. In section 4, we further comment on the scalings of the upper bound in Eq.(16) and its dependence on the observable and encoding strategy. We provide the proof of Theorem 2 and its formal version in Appendix B.2 for different norms and numerically evaluate the Monomial bound in Eq.(16) for different encoding strategies with different spectrum distributions in Section 5.3.

Here we note that the upper bounds in Theorem 2 are looser for local observables compared to global ones, in general. This observation does not come as a surprise since the obtained bound as a function of the *global* ε expressivity measure of the circuit does not capture the observable-circuit interaction in finer detail. Specifically, the interaction between an m -local observable and the remainder of the circuit is captured by the backward light cone of the observable, i.e., the sub-circuit containing all blocks with at least one qubit causally connected to the local observable input qubits. In the next section, we explore the variance of the Fourier coefficients by taking into account the observable locality.

3.3 Trainable layers as local 2-design blocks

In this section, we consider a brickwise circuit architecture formed by trainable local 2-design blocks and local encoding blocks, which has been previously studied in [27]. As depicted in Fig.4, the n -qubit circuit is made of layers of m -qubit unitaries (trainable and encoding unitaries) acting on alternating groups of m neighboring qubits. We consider that each of the trainable blocks forms an exact 2-design on the m -qubit subsystem on which it acts non-trivially. This setting is a special case of the *global* ε -approximate 2-design [28]. However, it will give us more accurate results when one considers local observables acting non-trivially on an m -qubit subsystem s_k of the form $O = \hat{O}_{s_k} \otimes \mathbb{1}_{\overline{s_k}}$. Indeed, by considering this circuit architecture, the backward *causal light cone* of such local observables is well defined and one can easily notice that the effective model spectrum will be reduced. In this setting, we derive the expression of the Fourier coefficients variance in Appendix B.4.3. However, since the obtained expression is quite cumbersome, we present in the

following theorem an upper bound on the Fourier coefficients variance for two different assumptions on the local observable.

Theorem 3 (Fourier coefficients variance decay with brickwise local 2-design circuit). *Consider a quantum model of the form in Eq.(1) and a parametrized circuit of the form in Eq.(2) using a brickwise architecture with $L = 1$ layers and observable $O = \hat{O}_{s_k} \otimes \mathbb{1}_{\overline{s_k}}$ acting non trivially on the m -qubit subsystem s_k . Assume that each trainable m -qubit unitary forms a local 2-design. The variance of each Fourier coefficient $c_\omega(\theta)$ for the frequencies $\omega \in \Omega$ appearing in the model Fourier decomposition in Eq.(3) is upper bounded as*

1. *If $\|\hat{O}_{s_k}\|_2^2 \leq 2^m$, we have*

$$\text{Var}[c_\omega] \leq \left(\frac{2^{m+1}}{2^{2m} - 1} \right)^{2L_2} |R_{E_k}(\omega)|^2. \quad (17)$$

2. *If $\hat{O}_{s_k}$ is a projector of rank r , we have*

$$\text{Var}[c_\omega] \leq \left(\frac{2^{m+1}}{2^{2m} - 1} \right)^{2L_2} \left(\frac{r}{2^m} \right)^2 |R_{E_k}(\omega)|^2. \quad (18)$$

Here $R_{E_k}(\omega)$ is the frequency generator obtained from the encoding blocks inside the observable backward light cone $\mathcal{L}_k$ (acting non trivially on $\mathcal{S}_{E_k}$) and L_2 is the depth of the post-encoding parameterized block.

Theorem 3 gives us an upper bound on the variance of the Fourier coefficients while considering circuits made of local 2-design blocks. Once again, we observe that this quantity is constrained by the frequency redundancy. In addition, this result indicates that the vanishing Fourier coefficient phenomenon could depend on the circuit depth while considering a local observable. Moreover, the bound in Theorem 3 implies that a frequency with a relatively low frequency (i.e. $|R_{E_k}(\omega)| = \mathcal{O}(1)$) will suffer from exponentially vanishing variance for a depth L_2 linear in n . The proof of this Theorem is given in Appendix B.4.

4 Discussing the quantum Fourier model constraints

In this work, we established a connection between the spectrum redundancies and the statistical




Figure 4: **Brickwise circuit architecture made of local blocks acting on neighboring qubits.** As shown, L_1 is the depth of the pre-encoding trainable block and L_2 is the depth of the post-encoding one. We consider an m -local observable acting non-trivially on subsystem s_k and we denote its backward light cone by the subcircuit $\mathcal{L}_k$. We denote by $S_{\mathcal{L}_k}$ the subsystem on which the backward light cone acts non-trivially and S_{E_k} the subspace on which the encoding layer (made of the green blocks) acts non trivially inside the light cone. we also define $S_{\overline{E_k}}$ as the compliment of S_{E_k} in $S_{\mathcal{L}_k}$.

behaviour of Fourier coefficients for arbitrary trainable unitaries, *on average*. Namely, we showed an *inductive bias* of the Fourier model where the variance of a Fourier coefficient is upper bounded by a polynome in its redundancy. we further introduced the concept of *vanishing expressivity* where the variance of some Fourier coefficient is exponentially vanishing in the number of qubits. In this section, we further discuss these phenomena and study their implications on model design guidelines in 4.1. In addition, we provide a generic bound on the 2-norm of the Fourier coefficients vector and briefly discuss controllability-related constraints on the Fourier coefficients in 4.3. Finally, we discuss the limitations of the framework and the assumptions we considered and thus the limitations of the obtained results in 4.4.

4.1 Vanishing Fourier coefficients and vanishing model

In this section, we discuss the *vanishing expressivity* phenomenon, whereby Fourier coefficients variance decay exponentially in the system size. Specifically, we further discuss the scaling of the upper bounds established in Theorem 2 with respect to the different quantities of interest. Moreover, we relate this analysis of the Fourier coefficients decay to the

analysis of the full model decay. Indeed, we previously stressed in Section 2 that the vanishing expressivity phenomenon, whereby some Fourier coefficients exhibit exponentially decaying variance on average, is conceptually different from the model's exponential concentration introduced in Definition 2. Nevertheless, one may still wonder if these two phenomena are equivalent or if one implies the other.

Under the 2-design assumption on the trainable unitaries, we showed in Theorem 1 that the model's exponential concentration goes in hand with the exponential decay of all Fourier coefficients independently of the encoding strategy as detailed in Corollary 1. However, when using approximate 2-design trainable layers, the link is not trivial. To better understand the relation between the model and Fourier coefficients exponential decay beyond the 2-design assumption, we provide in the following Corollary an upper bound on the model's variance when using approximate 2-design trainable layers. Ultimately, we are interested in identifying regimes where the model's variance is not exponentially vanishing whereas all or some of its Fourier coefficients suffer from exponential concentration.

Corollary 2. Consider a quantum model $f(x, \theta)$ of the form in Eq.(1) and a parametrized circuit

of the form in Eq.(2) with $L = 1$ layers and fixed encoding Hamiltonians. We assume that each of the trainable layer $W^l(\theta)$, $l \in \{1, 2\}$ form independently an ε_M -approximate 2-design according to the monomial definition introduced in Definition 5. For a fixed $x \in \mathcal{X}$, the variance of the model $f(x, \theta)$ is upper bounded as

$$\text{Var}_\theta[f(x, \theta)] \in \mathcal{O}\left(\|O\|_2^2 \varepsilon_M\right). \quad (19)$$

Corollary 2 establishes an upper bound on the full model variance similar to the one given in Theorem 7 for each Fourier coefficient. We note that this results have been established in previous works [6, 26, 23] but we adapt it here to the monomial distance ε_M to a 2-design defined in Eq.(B30). The proof of this Corollary is detailed in Appendix B.3.

By combining Theorem 2 and Corollary 2, we can capture scenarios where relatively low redundant frequencies are vanishing whereas there is a leeway for the global model to not be.

Specifically, if we consider frequencies with redundancies scaling at most polynomially in system size, i.e. $|R(\omega)| \in \mathcal{O}(\text{poly}(n))$ (or equivalently $|\tilde{R}(\omega)| \in \mathcal{O}\left(\frac{\text{poly}(n)}{d^2}\right)$), then the corresponding coefficient variance upper bound in Eq.(16) scales as

$$\text{Var}[c_\omega] \in \mathcal{O}\left(\text{poly}(n) \frac{\|O\|_2^2 \varepsilon_M}{d^2} (1 + \varepsilon_M)\right). \quad (20)$$

This implies that as long as $\|O\|_2^2 \varepsilon_M \in \mathcal{O}(\text{poly}(n))$, the Fourier coefficient with redundancies $|R(\omega)| \in \mathcal{O}(\text{poly}(n))$ suffer from exponentially decaying variance. On the other hand, Corollary 2 provides guarantees of the exponential decay of the full model only when $\|O\|_2^2 \varepsilon_M \in \mathcal{O}(1/d)$. Consequently, we see that there may be a leeway for the global model to be non vanishing while Fourier coefficients with polynomially big redundancies suffer from exponential concentration for a reasonably wide ε_M range as depicted in Fig.5.

A straightforward construction of an *expressive* model is to consider an encoding strategy where the size of the spectrum is exponential in the number of qubits and hence less prone to classical dequantization [10, 13]. This implies that the spectrum is weakly degenerate with many frequencies ω such that $|R(\omega)| = \Theta(1)$. This is indeed the case for the exponential encoding and

the Golomb encoding (See Appendix A.2). In the latter setup, we have that $|R(\omega)| = 1 \quad \forall \omega \in \Omega^*$.

Consequently, although the quantum model has theoretically access to an exponential number of frequencies, the contribution of each frequency is vanishing.

When considering general encoding strategies, Theorem 7 implies that frequencies with low redundancies are more likely to suffer from exponential concentration, limiting the expressivity of the quantum model. Specifically, for fixed trainable unitaries and thus fixed ε_M , the upper bound on the Fourier coefficient variance allows high redundant frequencies to possibly escape exponential concentration while the low redundant ones will exhibit vanishing variance, leading to the *vanishing expressivity* phenomenon.

4.2 Impact of the vanishing expressivity phenomena on training and dequantization

In the previous section, we compared the exponential concentration of the Fourier coefficients and the exponential concentration of the full model and focused on the settings where these two phenomena can happen independently. Indeed, we should further emphasize that the interpretation of these two behaviors are fundamentally different. Specifically, since the Fourier coefficients are not directly measured to evaluate the model gradient's, the statement that they initially suffer from exponential concentration cannot be directly related to a resource problem (i.e. finite number of shots) as in other exponential concentration analyses of the whole model. This means that the vanishing Fourier coefficients phenomenon cannot be directly related to trainability issues. This observation justifies why we opted to denote this behavior by *vanishing expressivity* with the intuition that frequencies with vanishing coefficients will have a negligible contribution to the quantum model.

However, one can plausibly think of a scenario where the signal from each Fourier coefficient is exponentially small. However, it can give rise to a significant signal when merged together. Hence, the exponential concentration of Fourier coefficients does not necessarily imply a constraint on the expressivity of the quantum Fourier model in this case.




Figure 5: **Illustration of the vanishing model and vanishing Fourier coefficient phenomena according to the ε_M distance to a 2-design.** Considering coefficients c_ω with redundancies $|R(\omega)| = \mathcal{O}(\text{poly}(n))$, one can notice that the vanishing expressivity phenomenon can happen outside of the regime with guaranteed exponential concentration of the full model (yellow part). Specifically, the dashed gray part corresponds to ε_M range where the quantum model f is proven to be vanishing while the black one corresponds to the vanishing Fourier coefficients regime. The yellow part indicates the regime where the Fourier coefficients are vanishing but not necessarily the case for the corresponding model. Finally, the behavior of the model and its coefficients is unknown in the blue part.

Moreover, the analysis of the vanishing expressivity phenomenon holds on average when the parameterized unitaries form approximate 2 design but it gives no guarantees on the model’s effective expressivity during the training stage. Consequently, it is possible to start with exponentially small contributions of the Fourier coefficients then reach all of the theoretically accessible frequencies, given that we can efficiently train the quantum model.

Impact of vanishing expressivity on training. In order to have a better understanding of the consequences of vanishing Fourier coefficients, we go beyond our analytical results and perform numerical simulations to study the impact of the vanishing expressivity phenomenon on the final trained model. Specifically, we consider the task of training a fixed quantum model (fixed trainable unitaries and encoding unitaries) to fit two different sinusoidal functions with two different target frequencies, one being highly redundant in the quantum Fourier model spectrum and the other having a relatively low redundancy as depicted in Fig.6. Then the training results are presented, showing that the model manages to reach the high redundant frequency but not the low redundant one. This result supports indeed the intuition that frequencies with initially vanishing Fourier coefficients are harder to reach. We note however that this behavior can also be due to controllability issues in which there is no configuration of parameters such that the low redundant frequency does actually have a non zero weight. In the next section, we discuss the controllability issue in more details.

Impact of vanishing expressivity on RFF⁴ based

⁴Random Fourier Features.

dequantization. For RFF-based dequantization schemes [13] of QFMs, finding the optimal frequency distribution to build the classical surrogate and hence to dequantize the quantum model, requires knowledge about the spectral properties of the final model. With our results, we can propose the frequency distribution given by the redundancies as described in section 2 as a natural distribution that encodes the bias in the quantum model. However, this choice is based on the assumption that the final model’s spectral properties will inherit the spectral properties of the average case model (initial model with random parameter initialization). Although, this assumption is not guaranteed to be fulfilled in general, the numerics in Fig.6 show that the decaying Fourier coefficients in “average” models persists in the final trained model.

4.3 Fourier Norm Bound and Controllability constraints

In this work, we focused on studying variances by considering a uniform distribution over the parameter vector θ . In this Section, we point out that additional constraints can occur due to the lack of model *controllability*, defined as the number of Fourier coefficient that one can independently control by tuning the trainable parameter vector.

First, we establish a generic constraint on the quantum model Fourier coefficients that holds for any Hamiltonian encoding scheme and that is independent of the trainable unitaries distribution in the following theorem.

Theorem 4 (Fourier Norm Bound). *Consider a quantum model $f(x, \theta)$ of the form in Eq.(1) using an observable O and a parametrized circuit*




Figure 6: **Impact of the vanishing expressivity phenomenon on a trained QFM.** We study the expressive capability of a trained QFM acting on $n = 12$ qubits using the exponential encoding strategy with a) a spectrum Ω and frequency redundancies $|R(\omega)|$. The model is trained to fit two target frequencies with different redundancies in the model's spectrum: Target 1 corresponds to a high redundant frequency (in red), and Target 2 corresponds to a low redundant frequency (in purple). b) We plot the loss and the Fourier coefficients norm evolution fitting Target 1. In c), we plot the loss and the Fourier coefficients norm evolution fitting Target 2. The model succeeds in fitting the high redundant frequency but not the low redundant one.

of the form in Eq.(2) with $L \geq 1$ layers. Also assume that the encoding Hamiltonians are fixed, giving rise to a spectrum Ω . Then,

$$\forall x \in \mathbb{R}^d, \forall \theta \in \Theta, |f(x, \theta)|^2 \leq \|O\|_\infty^2 \quad (21)$$

$$\forall \theta \in \Theta, \sum_{\omega \in \Omega} |c_\omega(\theta)|^2 \leq \|O\|_\infty^2 \quad (22)$$

The first part of Theorem 4 is a trivial constraint that holds for any quantum model of the form in Eq.(1) even outside of the Fourier framework. This constraint has been mentioned in [13] to highlight the fact the a quantum Fourier model can not achieve any linear function in the Fourier basis given by its spectrum. The second part of the theorem is more subtle. While very similar to the Parseval identity, the bound in Eq.(22) holds for any real valued spectrum Ω . It shows that the 2-norm of the Fourier coefficient vector is upper bounded by the observable largest eigenvalue, introducing another generic constraint on the quantum model Fourier coefficients. We provide the proof of Theorem 4 in Appendix B.5.

In addition to the previous results, it is important to stress that limitations in the controllability of the trainable unitaries can affect the controllability of the Fourier Model. By considering the expanded expression of the quantum Fourier model given in Eq.(4), each

Fourier coefficient is defined as a sum and product of coefficients from the trainable unitary matrices.

According to the number of parameters, and the set of gates chosen, the number of independent matrix coefficients that one can freely control through the trainable parameters can fluctuate from one circuit to another. Previous works [6, 5] have studied the controllability of PQC, by analyzing the corresponding Dynamical Lie Algebra (DLA), or by considering the quantum Fisher Information matrix which characterizes the state controllability. In [29, 30], the authors highlight a connection between the maximal controllability of a PQC, i.e., the dimension of its DLA, and its capacity to be trained. Hence, this controllability notion will be key in characterizing the controllability of Fourier coefficients.

Namely, by looking more closely at the Fourier coefficient expression given in Eq.(4), one can simply observe that a pair of paths $(J, J') \in R(\omega)$ from the frequency generator defined in Definition 1 allocates coefficients of the trainable unitary matrices to its frequency. In addition, some unitary coefficients are shared among different Fourier coefficients as a consequence of some branches in the generating tree (see Fig.2) being shared

between different frequencies. Consequently, this can potentially create correlations between the Fourier coefficients. Therefore, if the trainable layers have a low controllability, it could lead to the impossibility of controlling independently a large number of Fourier coefficients. This is particularly important due to the fact that increasing the number of parameters seems to increase the control and decrease the distance to a 2-design (see for example the evolution of the distance to a 2-design in the case of Periodic Ansatz through the Theorem 1 in [6]).

In Fig.6, we offer some plots on the evolution of the Fourier coefficients when training a QFM to learn two sinusoidal functions. The first one has a target frequency with high redundancy in the quantum Fourier model, and the second one corresponds to a low redundant frequency. One can notice that, for the second target, the PQC takes more epochs to converge and fails to minimize the loss. In addition, one can notice that, during the training, the Fourier coefficients surrounding the target frequency are changing a lot, due to the lack of controllability of the Fourier coefficient.

4.4 Limitations of the Framework

In this Section, we discuss the limitations of our framework and the used assumptions to derive our main results presented in Section 3.

First, the statistical analysis of the Fourier coefficients established in this work holds when the trainable parameters are sampled uniformly and independently. Although we considered the case where the trainable unitaries form an approximate 2-design, the obtained constraints on the Fourier coefficients variance and their decay only hold on average. Consequently, extrapolating this average case behavior to the final trained model is not systematic. Indeed, this gap between average case and final model guarantees constrains the direct applicability of our results to rigorously study the efficiency of random Fourier Features based dequantization schemes [13]. While proving analytically the impact of the frequencies redundancies on the *effective expressivity* of the final trained model is a hard task, we provide in section 4.2 numerical evidence of this behavior. Moreover, the upper bounds presented in Theorem 2 are useful to establish generic theoretical guarantees on the

model's expressivity. However, they only apply to quantum models with a single uploading layer. In addition, estimating the ε -distance of the trainable unitaries to 2-designs and its scaling is not efficient in practice.

Finally, we would like to point out that our results could be extended to the case of subspace preserving quantum circuits. In this type of PQCs, one can restrict the computation to a particular subspace by using input states which lie in the subspace, reducing the dimension of the effective Hilbert space. These methods allow to avoid Barren Plateaus while considering subspaces of polynomial size [6, 29, 31, 30] but question the quantum advantage of such models [32, 33]. Considering subspace preserving unitaries, our results can easily be adapted. The dependency over the frequency distribution will still hold, but the value of d (the dimension of the Hilbert space) will be substituted by the dimension of the subspace. Therefore, models generated by subspace preserving circuits could exhibit similar *inductive bias* arising from the redundancy constraint on the variance of its Fourier coefficients.

5 Numerical Results

In this section, we simulate the distribution of the Fourier coefficients for several types of circuits and encoding strategies allowing us to compare the numerical simulations with our theoretical results presented in Section 3.

5.1 Methodology

A description of the different types of parameterized circuit architectures used can be found in Appendix D. Each circuit consists of a choice of encoding layers (encoding Hamiltonians), alternated with trainable layers, L times. Each trainable layer is a repetition of the same ansatz block (See Appendix D). Finally, these blocks can act globally on all qubits, or locally on m qubits, with m being constant.

The statistical properties of the PQC are obtained with the following method: for each ansatz (parameterized circuit's architecture), we pick at random a large number of parameter vectors θ s. For each vector, we evaluate the model $f(x, \theta)$ as many points $x \in \mathbb{R}$ as the Shannon




Figure 7: For $n = 5$ qubits, one circuit layer $L = 1$, five repetitions of the strongly entangling ansatz per trainable layer (see Appendix D) and global observable O_G ; relation between the variance of each Fourier coefficient $\text{Var}[c_\omega]$ and its redundancy $|R(\omega)|$. Values shown for two different encoding strategies, a) Pauli encoding and b) exponential encoding.




Figure 8: Variance of the Fourier coefficients against the number of qubits for a single circuit layer ($L = 1$). Five repetitions of the strongly entangling ansatz per trainable layer are used ensuring full connectivity. a) Case of local observable. b) Case of global observable. Each color corresponds to a certain frequency ω , squares of the same color correspond to the variance theory for 2-design from Theorem 1 and the dots correspond to the simulated variance.

criterion establishes. We obtain the model's Fourier coefficients $\{c_\omega(\theta)\}_{\omega \in \Omega}$ by applying the discrete Fourier transform to the measurement output. We then aggregate the results for all θ s and compute the variance $\text{Var}(c_\omega)$ for each $\omega \in \Omega$.

As we are performing simulations, we are computing the wave function analytically and hence have access to the direct probabilities of each state. Therefore, the shot noise from measurement is not taken into account.

We use three distinct approaches for Hamiltonian encoding. The first is *Pauli encoding*, which is implemented here by applying the same single-qubit Pauli rotation gate with $H = \sigma_x$ applied on each available qubit, as seen in the inset within Fig.7a. Secondly, we use *exponential encoding* which, as described in Appendix A.2, generates an exponentially large spectrum Ω by also using single qubit rotations but with scaling factors introduced in [16] (see Fig.7b). Finally, we also consider

the *Golomb encoding* in the single circuit layer setting ($L = 1$) which is obtained by using a global Hamiltonian encoding whose diagonal elements are those of a perfect Golomb ruler [15].

Furthermore, two types of observables are considered. A global observable where the measurement is acting non-trivially and simultaneously on all of the qubits: $O_G = |0\rangle\langle 0|$ and a local observable where an average is taken over single qubit measurements $O_L = \frac{1}{n} \sum_j |0\rangle\langle 0|_j \otimes \mathbb{1}_{\bar{j}}$.

5.2 Trainable layers as global 2-design

We consider trainable layers that form an exact 2-design, by using five repetitions of the *Strongly Entangling* elementary ansatz described in Fig.17.

For a single-layer circuit ($L = 1$), Theorem 1 establishes a linear relationship between the variance of c_ω and the redundancy of the




Figure 9: Comparison between simulated (in green) and two theoretical values for the variance of c_ω : global 2-design setting (in blue) and ε -approximate 2 design setting (in pink). Parameterized layers are made of local blocks V made of 5 repetitions of the *Strongly Entangling* ansatz acting each on a) $m = 1$ qubits b) $m = 2$ qubits c) $m = 4$ qubits, with a Pauli encoding layer, and the global observable O_G .




Figure 10: Same setting as in Fig.9b but with a) Exponential Encoding b) Golomb encoding.

corresponding frequency. For $L > 1$, the relationship in Theorem 6 also involves the partial redundancies (See Definition 4).

In Fig.7, we show that the numerics match Theorem 1 for the Pauli and the exponential encodings respectively. Precisely, we see that the values of $\text{Var}[c_\omega]$ and $|R(\omega)|$ coincide, after a linear rescaling. Fig.7b shows that the exponential encoding strategy has a broader spectrum Ω than the Pauli encoding strategy in Fig.7a. In addition, both encoding methods display their predicted frequency distributions defined by $|R(\omega)|$ (See appendix A.2 for details).

In Fig.8, we simulate the same circuit, focusing on the Pauli encoding strategy, and observe the scaling of the Fourier coefficients variance with the number of qubits. In addition, we distinguish two cases: local observable in Fig.8a and global

observable in Fig.8b.

In both cases, we observe that the simulated variance and the one predicted by Theorem 1 match. While it is less obvious why we have this behavior for the local observable in Fig.8a, it is actually due to the fact that with the strongly entangling circuit, the local observable backward light cone covers all of the circuit. However, the difference between considering a local observable and a global one can be seen through the steepness of the slopes of each frequency variance as the number of qubits increases and through the number of qubits from which we start observing a strictly decreasing variance. Indeed, as predicted by Theorem 1, we observe that the Fourier coefficients variance vanish exponentially. Nonetheless, the vanishing phenomena for the high frequencies is not captured by the plots since




Figure 11: Comparison between simulated (in green) and two theoretical values for the variance of c_ω : global 2-design setting (in blue) and ε -approximate 2 design setting (in pink). Parameterized layers are made of 2-local blocks V made of 5 repetitions of the *Strongly Entangling* ansatz with the local observable O_L and a) Exponential encoding strategy b) Golomb encoding strategy.




Figure 12: Variance of the Fourier coefficients against the number of qubits for a single circuit layer ($L = 1$) made with the brickwise structure (Fig.4) by using repetitions of the Simplified Two Design ansatz (See Fig.18), Pauli encoding and $n = 2, 3, \dots, 12$ qubits. We gradually increase the depth of the trainable layers a)-f), alternating global and local observables. Each color corresponds to a certain frequency ω , squares of the same color correspond to the variance theory for 2-design of c_ω and the dots correspond to the simulated variance.

we stop the simulations at $n = 12$ qubits.

For the case of a reuploading circuit, we provide in Appendix C similar plots showing again that simulation and theory in Theorem 6 coincide. Precisely, we reproduce in Fig.15 similar plots as in Fig.7 for a reuploading model with $L = 2$ circuit layers.

5.3 Trainable layers as ε -approximate 2-design

In this section, we simulate parameterized circuits which gradually move away from being an exact 2-design. To do so, we decrease the number of qubits on which the trainable blocks, within a trainable layer, are acting on (or *connectivity*). We consider that each local trainable block (noted V in the figures) forms

a 2-design on the subset of qubits on which it acts non-trivially by using five repetitions of the strongly entangling ansatz applied to the corresponding m -qubit subsystem. We note that the smaller m (low connectivity), the entanglement between qubits decreases. We also observe numerically that the lower the connectivity, the higher the value of ε , which corresponds to the monomial distance from being a 2-design (see Definition 5). The numerical computation of this distance is expensive and hence we only compute it for $n = 4$ qubits. Specifically, we calculate the 2^{4n} -dimensional matrix $\mathcal{A}_{U(\Theta)}^{(2)}$ as given in Definition 5 and keep its biggest coefficient in absolute value. To do so, we obtain the Haar second-moment operator with the Weingarten Calculus [34] and




Figure 13: Variance of the Fourier coefficients in the local observable backward lightcone with Pauli encoding generated by three different methods: simulated variance in green, variance bound from Theorem 3 in pink and variance under the 2-design assumption on each of the lightcone trainable layers in blue. We gradually decrease the depth L_2 of the second trainable layer, from a)-c).




Figure 14: Variance of the Fourier coefficients in the local observable backward lightcone with $L_2 = 2$ generated by three different methods: simulated variance in green, variance bound from Theorem 3 in pink and variance under the 2-design assumption on each of the lightcone trainable layers in blue. a) Exponential encoding b) Golomb encoding.

the second-moment operator generated by the trainable unitary by computing the empirical average over different instances of $W(\theta)$ using randomly picked parameter vectors θ s. We ensured that the sample size used to compute the empirical average is large enough to have a low variance on the computed average values.

In Fig.9, the block V acts non trivially on $m = 1, m = 2$ and $m = 4$ qubits in Fig.9a, Fig.9b and Fig.9c respectively. For each case, we compare the simulated values of the variance of c_ω with two theoretical values: the upper bound from Theorem 7 (approximate 2-design), and the result from Theorem 1 (exact 2-design) using the global observable O_G . We notice that as

ε decreases, the simulated variance gets closer to the local 2-design result from Theorem 1. Moreover, the bound from Theorem 7 is satisfied and seems to be highly correlated with the simulated variance. In addition, it gets tighter as we approach the 2-design setting, exactly coinciding with the result from Theorem 1 when the trainable unitaries form an exact 2-design (See Fig.9c).

We also consider the exponential and Golomb encoding strategies with the same circuit architecture as the one used in Fig.9b ($m = 2$ and global observable O_G) to further assess the quality of the bound given in Theorem 7. In Fig.10, we see again that the bound and

the simulated variances are highly correlated, showing both a clear dependence on the redundancies.

In addition, if we consider the same setting but using the local observable O_L , we get the results plotted in Fig.11. First, we notice that the simulated spectrum does not cover all of the theoretical spectrum. This observation can be explained by the fact that the casual light cone of the local observable does not cover the encoding gates all at once, generating a sub-spectrum of the one observed in Fig.10 where the global observable O_G is used. Secondly, we see that the upper bound is still satisfied but is more loose compared to the global observable case.

5.4 Trainable layers made of Local 2-design blocks

In the following, we will be using a brickwise circuit made of $m = 2$ local trainable blocks forming an exact 2-design on the corresponding 2-qubit subsystem. We also consider a single layer circuit ($L = 1$) as depicted in Fig.4. To do so, we take trainable layers made of repetitions of the *Simplified Two Design Ansatz* described in Fig.18 and we refer to the repetition number by the depth of the trainable layer. In Fig.12, we gradually increase the depth of the trainable layers, each from 1 to 20, which should reach a global 2-design eventually for the number of qubits we are considering [28]. We also distinguish in this setting the case of local and global observables.

First, we observe that as the depth increases, the simulated variance gets closer to the 2 design theory with the global observable case converging more quickly, matching it exactly in Fig.12f and Fig.12h (depth 10 and 20 respectively for the global observable) and in Fig.12g (depth 20 for the local observable).

For the local observable with trainable layers of depth 1 (See Fig.12a), we observe the absence of some frequencies. This can be explained by the locality of the observable, the low connectivity of the trainable layers and the fact that the depth is smaller than $\log(n) \forall n > 2$. Indeed, for larger depths (bigger than $\log(n)$), we see that we do cover all of the theoretical spectrum because the light cone covers all of the encoding gates.

Secondly, the vanishing phenomenon is observed in all cases for the majority of the

frequencies but with different slopes except for the case of local observable and depth 1 in Fig.12a (consistent with [27]). For the same trainable layers depth, The steepness of the slope is smaller when considering the local observable and it increases with the depth for a fixed locality of the observable.

Besides, under this Brick-wise setting, we confirm the validity of the bound derived in Theorem 3 illustrated in Fig.13. Here we extract the lightcone for different depths (L_2) of the second trainable layer $W^{(2)}$ and plot the simulated variance of each Fourier coefficient generated by the light cone, the bound from Theorem 3 as well as the variance under the 2-design assumption on the light cone trainable layers from Eq.(B95). We perform the simulations in Fig.13 for the Pauli encoding strategy and replicate it for the exponential and Golomb encoding strategies in the fixed setting of $L_2 = 2$ in Fig.14.

As seen before, we observe that the bound from Theorem 3 is satisfied and seems highly correlated with the simulated variance for all the encoding strategies considered.

Finally, when we assume that the restriction of each of the trainable layers to the light cone forms a 2-design (on the subsystems on which it acts non trivially), we observe that the Fourier coefficients variance in this setting is below the simulated variance. This observation supports the fact that the variance in the *2-design lightcone* setting may provide a lower bound for the variance over the light cone. Thus, we leave this question for a future work.

6 Conclusion

In this work, we have studied the expressivity of parameterized quantum circuits (PQCs) used for learning purposes, through Fourier lens. Our results imply that the choice of the encoding Hamiltonians defines two important properties of the quantum model expressivity: the encoding *inductive bias* and the *vanishing expressivity* phenomena. The former captures the dependence of each Fourier coefficient's variance, and therefore its degree of freedom, on the corresponding frequency redundancy. The latter is a statement about the scaling of the Fourier coefficients variance as the number of

qubits grows. Specifically, we showed that models containing frequencies with relatively low redundancies suffer from a vanishing variance of those frequencies, reducing their expected expressivity.

The learning models obtained with PQC are therefore more constrained than expected. Even though this reduces their expressivity, it also indicates an inductive bias which might be specific to quantum models. It is left as an open question to study potential benefits of such bias as a source of quantum advantage. On the other hand, could one try to classically reproduce this bias, using the fact that this bias can be determined a priori by looking at the circuit? Overall, one can use these guidelines to design more expressive PQCs, and to further study what differentiates the quantum Fourier model from its classical equivalent.

7 Code availability

Code based on PennyLane [35] to reproduce the figures and analysis is available at the following repository: https://github.com/quantumsoftwarelab/DRM_fourier_expressivity.

Acknowledgments

This work is supported by the H2020-FETOPEN Grant PHOQUSING (GA no.: 899544), the Engineering and Physical Sciences Research Council (grants EP/T001062/1), and the Naval Group Centre of Excellence for Information Human factors and Signature Management (CEMIS).

References

- [1] Nathan Wiebe, Ashish Kapoor, and Krysta M. Svore. “Quantum deep learning” (2015). [arXiv:1412.3489](https://arxiv.org/abs/1412.3489).
- [2] M. Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C. Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R. McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick J. Coles. “Variational quantum algorithms”. *Nature Reviews Physics* **3**, 625–644 (2021).
- [3] Eric R. Anschuetz and Bobak T. Kiani. “Quantum variational algorithms are swamped with traps”. *Nature Communications* **13** (2022).
- [4] Andrew Arrasmith, Zoë Holmes, M Cerezo, and Patrick J Coles. “Equivalence of quantum barren plateaus to cost concentration and narrow gorges”. *Quantum Science and Technology* **7**, 045015 (2022).
- [5] Martin Larocca, Nathan Ju, Diego García-Martín, Patrick J. Coles, and M. Cerezo. “Theory of overparametrization in quantum neural networks”. *Nature Computational Science* **3**, 542–551 (2023).
- [6] Martin Larocca, Piotr Czarnik, Kunal Sharma, Gopikrishnan Muraleedharan, Patrick J. Coles, and M. Cerezo. “Diagnosing Barren Plateaus with Tools from Quantum Optimal Control”. *Quantum* **6**, 824 (2022).
- [7] Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven. “Barren plateaus in quantum neural network training landscapes”. *Nature Communications* **9**, 4812 (2018).
- [8] Christa Zoufal. “Generative Quantum Machine Learning” (2021). [arXiv:2111.12738](https://arxiv.org/abs/2111.12738) [quant-ph].
- [9] Maria Schuld, Ryan Sweke, and Johannes Jakob Meyer. “The effect of data encoding on the expressive power of variational quantum machine learning models”. *Physical Review A* **103**, 032430 (2021).
- [10] Jonas Landman, Slimane Thabet, Constantin Dalyac, Hela Mhiri, and Elham Kashefi. “Classically approximating variational quantum machine learning with random fourier features” (2022). [arXiv:2210.13200](https://arxiv.org/abs/2210.13200).
- [11] Franz J. Schreiber, Jens Eisert, and Johannes Jakob Meyer. “Classical surrogates for quantum learning models”. *Physical Review Letters* **131**, 100803 (2023).
- [12] Sofiene Jerbi, Lukas J. Fiderer, Hendrik Poulsen Nautrup, Jonas M. Kübler, Hans J. Briegel, and Vedran Dunjko. “Quantum machine learning beyond kernel methods”. *Nature Communications* **14**, 517 (2023).

- [13] Ryan Sweke, Erik Recio-Armengol, Sofiene Jerbi, Elies Gil-Fuster, Bryce Fuller, Jens Eisert, and Johannes Jakob Meyer. “Potential and limitations of random Fourier features for dequantizing quantum machine learning”. *Quantum* **9**, 1640 (2025).
- [14] Adrián Pérez-Salinas, Alba Cervera-Lierta, Elies Gil-Fuster, and José I. Latorre. “Data re-uploading for a universal quantum classifier”. *Quantum* **4**, 226 (2020).
- [15] Evan Peters and Maria Schuld. “Generalization despite overfitting in quantum machine learning models”. *Quantum* **7**, 1210 (2023).
- [16] S. Shin, Y. S. Teo, and H. Jeong. “Exponential data encoding for quantum supervised learning” (2022).
- [17] Ben Jaderberg, Antonio A. Gentile, Youssef Achari Berrada, Elvira Shishenina, and Vincent E. Elfving. “Let Quantum Neural Networks Choose Their Own Frequencies” (2023). 1 citations (Semantic Scholar/arXiv) [2024-02-22] 1 citations (Semantic Scholar/DOI) [2024-02-22] arXiv:2309.03279 [quant-ph].
- [18] Sofiene Jerbi, Casper Gyurik, Simon C. Marshall, Riccardo Molteni, and Vedran Dunjko. “Shadows of quantum machine learning”. *Nature Communications* **15**, 5676 (2024).
- [19] Slimane Thabet, Léo Monbroussou, Elliott Z. Mamon, and Jonas Landman. “When quantum and classical models disagree: Learning beyond minimum norm least square” (2025). arXiv:2411.04940.
- [20] Matthias C. Caro, Elies Gil-Fuster, Johannes Jakob Meyer, Jens Eisert, and Ryan Sweke. “Encoding-dependent generalization bounds for parametrized quantum circuits”. *Quantum* **5**, 582 (2021).
- [21] Alice Barthe and Adrián Pérez-Salinas. “Gradients and frequency profiles of quantum re-uploading models”. *Quantum* **8**, 1523 (2024).
- [22] Weijie Xiong, Giorgio Facelli, Mehrad Sahebi, Owen Agnel, Thiparat Chotibut, Supanut Thanasilp, and Zoë Holmes. “On fundamental aspects of quantum extreme learning machines”. *Quantum Machine Intelligence* **7**, 20 (2025).
- [23] Zoë Holmes, Kunal Sharma, M. Cerezo, and Patrick J. Coles. “Connecting ansatz expressibility to gradient magnitudes and barren plateaus”. *PRX Quantum* **3**, 010313 (2022).
- [24] Sukin Sim, Peter D. Johnson, and Alan Aspuru-Guzik. “Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms”. *Advanced Quantum Technologies* **2**, 1900070 (2019).
- [25] Richard A. Low. “Pseudo-randomness and Learning in Quantum Computation” (2010). 68 citations (Semantic Scholar/arXiv) [2024-01-13] arXiv:1006.5227 [quant-ph].
- [26] Martín Larocca, Supanut Thanasilp, Samson Wang, Kunal Sharma, Jacob Biamonte, Patrick J. Coles, Lukasz Cincio, Jarrod R. McClean, Zoë Holmes, and M. Cerezo. “Barren plateaus in variational quantum computing”. *Nature Reviews Physics* **7**, 174–189 (2025).
- [27] M. Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J. Coles. “Cost Function Dependent Barren Plateaus in Shallow Parametrized Quantum Circuits”. *Nature Communications* **12**, 1791 (2021).
- [28] Aram W. Harrow and Saeed Mehraban. “Approximate Unitary t-Designs by Short Random Quantum Circuits Using Nearest-Neighbor and Long-Range Gates”. *Communications in Mathematical Physics* **401**, 1531–1626 (2023).
- [29] Enrico Fontana, Dylan Herman, Shouvanik Chakrabarti, Niraj Kumar, Romina Yalovetzky, Jamie Heredge, Shree Hari Sureshbabu, and Marco Pistoia. “Characterizing barren plateaus in quantum ansätze with the adjoint representation”. *Nature Communications* **15**, 7171 (2024).
- [30] Michael Ragone, Bojko N. Bakalov, Frédéric Sauvage, Alexander F. Kemper, Carlos Ortiz Marrero, Martín Larocca, and M. Cerezo. “A Lie algebraic theory of barren plateaus for deep parameterized quantum circuits”. *Nature Communications* **15**, 7172 (2024).

- [31] Léo Monbroussou, Elliott Z. Mamon, Jonas Landman, Alex B. Grilo, Romain Kukla, and Elham Kashefi. “Trainability and Expressivity of Hamming-Weight Preserving Quantum Circuits for Machine Learning”. *Quantum* **9**, 1745 (2025).
- [32] Eric R. Anschuetz, Andreas Bauer, Bobak T. Kiani, and Seth Lloyd. “Efficient classical algorithms for simulating symmetric quantum systems”. *Quantum* **7**, 1189 (2023).
- [33] M. Cerezo, Martin Larocca, Diego García-Martín, N. L. Diaz, Paolo Braccia, Enrico Fontana, Manuel S. Rudolph, Pablo Bermejo, Aroosa Ijaz, Supanut Thanasilp, Eric R. Anschuetz, and Zoë Holmes. “Does provable absence of barren plateaus imply classical simulability? Or, why we need to rethink variational quantum computing” (2023). 13 citations (Semantic Scholar/arXiv) [2024-03-14] 13 citations (Semantic Scholar/DOI) [2024-03-14] arXiv:2312.09121 [quant-ph, stat].
- [34] Antonio Anna Mele. “Introduction to Haar Measure Tools in Quantum Information: A Beginner’s Tutorial” (2023). arXiv:2307.08956 [quant-ph].
- [35] Ville Bergholm, Josh Izaac, Maria Schuld, Christian Gogolin, Shahnawaz Ahmed, Vishnu Ajith, M. Sohaib Alam, Guillermo Alonso-Linaje, B. AkashNarayanan, Ali Asadi, Juan Miguel Arrazola, Utkarsh Azad, Sam Banning, Carsten Blank, Thomas R. Bromley, Benjamin A. Cordier, Jack Ceroni, Alain Delgado, Olivia Di Matteo, Amintor Dusko, Tanya Garg, Diego Guala, Anthony Hayes, Ryan Hill, Aroosa Ijaz, Theodor Isacsson, David Ittah, Soran Jahangiri, Prateek Jain, Edward Jiang, Ankit Khandelwal, Korbinian Kottmann, Robert A. Lang, Christina Lee, Thomas Loke, Angus Lowe, Keri McKiernan, Johannes Jakob Meyer, J. A. Montañez-Barrera, Romain Moyard, Zeyue Niu, Lee James O’Riordan, Steven Oud, Ashish Panigrahi, Chae-Yeun Park, Daniel Polatajko, Nicolás Quesada, Chase Roberts, Nahum Sá, Isidor Schoch, Borun Shi, Shuli Shu, Sukin Sim, Arshpreet Singh, Ingrid Strandberg, Jay Soni, Antal Száva, Slimane Thabet, Rodrigo A. Vargas-Hernández, Trevor Vincent, Nicola Vitucci, Maurice Weber, David Wierichs, Roeland Wiersema, Moritz Willmann, Vincent Wong, Shaoming Zhang, and Nathan Killoran. “PennyLane: Automatic differentiation of hybrid quantum-classical computations” (2022). 0 citations (Semantic Scholar/arXiv) [2024-03-12] arXiv:1811.04968 [physics, physics:quant-ph].
- [36] Motohisa Fukuda, Robert Koenig, and Ion Nechita. “RTNI - A symbolic integrator for Haar-random tensor networks”. *Journal of Physics A: Mathematical and Theoretical* **52**, 425303 (2019).
- [37] Maria Schuld, Alex Bocharov, Krysta Svore, and Nathan Wiebe. “Circuit-centric quantum classifiers”. *Physical Review A* **101**, 032308 (2020).

A Appendix: Preliminaries

A.1 Quantum Models with Hamiltonian encoding as Large Fourier series

For completeness, we revisit the proof in [9], showing that a Quantum model with Hamiltonian encoding can be expressed as a truncated Fourier series of the form in Eq.(1), with a spectrum constructed from the encoding Hamiltonians eigenvalues defined in Eq.(5).

We start by detailing the proof for the one-dimensional input case ($x \in \mathbb{R}$) and generalize it afterwards to the D -dimensional input case ($x \in \mathbb{R}^D$).

To do so, we consider an L layered ansatz of the following form $U(x; \theta) = W^{(L+1)}(\theta)S^{(L)}(x) \dots S^{(1)}(x)W^{(1)}(\theta)$ as described in Eq.(2). In what follows, we drop the explicit dependence of the trainable layers $W^{(l)}(\theta)$ on the parameter vector θ for simplicity and recall that the trainable and encoding unitaries are d -dimensional matrices acting on n qubits ($d = 2^n$).

We consider the Hamiltonian encoding scheme where the classical input $x \in \mathbb{R}$ is encoded as the time evolution of some Hamiltonian. Thus, the encoding unitary in each layer l is of the form $S^l(x) = e^{-ixH_l} \quad \forall l \in \{1, \dots, L\}$. If one considers $H_l = P_l D_l P_l^\dagger$ and injects P_l and $P_l^\dagger$ in the parameterized unitaries W^{l+1} and W^l respectively, then $S^l(x)$ can be rewritten in the following form $S^l(x) = \text{diag}(e^{-ix\lambda_1^l}, \dots, e^{-ix\lambda_d^l})$ where λ_i^l s are eigenvalues of the underlying d -dimensional encoding Hamiltonian H_l .

We start by applying the first layer $S^{(1)}(x)W^{(1)}$ on the $|0\rangle$ computational basis state and we iterate through the remaining layers to obtain $|\psi(x; \theta)\rangle = U(x; \theta)|0\rangle$.

$$\begin{aligned} S^{(1)}(x)W^{(1)}|0\rangle &= \sum_{j_1=1}^d W_{j_1 1}^{(1)} e^{-ix\lambda_{j_1}^1} |j_1\rangle \\ S^{(2)}(x)W^{(2)}S^{(1)}(x)W^{(1)}|0\rangle &= \sum_{j_2=1}^d \sum_{j_1=1}^d W_{j_2 j_1}^{(2)} W_{j_1 1}^{(1)} e^{-ix(\lambda_{j_1}^1 + \lambda_{j_2}^2)} |j_2\rangle \\ &\vdots \\ W^{(L+1)} \prod_{l=1}^L S^{(l)}(x)W^{(l)}|0\rangle &= \sum_{k=1}^d \sum_{\mathbf{J} \in [d]^L} W_{k j_L}^{(L+1)} \dots W_{j_1 1}^{(1)} e^{-ix\Lambda_{\mathbf{J}}} |k\rangle \end{aligned}$$

where $\Lambda_{\mathbf{J}} = \sum_{j_l \in \mathbf{J}} \lambda_{j_l}$ and $\mathbf{J} = (j_1, \dots, j_L) \in [1, d]^L$.

The full quantum model can be written as a truncated Fourier series:

$$\begin{aligned} f(x; \theta) &= \langle 0|U(x; \theta)^\dagger O U(x; \theta)|0\rangle = \sum_{k, k'=1}^d \sum_{\mathbf{J}, \mathbf{J}' \in [d]^L} (W_{k' j'_L}^{(L+1)} \dots W_{j'_1 1}^{(1)})^* \langle k'|O|k\rangle (W_{k j_L}^{(L+1)} \dots W_{j_1 1}^{(1)}) e^{-ix(\Lambda_{\mathbf{J}} - \Lambda_{\mathbf{J}'})} \\ &= \sum_{\omega \in \Omega} \sum_{\mathbf{J}, \mathbf{J}' \in R(\omega)} c_\omega e^{-ix\omega} \end{aligned}$$

Hence, the expression of the Fourier coefficient is obtained by grouping multi-indices $(\mathbf{J}, \mathbf{J}')$ that lead to the same frequency ω (i.e. $\Lambda_{\mathbf{J}} - \Lambda_{\mathbf{J}'} = \omega$), i.e. in the frequency generator $R(\omega)$:

$$c_\omega = \sum_{\mathbf{J}, \mathbf{J}' \in R(\omega)} \sum_{k, k'} W_{j'_1 0}^{(1)*} W_{j'_2 j'_1}^{(2)*} \dots W_{k' j'_L}^{(L+1)*} O_{k' k} W_{k j_L}^{(L+1)} \dots W_{j_2 j_1}^{(2)} W_{j_1 0}^{(1)} \quad (\text{A1})$$

with

$$\begin{aligned} \Omega &= \left\{ \sum_{l \in [L]} (\lambda_{j_l}^l - \lambda_{j'_l}^l) \mid (j_1, \dots, j_L), (j'_1, \dots, j'_L) \in [1, d]^L \right\} \\ R(\omega) &= \left\{ (\mathbf{J}, \mathbf{J}') \in [1, d]^L \times [1, d]^L \mid \sum_{l \in [L]} (\lambda_{j_l}^l - \lambda_{j'_l}^l) = \omega \right\} \end{aligned} \quad (\text{A2})$$

Here we recall that the trainable unitaries $W^{(l)}$ implicitly contain the projector P_l formed by the eigenvectors of the encoding Hamiltonian H_l . Thus, in the following section, whenever an assumption on the distribution of the trainable unitaries is made, we should take into account their multiplication by projectors P_l .

Moreover, we note that using encoding Hamiltonians that act non trivially only on a subset of the n qubits will increase the eigenvalues degeneracy. Hence the redundancies $|R(\omega)|$ of all the frequencies $\omega \in \Omega$ according to Definition 1 will be multiplied by the same factor.

A.1.1 Quantum Fourier model with high-dimensional input

We can now generalize the Quantum model Fourier Series representation to the D -dimensional input setting. Assume that for $x = (x_1, \dots, x_D) \in \mathbb{R}^D$, each component x_k is encoded as the time evolution of some Hamiltonian $H_l^{(k)}$ in the l^{th} uploading layer, such that $S^l(x) = \prod_{k=1}^D S^l(x_k) = \prod_{k=1}^D e^{-ix_k H_l^{(k)}}$.

In general, writing down the expression of a Fourier coefficient with high dimensional input data is intricate. For simplicity, we give the expression of the a Fourier coefficient c_ω for the frequency vector $\omega = (\omega_1, \dots, \omega_D)$ under the assumption that the encoding unitaries of the different components commute. Precisely, for each layer l , $\{S^l(x_k)\}_{k=1}^D$ mutually commute (i.e. equivalently $\{H_l^{(k)}\}_{k=1}^D$ mutually commute). Using this assumption, the global encoding unitary in the l^{th} layer $S^l(x)$ can be diagonalized as

$$S^l(x) = P_l \prod_{k=1}^D e^{-ix_k D_l^{(k)}} P_l^\dagger = P_l e^{-i \sum_{k=1}^D x_k D_l^{(k)}} P_l^\dagger = P_l \left(\sum_{i=1}^d \sum_{k=1}^D x_k \lambda_i^{k,l} |i\rangle\langle i| \right) P_l^\dagger, \quad (\text{A3})$$

where P_l is an orthogonal projector and each $D_l^{(k)} := \text{diag}(\lambda_1^{k,l}, \dots, \lambda_d^{k,l})$ is the diagonal matrix of eigenvalues of the Hamiltonian $H_l^{(k)}$. Similarly to the analysis done for a one-dimensional input vector, the quantum spectrum in the D -dimensional setting will be of the form

$$\Omega = \{\omega = (\omega_k)_{k=1}^D, \omega_k = \lambda_i^k - \lambda_j^k, \forall i, j \in [1, 2^n]\} \quad (\text{A4})$$

Here, we see that the maximal size of the spectrum is independent of the dimensionality D , i.e. $|\Omega| \leq 2^{2n}$. This independence is mainly due to the assumption that the encoding Hamiltonians commute. However, in the most general case, the D -dimensional spectrum can scale up to $\mathcal{O}(\prod_{k=1}^D |\Omega_k|)$ where Ω_k is the spectrum constructed solely from the eigenvalues of Hamiltonians H_k^l as defined in Eq.(A2).

Consequently, it can be easily seen that all of our theoretical results can generalize to high-dimensional input vectors when the encoding Hamiltonians commute within a single uploading layer. Precisely, there is a one to one correspondence between frequency vectors $\omega \in \Omega$ and computational basis states. Hence, the associated Fourier coefficient can be simply expressed in a very similar way to the one dimensional case in Eq.(A1). Consequently, the proofs of our main theorems trivially apply to this setting. However, our techniques no longer hold if we consider non commuting encoding unitaries.

A.2 Spectrum distribution

From the above spectrum construction, one can make several observations regarding the spectrum structure and its size by considering different Hamiltonian encoding schemes. We further comment on how the spectrum distribution interacts with the Fourier coefficient expression in Eq.(A1).

In what follows, we will focus on the one-dimensional setting ($x \in \mathbb{R}$).

Lemma 1. (*Sequential and parallel encoding*) Consider the spectrum Ω_{parallel} obtained from a re-uploading circuit with L layers and encoding Hamiltonians $H_1, \dots, H_L$ acting each on n -qubits, such that $H_l = P_l D_l P_l^\dagger, \forall l \in [L]$. The spectrum Ω_{parallel} is the same as the spectrum $\Omega_{\text{sequential}}$ obtained by considering a single layered circuit acting on nL -qubits with an encoding Hamiltonian of the form $H = P \sum_{l=1}^L D_l^{(nl, n(l+1))} P^\dagger$ where $D_l^{(nl, n(l+1))}$ acts non trivially on the subset of the total nL qubits indexed from nl to $n(l+1)$ and $P = P_1 \otimes P_2 \otimes \dots \otimes P_L$.

Proof. We consider a single layer reuploading circuit of the form in Eq.(2) ($L = 1$) with the following encoding layer unitary acting on nL qubits:

$$\begin{aligned} S(x) &:= e^{-ixH_1} \otimes e^{-ixH_2} \otimes \dots e^{-ixH_L} \\ &= P_1 e^{-ixD_1} P_1^\dagger \otimes P_2 e^{-ixD_2} P_2^\dagger \otimes P_L \dots e^{-ixD_L} P_L^\dagger \\ &= P e^{-ix \sum_{l=1}^L D_l^{(nl, n(l+1))}} P^\dagger \end{aligned}$$

where $H_l = P_l D_l P_l^\dagger$, $D_l = \text{diag}(\lambda_1^l, \dots, \lambda_d^l)$ contains the eigenvalues of the Hamiltonian H_l and $P = P_1 \otimes P_2 \otimes \dots \otimes P_L$. $\square$

This result was already proven in [9] to convey that we can use the encoding gates in sequence or in parallel (reuploading model) to get the same spectrum. In practice, we will have a combination of the *sequential* and the *parallel* encodings since each encoding layer in the reuploading model is usually made of one-qubit or two-qubit gates. Hence a more general form of the spectrum is the following:

$$\Omega = \left\{ \sum_{l \in [L]} \sum_{k \in [K_l]} (\lambda_{j_l, k}^{l, k} - \lambda_{j'_l, k}^{l, k}) \mid (j_{1,1}, \dots, j_{1,K}, \dots, j_{L,1}, \dots, j_{L,K}), (j'_{1,1}, \dots, j'_{1,K}, \dots, j'_{L,1}, \dots, j'_{L,K}) \in \llbracket 1, 2^{n/K} \rrbracket^{KL} \right\} \quad (\text{A5})$$

where K_l is the number of blocks of the l^{th} encoding layer and $\lambda_{j_l, k}^{l, k}$ is the eigenvalue of the local Hamiltonian $H_{l,k}$ acting non trivially on the k^{th} subsystem in layer l .

From the above construction, one can easily see that the spectrum Ω is conserved under permutation of the encoding layers and even under permutation of tensor product unitaries within the layers. However, the permutation invariance of the spectrum does not imply that we get the same quantum model under encoding blocks permutation since the expression of the Fourier coefficients in Eq.(A1) is not invariant under permutation of the indices in $(J, J') \in R(\omega)$.

We introduce the *partial* frequency generator as it will serve to derive Theorem 6 and to make comments on the spectrum distribution in general.

Definition 4 (Partial Frequency Generator). *Considering an L -layered quantum Fourier model as described in Eqs.(1-5), we denote by $R_h^l(\omega)$ the ensemble of eigenvalue indices giving rise to the frequency ω by considering only encoding layers from h to l . The partial redundancy $|R_h^l(\omega)|$ is the size of this set:*

$$R_h^l(\omega) = \left\{ ((j_h, \dots, j_l), (j'_h, \dots, j'_l)) \in \llbracket 1, d \rrbracket^{l-h+1} \times \llbracket 1, d \rrbracket^{l-h+1} \mid \sum_{k=h}^l (\lambda_{j_k}^k - \lambda_{j'_k}^k) = \omega \right\} \quad (\text{A6})$$

For simplicity, we denote the frequency generator $R_1^L(\omega)$ by $R(\omega)$ as defined in Definition 1 and the frequency generator from a single layer $R_l^l(\omega)$ by $R_l(\omega)$.

Lemma 2 (Recursive spectrum construction). *We recall that $R_h^L(\omega) = \{(J := (j_h, \dots, j_L), J' := (j'_h, \dots, j'_L)) \in [d]^{L-h+1} \times [d]^{L-h+1} ; \Lambda_J - \Lambda_{J'} = \omega\}$. We denote by Ω^l the spectrum obtained from the l^{th} encoding layer and its corresponding frequency generator $R_l(\omega)$ as defined in Definition 4. Consider a Quantum Fourier model on n -qubits with L layers and a spectrum Ω . For a fixed frequency $\omega \in \Omega$, we have the following recursive relation between $R_h^L(\omega)$ and $R_{h+1}^L(\omega)$:*

$$R_h^L(\omega) = \bigcup_{\substack{j_h, j'_h=1 \\ k_h = \lambda_{j_h} - \lambda_{j'_h}}}^d \{R_{h+1}^L(\omega - k_h) \times \{(j_h, j'_h)\}\}$$

Consequently, $|R_1^L(\omega)| = \sum_{k_L, \dots, k_2, k_1} \beta_{k_L} \dots \beta_{k_1} \delta_{\sum_{l=1}^L k_l}^\omega$ with $k_l \in \Omega_{\text{distinct}}^l, \forall l \in [L]$, Ω^l is the spectrum generated by the l^{th} encoding layer and $\beta_{k_l} := |R_l(k_l)|$ is the redundancy of the frequency k_l in Ω^l .

Proof. Let $(J, J') \in R_h^L(\omega)$,

$$\omega = \Lambda_J - \Lambda_{J'} = \sum_{l=h}^L (\lambda_{j_l} - \lambda_{j'_l}) = \sum_{l=h+1}^L (\lambda_{j_l} - \lambda_{j'_l}) + (\lambda_{j_h} - \lambda_{j'_h})$$

$$R_h^L(\omega) = \bigcup_{\substack{j_h, j'_h=1 \\ k_h = \lambda_{j_h} - \lambda_{j'_h}}}^d \{R_{h+1}^L(\omega - k_h) \times \{(j_h, j'_h)\}\}$$

$$|R_h^L(\omega)| = \sum_{\substack{j_h, j'_h=1 \\ k_h = \lambda_{j_h} - \lambda_{j'_h}}}^d |R_{h+1}^L(\omega - k_h)| = \sum_{k_h \in \Omega_{distinct}^h} |R_h(k_h)| |R_{h+1}^L(\omega - k_h)| = \sum_{k_h \in \Omega_{distinct}^h} \beta_{k_h} |R_{h+1}^L(\omega - k_h)|$$

Therefore we obtain : $|R_1^L(\omega)| = \sum_{k_L, \dots, k_2} \beta_{k_L} \dots \beta_{k_2} |R_1(\omega - \sum_{l=2}^L k_l)| = \sum_{k_L, \dots, k_2, k_1} \beta_{k_L} \dots \beta_{k_1} \delta_{\sum_{l=1}^L k_l}^\omega$ with each $k_l \in \Omega_{distinct}^l$, Ω^l is the spectrum generated by the l^{th} encoding layer and β_{k_l} is the redundancy of the frequency k_l in Ω^l . $\square$

Whether we are considering sequential encoding, parallel encoding or both, the result of Lemma 2 enables us to track the evolution of the paths in the quantum tree (See Fig.2) leading to a certain frequency and mainly to characterize the size of the spectrum $|\Omega|$ and the frequency redundancies $|R(\omega)|$.

First, one can easily see that $\forall l \in [L], |R_l(0)| \geq 2^n$. Hence the recursive construction of the spectrum implies that all frequencies that can be generated from a subset of $\kappa < L$ layers have an exponential redundancy in the number of qubits and in the number of the remaining $(L - \kappa)$ layers. Specifically, for these frequencies we have $|R(\omega)| \geq 2^{n(L-\kappa)}$.

Consequently, adopting the reuploading scheme leads inevitably to exponential redundancies in n as mentioned in [21]. Moreover, if one considers tensor product encoding unitaries within a single encoding layer (i.e. Pauli rotations), then the frequencies already generated by few blocks will have exponential redundancies in n .

We give some examples of encoding Hamiltonians and explicit the obtained spectrum, its size and the redundancies scaling in each case. For more details, we refer the reader to a discussion of the degeneracy of a Quantum Fourier model spectrum in [15] (See Appendix C).

A.2.1 Pauli encoding

The Pauli encoding scheme consists in using Pauli matrices $\{\sigma_x, \sigma_y, \sigma_z\}$ as encoding Hamiltonians, such that a single encoding layer is a tensor product of Pauli rotations parameterized by x : $S^l(x) := \otimes_{i=1}^n e^{-ix\sigma/2}$, $\sigma \in \{\sigma_x, \sigma_y, \sigma_z\}$. In this case, all eigenvalues are $\lambda = \pm 1/2$. Hence by using Eq.A5, the set of frequencies we obtain by applying Pauli rotations on all n -qubits in each of the L layers are consecutive integer frequencies from $-nL$ to nL . Hence the size of the spectrum $|\Omega| = 2nL$ and $|R(k)| = \binom{2nL}{nL-|k|} \forall k \in \llbracket -nL, nL \rrbracket$. The redundancy value can be obtained by using Lemma 2.

A.2.2 Exponential encoding

If we consider single qubit Pauli rotations for encoding with a scaling coefficient, the largest spectrum we can get is of size $|\Omega| = 3^{nL}$ since the spectrum of a scaled Pauli matrix is of the form $\{-\lambda, \lambda\}$. One spectrum that reaches this limit is the set of consecutive integer frequencies from $-\frac{3^{nL}}{2}$ to $\frac{3^{nL}}{2}$ with $|\Omega| = 3^{nL}$. This spectrum is obtained by using the scaling coefficient 3^{jl} in the Pauli rotation gate acting on qubit j in layer l . This scheme was introduced in [16] under the name of *exponential encoding*, which refers to the exponential size of the spectrum in the number of qubits/layers. We note that $|R(\omega = k)| = 2^{nL-T(|k|)}$ where $T(|k|)$ takes values in $\llbracket 0, nL \rrbracket$. More details about the proof are provided in [15, 16].

A.2.3 Non degenerate encoding (Golomb encoding)

As mentioned in Lemma 2, with the reuploading scheme we will certainly have frequencies with redundancies exponential in the number of layers. However to surpass the limit $|\Omega^l| = 3^n$ per layer l , one needs to consider non separable encoding unitaries in a single layer. A perfectly non degenerate spectrum for a single encoding layer ($|R(\omega)| = 1 \forall \omega \in \Omega \setminus \{0\}$) was proposed in [15] by setting the diagonal of the data-encoding Hamiltonian to the elements of a Golomb ruler. However, we stress again that using this scheme in the reuploading model will give rise to redundancies that are still exponential in nL for frequencies generated from few layers since the null frequency in a single layer has redundancy $|R(0)| = 2^n$.

B Appendix: Proofs

B.1 Fourier coefficients variance with 2-design trainable unitaries

We present the formal version of Theorem 1 showcasing an exact expression for the expectation and variance of Fourier coefficients corresponding to a single layer model.

Theorem 5. *Consider a quantum model of the form in Eq.(1) and a parametrized circuit of the form in Eq.(2) with $L = 1$ layers and fixed encoding Hamiltonians resulting in a spectrum Ω . We assume that each of the trainable layers $W^l(\theta)$, $l \in \{1, 2\}$ form independently a 2-design. The expectation and variance of each Fourier coefficient $c_\omega(\theta)$ for the frequencies $\omega \in \Omega$ appearing in the model Fourier decomposition in Eq.(3) are given by*

$$\begin{aligned} \mathbb{E}_\theta[c_\omega(\theta)] &= \frac{\text{Tr}(O)}{d} \delta_\omega^0, \\ \text{Var}_\theta[c_\omega(\theta)] &= \frac{d\|O\|_2^2 - \text{Tr}(O)^2}{d(d^2 - 1)} \frac{|R(\omega)|}{d(d+1)} + \frac{\text{Tr}(O)^2 - d\|O\|^2}{d^2(d^2 - 1)} \delta_\omega^0, \end{aligned} \quad (\text{B1})$$

where we recall that $d = 2^n$.

The following Theorem provides a generalization of Theorem 5 to the case of reuploading models with $L \geq 1$ layers. Indeed, for $L = 1$, we retrieve the results in Theorem 5. Hence, the proof of Theorem 6 covers also the proof of Theorem 5.

Theorem 6. *Consider a quantum model of the form in Eq.(1) and a parametrized circuit of the form in Eq.(2) with $L \geq 1$ layers and fixed encoding Hamiltonians resulting in a spectrum Ω . We assume that each of the trainable layers $W^l(\theta)$, $l \in \{1, \dots, L\}$ form independently a 2-design. The expectation and variance of each Fourier coefficient $c_\omega(\theta)$ for the frequencies $\omega \in \Omega$ appearing in the model Fourier decomposition in Eq.(3) are given by*

$$\begin{aligned} \mathbb{E}_\theta[c_\omega(\theta)] &= \frac{\text{Tr}(O)}{d} \delta_\omega^0, \\ \text{Var}_\theta[c_\omega(\theta)] &\simeq \frac{d\|O\|_2^2 - \text{Tr}(O)^2}{d(d^2 - 1)} \left[\frac{|R_1^L(\omega)| - |R_2^L(\omega)|}{d(d+1)(d^2 - 1)^{L-1}} + \sum_{j=3}^L \frac{|R_j^L(\omega)|}{d(d^2 - 1)^{L-j+2}} \right] + \frac{\text{Tr}(O)^2 - d\|O\|^2}{d^2(d^2 - 1)} \delta_\omega^0. \end{aligned} \quad (\text{B2})$$

Here $R_j^l(\omega)$ denotes the partial spectrum formed from the j^{th} up to the l^{th} encoding layers as defined in Definition 4.

Proof. We begin recall the expression of the Fourier coefficient derived in Appendix A.1.

$$c_\omega = \sum_{J, J' \in R(\omega)} \sum_{k, k'} W_{j_1' 0}^{(1)*} W_{j_2' j_1'}^{(2)*} \dots W_{k' j_L'}^{(L+1)*} O_{k' k} W_{k j_L}^{(L+1)} \dots W_{j_2 j_1}^{(2)} W_{j_1 0}^{(1)} \quad (\text{B3})$$

We start by proving the expression of the expectation value of a Fourier coefficient and note that it is sufficient to establish this result under the 1-design hypothesis. To do so, we apply the Weingarten formula for the first moment [34] and obtain

$$\begin{aligned}\mathbb{E}_{W^{(1)}, \dots, W^{(L+1)} \sim U(N)} [c_\omega] &= \sum_{\substack{k, k' \\ J, J' \in R(\omega)}} \frac{\delta_{j_1}^{j'_1} \delta_{j_2}^{j'_2} \dots \delta_{j_L}^{j'_L} \delta_k^{k'} O_{k'k}}{d^{L+1}} = \sum_{\substack{k \\ J, J' \in R(\omega)}} \frac{\delta_{j_1}^{j'_1} \delta_{j_2}^{j'_2} \dots \delta_{j_L}^{j'_L} O_{kk}}{d^{L+1}} = \sum_{J, J' \in R(\omega)} \delta_J^{J'} \frac{\text{Tr}(O)}{d^{L+1}} \\ &= \frac{\text{Tr}(O)}{d} \delta_\omega^0\end{aligned}\quad (\text{B4})$$

The variance of a Fourier coefficient for a reuploading VQC (i.e $L > 1$) is obtained recursively starting from the variance of a single-layered circuit and using the recursive relation between the partial redundancies $R_h^L(\omega)$ established in Lemma 2.

From Eq.(A1), we give the expression of $\mathbb{E}_{W^{(1)}, \dots, W^{(L+1)} \sim U(N)} [|c_\omega|^2]$

$$\mathbb{E}_{W^{(1)}, \dots, W^{(L+1)} \sim U(N)} [|c_\omega|^2] = \sum_{\substack{I, I' \in R(\omega) \\ J, J' \in R(\omega)}} \prod_{l=1}^{L+1} X_l(I, I', J, J') \quad (\text{B5})$$

where we introduce the shorthands

$$X_l(I, I', J, J') = \begin{cases} \mathbb{E} \left[W_{j_1 1}^{(1)} W_{i'_1 1}^{(1)} W_{j'_1 1}^{(1)*} W_{i_1 1}^{(1)*} \right] & \text{if } l = 1, \\ \mathbb{E} \left[W_{j_l j_{l-1}}^{(l)} W_{i'_l i'_{l-1}}^{(l)} W_{j'_l j'_{l-1}}^{(l)*} W_{i_l i_{l-1}}^{(l)*} \right] & \text{if } 2 \leq l \leq L, \\ \sum_{h, h'} \mathbb{E} \left[W_{k j_L}^{(L+1)} W_{h' i'_L}^{(L+1)} W_{k' j'_L}^{(L+1)*} W_{h i_L}^{(L+1)*} \right] O_{k'k} O_{hh'} & \text{if } l = L+1. \end{cases} \quad (\text{B6})$$

Using the Weingarten calculus for first and second moment monomials [34], we compute independently the terms $X_l(I, I', J, J')$ appearing in the sum of Eq.(B5).

- The first term depends on the parameterized part of the first layer $W^{(1)}$.

$$X_1(I, I', J, J') = \frac{1}{d(d+1)} \left(\delta_{j_1}^{j'_1} \delta_{i'_1}^{j'_1} + \delta_{i_1}^{j_1} \delta_{i'_1}^{j'_1} \right) \quad (\text{B7})$$

- The last term depends on the parameterized part $W^{(L+1)}$ and the observable O .

$$\begin{aligned}X_{L+1}(I, I', J, J') &= \delta_{j_L}^{j'_L} \delta_{i'_L}^{j'_L} \left(\frac{d \text{Tr}(O)^2 - \|O\|^2}{d(d^2-1)} \right) + \delta_{i_L}^{j_L} \delta_{i'_L}^{j'_L} \left(\frac{d \|O\|^2 - \text{Tr}(O)^2}{d(d^2-1)} \right) \\ &= C_1 \delta_{j_L}^{j'_L} \delta_{i'_L}^{j'_L} + C_2 \delta_{i_L}^{j_L} \delta_{i'_L}^{j'_L}\end{aligned}\quad (\text{B8})$$

where $C_1 = \frac{d \text{Tr}(O)^2 - \|O\|^2}{d(d^2-1)}$ and $C_2 = \frac{d \|O\|^2 - \text{Tr}(O)^2}{d(d^2-1)}$.

- For the remaining layers, we get $\forall l \notin \{1, L+1\}$

$$X_l(I, I', J, J') = \frac{\delta_{j_l}^{j'_l} \delta_{i'_l}^{j'_l} \delta_{j_{l-1}}^{j'_{l-1}} \delta_{i'_{l-1}}^{j'_{l-1}} + \delta_{i_l}^{j_l} \delta_{i'_l}^{j'_l} \delta_{i_{l-1}}^{j_{l-1}} \delta_{i'_{l-1}}^{j'_{l-1}}}{d^2-1} - \frac{\delta_{j_l}^{j'_l} \delta_{i'_l}^{j'_l} \delta_{i_{l-1}}^{j_{l-1}} \delta_{i'_{l-1}}^{j'_{l-1}} + \delta_{i_l}^{j_l} \delta_{i'_l}^{j'_l} \delta_{j_{l-1}}^{j_{l-1}} \delta_{j'_{l-1}}^{j'_{l-1}}}{d(d^2-1)}$$

We denote by c_ω^L the Fourier coefficient if we consider the full circuit up to the L layer and we establish in what follows a recursive relation between $\mathbb{E}[|c_\omega^L|^2]$ and $\mathbb{E}[|c_\omega^{L-1}|^2]$ where c_ω^{L-1} is a Fourier coefficient of the quantum model generated by the first $L-1$ layers.

We begin by recalling that the expectation value of the modulus squared of the Fourier coefficients c_ω^L and c_ω^{L-1} can be written as

$$\mathbb{E}_{W^{(1)}, \dots, W^{(L+1)} \sim U(N)} [|c_\omega^L|^2] = \sum_{\substack{J, J' \in R_1^L(\omega) \\ I, I' \in R_1^L(\omega)}} \left(\prod_{l=1}^{L-1} X_l(J, J', I, I') \right) X_L(J, J', I, I') \left(C_1 \delta_{j_L}^{j'_L} \delta_{i_L}^{i'_L} + C_2 \delta_{i_L}^{j'_L} \delta_{j_L}^{i'_L} \right) \quad (\text{B9})$$

$$\mathbb{E}_{W^{(1)}, \dots, W^L \sim U(N)} [|c_\omega^{L-1}|^2] = \sum_{\substack{J, J' \in R_1^{L-1}(\omega) \\ I, I' \in R_1^{L-1}(\omega)}} \left(\prod_{l=1}^{L-1} X_l(J, J', I, I') \right) (C_1 \delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}} + C_2 \delta_{i_{L-1}}^{j'_{L-1}} \delta_{j_{L-1}}^{i'_{L-1}}) \quad (\text{B10})$$

Now, we exploit the recursive relation between the frequency generators up to layer $L-1$ and layer L derived in Lemma 2 and expand the term $X_L(J, J', I, I') (C_1 \delta_{j_L}^{j'_L} \delta_{i_L}^{i'_L} + C_2 \delta_{i_L}^{j'_L} \delta_{j_L}^{i'_L})$ in Eq.(B9) to make appear the expression of $\mathbb{E}_{W^{(1)}, \dots, W^{(L)} \sim U(N)} |c_\omega^{L-1}|^2$. For ease of notation, we introduce the shorthand $\tilde{X}_{L-1} := \prod_{l=1}^{L-1} X_l(J, J', I, I')$ and drop the explicit dependence on the multi-indices J, J', I, I' . Moreover, we introduce the shorthands $A_1^L(\omega)$ and $A_2^L(\omega)$ obtained as follows.

$$\mathbb{E} [|c_\omega^L|^2] = \sum_{\substack{J, J' \in R_1^L(\omega) \\ I, I' \in R_1^L(\omega)}} \tilde{X}_{L-1} \mathbb{E} \left[W_{j_L j_{L-1}}^{(L)} W_{i'_L i'_{L-1}}^{(L)} W_{j'_L j'_{L-1}}^{(L)*} W_{i_L i_{L-1}}^{(L)*} \right] (C_1 \delta_{j_L}^{j'_L} \delta_{i_L}^{i'_L} + C_2 \delta_{i_L}^{j'_L} \delta_{j_L}^{i'_L}) \quad (\text{B11})$$

$$= C_1 A_1^L(\omega) + C_2 A_2^L(\omega) \quad (\text{B12})$$

where we define

$$A_1^L(\omega) = \sum_{\substack{J, J' \in R_1^L(\omega) \\ I, I' \in R_1^L(\omega)}} \tilde{X}_{L-1} \delta_{j_L}^{j'_L} \delta_{i_L}^{i'_L} \mathbb{E} \left[W_{j_L j_{L-1}}^{(L)} W_{i'_L i'_{L-1}}^{(L)} W_{j'_L j'_{L-1}}^{(L)*} W_{i_L i_{L-1}}^{(L)*} \right] \quad (\text{B13})$$

$$A_2^L(\omega) = \sum_{\substack{J, J' \in R_1^L(\omega) \\ I, I' \in R_1^L(\omega)}} \tilde{X}_{L-1} \delta_{i_L}^{j'_L} \delta_{j_L}^{i'_L} \mathbb{E} \left[W_{j_L j_{L-1}}^{(L)} W_{i'_L i'_{L-1}}^{(L)} W_{j'_L j'_{L-1}}^{(L)*} W_{i_L i_{L-1}}^{(L)*} \right] \quad (\text{B14})$$

In what follows, we focus on deriving recursive relations for $A_1^L(\omega)$ and $A_2^L(\omega)$ separately. Hence, we start with the first step of the recursion and compute both quantities for $L = 1$.

We first show that $A_1^1(\omega) = \delta_\omega^0$.

$$A_1^1(\omega) = \sum_{\substack{j_1, j'_1 \in R_1(\omega) \\ i_1, i'_1 \in R_1(\omega)}} \mathbb{E} \left[W_{j_1 1}^{(1)} W_{i'_1 1}^{(1)} W_{j'_1 1}^{(1)*} W_{i_1 1}^{(1)*} \right] \delta_{j_1}^{j'_1} \delta_{i_1}^{i'_1} \quad (\text{B15})$$

$$= \sum_{\substack{j_1, j'_1 \in R_1(\omega) \\ i_1, i'_1 \in R_1(\omega)}} \left(\frac{\delta_{j_1}^{j'_1} \delta_{i_1}^{i'_1} + \delta_{i_1}^{j'_1} \delta_{j_1}^{i'_1}}{d(d+1)} \right) \delta_{j_1}^{j'_1} \delta_{i_1}^{i'_1} \quad (\text{B16})$$

$$= \sum_{\substack{j_1, j'_1 \in R_1(\omega) \\ i_1, i'_1 \in R_1(\omega)}} \frac{\delta_{j_1}^{j'_1} \delta_{i_1}^{i'_1}}{d(d+1)} + \sum_{\substack{j_1, j'_1 \in R_1(\omega) \\ i_1, i'_1 \in R_1(\omega)}} \frac{\delta_{j_1}^{j'_1} \delta_{i_1}^{i'_1} \delta_{i_1}^{j_1} \delta_{j_1}^{i'_1}}{d(d+1)} \quad (\text{B17})$$

$$= \sum_{\substack{j_1, j_1 \in R_1(\omega) \\ i_1, i_1 \in R_1(\omega)}} \frac{1}{d(d+1)} + \sum_{\substack{j_1, j_1 \in R_1(\omega)}} \frac{1}{d(d+1)} \quad (\text{B18})$$

$$= \frac{d^2}{d(d+1)} \delta_\omega^0 + \frac{d}{d(d+1)} \delta_\omega^0 = \delta_\omega^0 \quad (\text{B19})$$

For $A_2^1(\omega)$, we obtain

$$A_2^1(\omega) = \sum_{\substack{j_1, j'_1 \in R_1(\omega) \\ i_1, i'_1 \in R_1(\omega)}} \mathbb{E} \left[W_{j_1 1}^{(1)} W_{i'_1 1}^{(1)} W_{j'_1 1}^{(1)*} W_{i_1 1}^{(1)*} \right] \delta_{i_1}^{j_1} \delta_{i'_1}^{j'_1} = \sum_{\substack{j_1, j'_1 \in R_1(\omega) \\ i_1, i'_1 \in R_1(\omega)}} \frac{\delta_{j_1}^{j'_1} \delta_{i_1}^{i'_1} + \delta_{i_1}^{j_1} \delta_{i'_1}^{j'_1}}{d(d+1)} \delta_{i_1}^{j_1} \delta_{i'_1}^{j'_1} = \frac{\delta_\omega^0}{d+1} + \frac{|R_1(\omega)|}{d(d+1)} \quad (\text{B20})$$

The recursive relation between $A_1^L(\omega)$ and $A_1^{L-1}(\omega)$ is obtained as follows.

$$A_1^L(\omega) = \sum_{\substack{J, J' \in R_1^L(\omega) \\ I, I' \in R_1^L(\omega)}} \tilde{X}_{L-1} \delta_{JL}^{j'_L} \delta_{i'_L}^{j_L} \mathbb{E} \left[W_{j_L j_{L-1}}^{(L)} W_{i'_L i'_{L-1}}^{(L)} W_{j'_L j'_{L-1}}^{(L)*} W_{i_L i_{L-1}}^{(L)*} \right] \quad (\text{B21})$$

Now, we invoke the recursive relation between R_1^L and R_1^{L-1} established in Lemma 2.

$$\begin{aligned} A_1^L(\omega) &= \sum_{\substack{j_L, j'_L=1 \\ i_L, i'_L=1 \\ J, J' \in R_1^{L-1}(\omega-k_L) \times (j_L, j'_L) \\ I, I' \in R_1^{L-1}(\omega-k_L) \times (i_L, i'_L)}} \tilde{X}_{L-1} \delta_{j_L}^{j'_L} \delta_{i'_L}^{j_L} \left(\frac{\delta_{j_L}^{j'_L} \delta_{i'_L}^{j_L} \delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}} + \delta_{i'_L}^{j_L} \delta_{i'_L}^{j'_L} \delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}}}{d^2 - 1} - \frac{\delta_{j_L}^{j'_L} \delta_{i'_L}^{j_L} \delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}} + \delta_{i'_L}^{j_L} \delta_{i'_L}^{j'_L} \delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}}}{d(d^2 - 1)} \right) \\ &= \sum_{\substack{j_L=1 \\ i_L=1}}^d \tilde{X}_{L-1} \left(\frac{\delta_{j_{L-1}}^{j'_{L-1}} \delta_{i'_{L-1}}^{j_{L-1}} + \delta_{i'_L}^{j_L} \delta_{i'_L}^{j'_L} \delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}}}{d^2 - 1} - \frac{\delta_{j_{L-1}}^{j'_{L-1}} \delta_{i'_{L-1}}^{j_{L-1}} + \delta_{i'_L}^{j_L} \delta_{i'_L}^{j'_L} \delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}}}{d(d^2 - 1)} \right) \\ &= \sum_{j_L=1}^d \sum_{i_L=j_L} \sum_{\substack{J, J' \in R_1^{L-1}(\omega-0) \times (j_L, j_L) \\ I, I' \in R_1^{L-1}(\omega-0) \times (j_L, j_L)}} \tilde{X}_{L-1} \left(\frac{\delta_{j_{L-1}}^{j'_{L-1}} \delta_{i'_{L-1}}^{j_{L-1}} + \delta_{i'_L}^{j_L} \delta_{i'_L}^{j'_L} \delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}}}{d^2 - 1} - \frac{\delta_{j_{L-1}}^{j'_{L-1}} \delta_{i'_{L-1}}^{j_{L-1}} + \delta_{i'_L}^{j_L} \delta_{i'_L}^{j'_L} \delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}}}{d(d^2 - 1)} \right) \\ &\quad + \sum_{j_L=1}^d \sum_{i_L \neq j_L} \sum_{\substack{J, J' \in R_1^{L-1}(\omega-0) \times (j_L, j_L) \\ I, I' \in R_1^{L-1}(\omega-0) \times (i_L, i_L)}} \tilde{X}_{L-1} \left(\frac{\delta_{j_{L-1}}^{j'_{L-1}} \delta_{i'_{L-1}}^{j_{L-1}}}{d^2 - 1} - \frac{\delta_{i'_L}^{j_L} \delta_{i'_L}^{j'_L} \delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}}}{d(d^2 - 1)} \right) \\ &= \sum_{j_L=1}^d \sum_{\substack{J, J' \in R_1^{L-1}(\omega-0) \times (j_L, j_L) \\ I, I' \in R_1^{L-1}(\omega-0) \times (j_L, j_L)}} \tilde{X}_{L-1} \left(\frac{\delta_{j_{L-1}}^{j'_{L-1}} \delta_{i'_{L-1}}^{j_{L-1}} + \delta_{i'_L}^{j_L} \delta_{i'_L}^{j'_L} \delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}}}{d(d+1)} \right) \\ &\quad + \sum_{j_L=1}^d \sum_{i_L \neq j_L} \sum_{\substack{J, J' \in R_1^{L-1}(\omega-0) \times (j_L, j_L) \\ I, I' \in R_1^{L-1}(\omega-0) \times (i_L, i_L)}} \tilde{X}_{L-1} \left(\frac{\delta_{j_{L-1}}^{j'_{L-1}} \delta_{i'_{L-1}}^{j_{L-1}}}{d^2 - 1} - \frac{\delta_{i'_L}^{j_L} \delta_{i'_L}^{j'_L} \delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}}}{d(d^2 - 1)} \right) \\ &= d \sum_{\substack{J, J' \in R_1^{L-1}(\omega-0) \\ I, I' \in R_1^{L-1}(\omega-0)}} \tilde{X}_{L-1} \left(\frac{\delta_{j_{L-1}}^{j'_{L-1}} \delta_{i'_{L-1}}^{j_{L-1}} + \delta_{i'_L}^{j_L} \delta_{i'_L}^{j'_L} \delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}}}{d(d+1)} \right) + d(d-1) \sum_{\substack{J, J' \in R_1^{L-1}(\omega-0) \\ I, I' \in R_1^{L-1}(\omega-0)}} \tilde{X}_{L-1} \left(\frac{\delta_{j_{L-1}}^{j'_{L-1}} \delta_{i'_{L-1}}^{j_{L-1}}}{d^2 - 1} - \frac{\delta_{i'_L}^{j_L} \delta_{i'_L}^{j'_L} \delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}}}{d(d^2 - 1)} \right) \\ &= \frac{1}{d+1} [A_1^{L-1}(\omega) + A_2^{L-1}(\omega)] + \frac{d}{d+1} A_1^{L-1}(\omega) - \frac{1}{d+1} A_2^{L-1}(\omega) \\ &= A_1^{L-1}(\omega) \end{aligned} \quad (\text{B22})$$

We showed that $\forall \omega \in \Omega, \forall L \geq 2, A_1^L(\omega) = A_1^{L-1}(\omega)$. Moreover, we have $A_1^1(\omega) = \delta_\omega^0$. Hence, we obtain

$$\forall \omega \in \Omega, \forall L \geq 1, A_1^L(\omega) = A_1^1(\omega) = \delta_\omega^0. \quad (\text{B23})$$

This implies that apart from the frequency zero, the term which contributes to the variance expression of a Fourier coefficient is $A_2^L(\omega)$.

For the recursive relation between $A_2^L(\omega)$ and $A_2^{L-1}(\omega)$, we get the following:

$$\begin{aligned}
A_2^L(\omega) &= \sum_{\substack{J, J' \in R_1^L(\omega) \\ I, I' \in R_1^L(\omega)}} \tilde{X}_{L-1} \delta_{i_L}^{j_L} \delta_{i'_L}^{j'_L} \mathbb{E} \left[W_{j_L j_{L-1}}^{(L)} W_{i'_L i'_{L-1}}^{(L)} W_{j'_L j'_{L-1}}^{(L)*} W_{i_L i_{L-1}}^{(L)*} \right] \\
&= \sum_{\substack{i_L, i'_L=1 \\ j_L, j'_L=1 \\ I, I' \in R_1^{L-1}(\omega - k_L) \times (i_L, i'_L) \\ J, J' \in R_1^{L-1}(\omega - k_L) \times (j_L, j'_L)}} \tilde{X}_{L-1} \delta_{i_L}^{j_L} \delta_{i'_L}^{j'_L} \left(\frac{\delta_{j_L}^{j'_L} \delta_{i_L}^{i'_{L-1}} \delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}} + \delta_{i_L}^{j'_L} \delta_{i'_L}^{j_{L-1}} \delta_{i_{L-1}}^{j'_{L-1}}}{d^2 - 1} - \frac{\delta_{j_L}^{j'_L} \delta_{i_L}^{i'_{L-1}} \delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}} + \delta_{i_L}^{j'_L} \delta_{i'_L}^{j_{L-1}} \delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}}}{d(d^2 - 1)} \right) \\
&= \sum_{j_L, j'_L=1}^d \sum_{J, J' \in R_1^{L-1}(\omega - k_L) \times (j_L, j'_L)} \tilde{X}_{L-1} \left(\frac{\delta_{j_L}^{j'_L} \delta_{i_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}} + \delta_{i_{L-1}}^{j_{L-1}} \delta_{i'_{L-1}}^{j'_{L-1}}}{d^2 - 1} - \frac{\delta_{j_L}^{j'_L} \delta_{i_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}} + \delta_{i_{L-1}}^{j_{L-1}} \delta_{i'_{L-1}}^{j'_{L-1}}}{d(d^2 - 1)} \right) \\
&= \sum_{j_L=1}^d \sum_{j'_L=j_L} \sum_{J, J' \in R_1^{L-1}(\omega - k_L) \times (j_L, j'_L)} \tilde{X}_{L-1} \left(\frac{\delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}} + \delta_{i_{L-1}}^{j_{L-1}} \delta_{i'_{L-1}}^{j'_{L-1}}}{d^2 - 1} - \frac{\delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}} + \delta_{i_{L-1}}^{j_{L-1}} \delta_{i'_{L-1}}^{j'_{L-1}}}{d(d^2 - 1)} \right) \\
&\quad + \sum_{j_L=1}^d \sum_{j'_L \neq j_L} \sum_{J, J' \in R_1^{L-1}(\omega - k_L) \times (j_L, j'_L)} \tilde{X}_{L-1} \left(\frac{\delta_{i_{L-1}}^{j'_{L-1}} \delta_{i'_{L-1}}^{j_{L-1}}}{d^2 - 1} - \frac{\delta_{j_{L-1}}^{j'_{L-1}} \delta_{i'_{L-1}}^{i_{L-1}}}{d(d^2 - 1)} \right) \\
&= \sum_{j_L=1}^d \sum_{j'_L=j_L} \sum_{J, J' \in R_1^{L-1}(\omega - 0) \times (j_L, j'_L)} \tilde{X}_{L-1} \left(\frac{\delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}} + \delta_{i_{L-1}}^{j_{L-1}} \delta_{i'_{L-1}}^{j'_{L-1}}}{d(d+1)} \right) \\
&\quad + \sum_{j_L=1}^d \sum_{j'_L \neq j_L} \sum_{J, J' \in R_1^{L-1}(\omega - k_L) \times (j_L, j'_L)} \tilde{X}_{L-1} \left(\frac{\delta_{i_{L-1}}^{j'_{L-1}} \delta_{i'_{L-1}}^{j_{L-1}}}{d^2 - 1} - \frac{\delta_{j_{L-1}}^{j'_{L-1}} \delta_{i'_{L-1}}^{i_{L-1}}}{d(d^2 - 1)} \right) \\
&= d \sum_{J, J' \in R_1^{L-1}(\omega - 0)} \tilde{X}_{L-1} \left(\frac{\delta_{j_{L-1}}^{j'_{L-1}} \delta_{i_{L-1}}^{i'_{L-1}} + \delta_{i_{L-1}}^{j_{L-1}} \delta_{i'_{L-1}}^{j'_{L-1}}}{d(d+1)} \right) + \sum_{\substack{j_L, j'_L \\ j'_L \neq j_L}} \sum_{J, J' \in R_1^{L-1}(\omega - k_L)} \tilde{X}_{L-1} \left(\frac{\delta_{i_{L-1}}^{j'_{L-1}} \delta_{i'_{L-1}}^{j_{L-1}}}{d^2 - 1} - \frac{\delta_{j_{L-1}}^{j'_{L-1}} \delta_{i'_{L-1}}^{i_{L-1}}}{d(d^2 - 1)} \right) \\
&= \frac{1}{d+1} [A_1^{L-1}(\omega) + A_2^{L-1}(\omega)] + \sum_{\substack{j_L, j'_L \\ j'_L \neq j_L}} \left[\frac{1}{d^2 - 1} A_2^{L-1}(\omega - k_L) - \frac{1}{d(d^2 - 1)} A_1^{L-1}(\omega - k_L) \right] \\
&= \frac{1}{d+1} [\delta_\omega^0 + A_2^{L-1}(\omega)] + \sum_{\substack{j_L, j'_L \\ j'_L \neq j_L}} \frac{A_2^{L-1}(\omega - k_L)}{d^2 - 1} - \frac{|R_L(\omega)| - d\delta_\omega^0}{d(d^2 - 1)} \\
&= \frac{d-1}{d^2 - 1} A_2^{L-1}(\omega) + \sum_{\substack{j_L, j'_L \\ j'_L \neq j_L}} \frac{A_2^{L-1}(\omega - k_L)}{d^2 - 1} + \frac{d\delta_\omega^0}{d^2 - 1} - \frac{|R_L(\omega)|}{d(d^2 - 1)} \\
&= \sum_{j_L, j'_L} \frac{A_2^{L-1}(\omega - k_L)}{d^2 - 1} - \frac{A_2^{L-1}(\omega)}{d^2 - 1} + \frac{d\delta_\omega^0}{d^2 - 1} - \frac{|R_L(\omega)|}{d(d^2 - 1)} \\
&= \sum_{k_L \in \Omega_{distinct}^L} \frac{|R_L(k_L)| - \delta_{k_L}^0}{d^2 - 1} A_2^{L-1}(\omega - k_L) + \frac{d\delta_\omega^0}{d^2 - 1} - \frac{|R_L(\omega)|}{d(d^2 - 1)}
\end{aligned} \tag{B24}$$

Thus, the final recursion formula for $A_2^L(\omega)$ is the following:

$$\begin{aligned} \forall \omega, \forall L \geq 2, A_2^L(\omega) &= \sum_{k_L \in \Omega_{distinct}^L} \frac{|R_L(k_L)| - \delta_{k_L}^0}{d^2 - 1} A_2^{L-1}(\omega - k_L) + \frac{d\delta_\omega^0}{d^2 - 1} - \frac{|R_L(\omega)|}{d(d^2 - 1)} \\ &= \sum_{k_L \in \Omega_{distinct}^L} \alpha_{k_L} A_2^{L-1}(\omega - k_L) + \frac{d\delta_\omega^0}{d^2 - 1} - \frac{|R_L(\omega)|}{d(d^2 - 1)} \end{aligned}$$

We iterate through L to find the expression of $A_2^L(\omega)$:

$$\begin{aligned} A_2^L(\omega) &= \sum_{k_L, \dots, k_2} \alpha_{k_L} \dots \alpha_{k_2} \left[\frac{\delta_{\omega - \sum_{l=2}^L k_l}^0}{d+1} + \frac{|R_1(\omega - \sum_{l=2}^L k_l)|}{d(d+1)} \right] \\ &+ \sum_{k_L, \dots, k_3} \alpha_{k_L} \dots \alpha_{k_3} \left[\frac{d\delta_{\omega - \sum_{l=3}^L k_l}^0}{d^2 - 1} - \frac{|R_2(\omega - \sum_{l=3}^L k_l)|}{d(d^2 - 1)} \right] \\ &+ \dots + \sum_{k_L, \dots, k_h} \alpha_{k_L} \dots \alpha_{k_h} \left[\frac{d\delta_{\omega - \sum_{l=h}^L k_l}^0}{d^2 - 1} - \frac{|R_{h-1}(\omega - \sum_{l=h}^L k_l)|}{d(d^2 - 1)} \right] + \dots \\ &+ \sum_{k_L} \alpha_{k_L} \left[\frac{d\delta_{\omega - k_L}^0}{d^2 - 1} - \frac{|R_{L-1}(\omega - k_L)|}{d(d^2 - 1)} \right] + \frac{d\delta_\omega^0}{d^2 - 1} - \frac{|R_L(\omega)|}{d(d^2 - 1)} \end{aligned} \quad (\text{B25})$$

We have $\alpha_{k_l} = \frac{(|R_l(k_l)| - \delta_{k_l}^0)}{d^2 - 1} = \frac{(\beta_{k_l} - \delta_{k_l}^0)}{d^2 - 1} \simeq \frac{\beta_{k_l}}{d^2 - 1}$ so we can replace α_{k_l} with $\frac{\beta_{k_l}}{d^2 - 1}$ in the above expression since $\beta_0 \geq d$.

In this step, we use the expression of the recursive redundancy relation in Lemma 2 to simplify the generic term in the summation of Eq.(B25):

$$\begin{aligned} \frac{1}{(d^2 - 1)^{L-h+1}} \sum_{k_L, \dots, k_h} \beta_{k_L} \dots \beta_{k_h} \left[\frac{d\delta_{\omega - \sum_{l=h}^L k_l}^0}{d^2 - 1} - \frac{|R_{h-1}(\omega - \sum_{l=h}^L k_l)|}{d(d^2 - 1)} \right] \\ = \frac{d}{(d^2 - 1)^{L-h+2}} |R_h^L(\omega)| - \frac{1}{d(d^2 - 1)^{L-h+2}} |R_{h-1}^L(\omega)| \end{aligned} \quad (\text{B26})$$

Finally, we inject Eq.(B26) in Eq.(B25) and get:

$$\begin{aligned} A_2^L(\omega) &= \frac{1}{(d^2 - 1)^{L-1}} \sum_{k_L, \dots, k_2} \beta_{k_L} \dots \beta_{k_2} \left[\frac{\delta_{\omega - \sum_{l=2}^L k_l}^0}{d+1} + \frac{|R_1(\omega - \sum_{l=2}^L k_l)|}{d(d+1)} \right] \\ &+ \sum_{h=3}^L \left[\frac{d}{(d^2 - 1)^{L-h+2}} |R_h^L(\omega)| - \frac{1}{d(d^2 - 1)^{L-h+2}} |R_{h-1}^L(\omega)| \right] \\ &+ \frac{d\delta_\omega^0}{d^2 - 1} - \frac{|R_L(\omega)|}{d(d^2 - 1)} \\ &= \frac{1}{d(d+1)(d^2 - 1)^{L-1}} |R_1^L(\omega)| + \left(\frac{1}{(d+1)(d^2 - 1)^{L-1}} - \frac{1}{d(d^2 - 1)^{L-1}} \right) |R_2^L(\omega)| \\ &+ \sum_{h=3}^{L-1} \left(\frac{d}{(d^2 - 1)^{L-h+2}} - \frac{1}{d(d^2 - 1)^{L-h+1}} \right) |R_h^L(\omega)| + \left(\frac{d}{(d^2 - 1)^2} - \frac{1}{d(d^2 - 1)} \right) |R_L(\omega)| + \frac{d\delta_\omega^0}{d^2 - 1} \\ &= \frac{1}{d(d+1)(d^2 - 1)^{L-1}} |R_1^L(\omega)| - \frac{1}{d(d+1)(d^2 - 1)^{L-1}} |R_2^L(\omega)| + \sum_{h=3}^L \frac{1}{d(d^2 - 1)^{L-h+2}} |R_h^L(\omega)| + \frac{d\delta_\omega^0}{d^2 - 1} \end{aligned} \quad (\text{B27})$$

By combining Eq.(B23) and Eq.(B27), we retrieve the result from Theorem 6:

$$\mathbb{E}[|c_\omega^L|^2] = C_1 A_1^L(\omega) + C_2 A_2^L(\omega) = C_1 \delta_\omega^0 + C_2 A_2^L(\omega)$$

□

B.2 Fourier coefficients variance with approximate 2-design trainable unitaries

We present the formal version of Theorem 2.

Theorem 7 (Fourier coefficients variance with approximate 2-design trainable unitaries, Formal). *Consider a quantum model of the form in Eq.(1) and a parametrized circuit of the form in Eq.(2) with $L = 1$ layers and fixed encoding Hamiltonians resulting in a spectrum Ω . We assume that each of the trainable layers $W^l(\theta)$, $l \in \{1, 2\}$ form independently an ε -approximate 2-design. The variance of each Fourier coefficient $c_\omega(\theta)$ for the frequencies $\omega \in \Omega$ appearing in the model Fourier decomposition in Eq.(3) is upper bounded as*

$$\text{Var}_\theta[c_\omega(\theta)] \leq \text{Var}_{2\text{-design}}[c_\omega(\theta)] + Q(\varepsilon, |R(\omega)|) \quad (\text{B28})$$

with $\text{Var}_{2\text{-design}}[c_\omega]$ the variance of a Fourier coefficient under the 2-design assumption given in Theorem 1 and Q is a polynomial in ε and the frequency redundancy $|R(\omega)|$. Moreover, the polynomial Q is defined as follows for different measures of approximate 2-design

1. For diamond norm-based approximate 2-design, we have

$$\begin{aligned} Q(\varepsilon_\diamond, |R(\omega)|) &= C_1 \varepsilon_\diamond + \|O\|_1^2 \varepsilon_\diamond^2 \\ &\quad + \frac{\|O\|_\infty^2}{d(d+1)} |R(\omega)| \varepsilon_\diamond \end{aligned} \quad (\text{B29})$$

2. For monomial-based approximate 2-design, we have

$$\begin{aligned} Q(\varepsilon_M, |R(\omega)|) &= \left(\frac{C_1}{d^2} + \frac{C_2}{d(d+1)} \right) \varepsilon_M |R(\omega)| \\ &\quad + \frac{C_2}{d^2} (\varepsilon_M |R(\omega)|)^2 \end{aligned} \quad (\text{B30})$$

3. For spectral norm-based approximate 2-design, we have

$$\begin{aligned} Q(\varepsilon_\infty, |R(\omega)|) &= \left(C_1 + \frac{\|O\|_2^2}{d(d+1)} \right) \varepsilon_\infty \sqrt{|R(\omega)|} \\ &\quad + \|O\|_2^2 \varepsilon_\infty^2 |R(\omega)| \end{aligned} \quad (\text{B31})$$

Here, we use the shorthand $\varepsilon_\diamond := \|\mathcal{A}^{(2)}\|_\diamond$ for the diamond norm defined in Eq.(B38), $\varepsilon_\infty := \|O\|_\infty$ for the spectral norm and $\varepsilon_M := d^2 \max_{i,j} |\mathcal{A}^{(2)}|_{i,j}$. The constants C_1 and C_2 are defined as follows $C_1 = \frac{d\|O\|^2 - \text{Tr}(O)^2}{d(d^2-1)}$, $C_2 = \sum_{l,k} \frac{[O^{\otimes 2}]_{l,k}}{d^2}$.

Here we provide an upper bound on the variance of each Fourier coefficient in the setting of a single layered QFM when each of the two parameterized layers form an ε -approximate 2 design according to the diamond norm, spectral norm and the monomial definition introduced in Definition 5.

We note that if we assume that the parameterized layers form each a 1-design, then the expectation value of each Fourier coefficient is zero except for the null frequency as mentioned in B.1.

The main takeaway from this result is that the bound is polynomial in the frequency redundancy. This shows that the degree to which a Fourier coefficient concentrates around its mean (which is 0 under the 1-design hypothesis for non null frequencies) is restricted by its frequency in the spectrum construction, thus introducing an encoding dependent *inductive bias*. The proof of this Theorem follows a similar proof to the one given in [23] where the authors used the same techniques to compute a bound on the variance of the quantum model's gradient.

Proof. Let us first recall the Fourier coefficient expression for a single encoding layer ($L = 1$ in Eq.(A1)):

$$\begin{aligned}
c_\omega &= \sum_{j_1, j'_1 \in R(\omega)} \sum_{k, k'} W_{j'_1 0}^{(1)*} W_{j'_2 j'_1}^{(2)*} O_{k' k} W_{j_2 j_1}^{(2)} W_{j_1 0}^{(1)} \\
&= \sum_{j_1, j'_1 \in R(\omega)} \left(W^{(1)} |0\rangle \langle 0| W^{(1)\dagger} \right)_{j_1, j'_1} \left(W^{(2)\dagger} O W^{(2)} \right)_{j'_1, j_1} \\
&= \sum_{j_1, j'_1 \in R(\omega)} \text{Tr} \left[W^{(1)} |0\rangle \langle 0| W^{(1)\dagger} |j'_1\rangle \langle j_1| \right] \text{Tr} \left[W^{(2)\dagger} O W^{(2)} |j_1\rangle \langle j'_1| \right]
\end{aligned} \tag{B32}$$

Then, the expectation of the modulus squared of the coefficient c_ω is given by:

$$\mathbb{E} [|c_\omega|^2] = \sum_{\substack{i_1, i'_1 \in R(\omega) \\ j_1, j'_1 \in R(\omega)}} \text{Tr} \left[\mathbb{E}_{W^{(1)}} \left[W^{(1)\otimes 2} |00\rangle \langle 00| W^{(1)\dagger \otimes 2} |j'_1 i_1\rangle \langle j_1 i'_1| \right] \right] \text{Tr} \left[\mathbb{E}_{W^{(2)}} \left[W^{(2)\dagger \otimes 2} O^{\otimes 2} W^{(2)\otimes 2} |j_1 i'_1\rangle \langle j'_1 i_1| \right] \right] \tag{B33}$$

where we use in the second equality the property $\text{Tr}[A] \times \text{Tr}[B] = \text{Tr}[A \otimes B]$.

We recall here that in Theorem 7 we consider that both parameterized unitaries have the same ε -distance to a 2-design. However the following proof holds if each of the parameterized unitaries have different distributions with different distances to a 2-design.

Expectation terms in Eq.(B33) can be written using the superoperator $\mathcal{A}_W^{(2)}(\cdot) := \int_{\text{Haar}} d\mu(W) W^{\otimes 2}(\cdot) (W^\dagger)^{\otimes 2} - \int_W dW W^{\otimes 2}(\cdot) (W^\dagger)^{\otimes 2}$. Moreover, we use the property $\int_W dW W^{\otimes 2}(\cdot) (W^\dagger)^{\otimes 2} = \int_W dW (W^\dagger)^{\otimes 2}(\cdot) W^{\otimes 2}$ [34] and obtain:

$$\begin{aligned}
\text{Tr} \left[\mathbb{E}_{W^{(1)}} \left[W^{(1)\otimes 2} |00\rangle \langle 00| W^{(1)\dagger \otimes 2} |j'_1 i_1\rangle \langle j_1 i'_1| \right] \right] &= \mathbb{E}_{\text{Haar}} \left[\text{Tr} \left[W^{(1)\otimes 2} |00\rangle \langle 00| W^{(1)\dagger \otimes 2} |j'_1 i_1\rangle \langle j_1 i'_1| \right] \right] \\
&\quad - \text{Tr} \left[\mathcal{A}_W^{(2)}(|00\rangle \langle 00|) |j'_1 i_1\rangle \langle j_1 i'_1| \right] \\
&= \frac{\delta_{j_1, j'_1} \delta_{i_1, i'_1} (\delta_{\omega, 0}) + \delta_{i_1, j_1} \delta_{i'_1, j'_1}}{d(d+1)} - \text{Tr} \left[\mathcal{A}_W^{(2)}(|00\rangle \langle 00|) |j'_1 i_1\rangle \langle j_1 i'_1| \right] \\
\text{Tr} \left[\mathbb{E}_{W^{(2)}} \left[W^{(2)\dagger \otimes 2} O^{\otimes 2} W^{(2)\otimes 2} |j_1 i'_1\rangle \langle j'_1 i_1| \right] \right] &= \mathbb{E}_{\text{Haar}} \left[\text{Tr} \left[W^{(2)\dagger \otimes 2} O^{\otimes 2} W^{(2)\otimes 2} |j_1 i'_1\rangle \langle j'_1 i_1| \right] \right] \\
&\quad - \text{Tr} \left[\mathcal{A}_W^{(2)}(O^{\otimes 2}) |j_1, i'_1\rangle \langle j'_1, i_1| \right] \\
&= K_1 \delta_{j_1, j'_1} \delta_{i_1, i'_1} + K_2 \delta_{i_1, j_1} \delta_{i'_1, j'_1} - \text{Tr} \left[\mathcal{A}_W^{(2)}(O^{\otimes 2}) |j_1, i'_1\rangle \langle j'_1, i_1| \right] \\
&= K_1 \delta_{j_1, j'_1} \delta_{i_1, i'_1} + K_2 \delta_{i_1, j_1} \delta_{i'_1, j'_1} - \text{Tr} \left[\mathcal{A}_W^{(2)}(|j_1, i'_1\rangle \langle j'_1, i_1|) O^{\otimes 2} \right]
\end{aligned} \tag{B34}$$

The expectation with respect to the Haar measure in both terms is simply computed using the Weingarten calculus formula for the second moment monomials [34] where

$$K_1 = \frac{d \text{Tr}(O)^2 - \|O\|_2^2}{d(d^2 - 1)} \quad \text{and} \quad K_2 = \frac{d \|O\|^2 - \text{Tr}(O)^2}{d(d^2 - 1)} \tag{B35}$$

By substituting the terms of Eq.(B34) in Eq.(B33), we get the following expression:

$$\begin{aligned}
\mathbb{E} [|c_\omega|^2] &= \mathbb{E}_{\text{Haar}} [|c_\omega|^2] - \sum_{\substack{i_1, i'_1 \in R(\omega) \\ j_1, j'_1 \in R(\omega)}} \left(\frac{\delta_{j_1, j'_1} \delta_{i_1, i'_1} (\delta_{\omega, 0}) + \delta_{i_1, j_1} \delta_{i'_1, j'_1}}{d(d+1)} \right) \text{Tr} [\mathcal{A}^{\otimes 2}(|j_1, i'_1\rangle \langle j'_1, i_1|) O^{\otimes 2}] \\
&\quad - \sum_{\substack{i_1, i'_1 \in R(\omega) \\ j_1, j'_1 \in R(\omega)}} \left(K_1 \delta_{j_1, j'_1} \delta_{i_1, i'_1} (\delta_{\omega, 0}) + K_2 \delta_{i_1, j_1} \delta_{i'_1, j'_1} \right) \text{Tr} [\mathcal{A}_W^{(2)}(|00\rangle \langle 00|) |j'_1, i_1\rangle \langle j_1, i'_1|] \\
&\quad + \sum_{\substack{i_1, i'_1 \in R(\omega) \\ j_1, j'_1 \in R(\omega)}} \text{Tr} [\mathcal{A}_W^{(2)}(|00\rangle \langle 00|) |j'_1, i_1\rangle \langle j_1, i'_1|] \text{Tr} [\mathcal{A}_W^{(2)\otimes 2}(|j_1, i'_1\rangle \langle j'_1, i_1|) O^{\otimes 2}]
\end{aligned} \tag{B36}$$

$\forall \omega \in \Omega \setminus \{0\}$, we can further simplify the expression in Eq.(B36),

$$\begin{aligned}
\mathbb{E} [|c_\omega|^2] &= \mathbb{E}_{\text{Haar}} [|c_\omega|^2] - \frac{1}{d(d+1)} \text{Tr} \left[\mathcal{A}_W^{(2)} \left(\sum_{j_1, j'_1 \in R(\omega)} |j_1, j'_1\rangle \langle j'_1, j_1| \right) O^{\otimes 2} \right] \\
&\quad - K_2 \text{Tr} \left[\mathcal{A}_W^{(2)}(|00\rangle \langle 00|) \sum_{j_1, j'_1 \in R(\omega)} |j'_1, j_1\rangle \langle j_1, j'_1| \right] \\
&\quad + \text{Tr} \left[\left(\mathcal{A}_W^{(2)}(|00\rangle \langle 00|) \otimes \mathcal{A}_W^{(2)}(O^{\otimes 2}) \right) \left(\sum_{\substack{i_1, i'_1 \in R(\omega) \\ j_1, j'_1 \in R(\omega)}} |j'_1, i_1\rangle \langle j_1, i'_1| \otimes |j_1, i'_1\rangle \langle j'_1, i_1| \right) \right]
\end{aligned} \tag{B37}$$

Different norms of $\mathcal{A}_W^{(2)}(X^{\otimes 2})$ can be used to further obtain an upper bound on $\mathbb{E} [|c_\omega|^2]$. The norm choice leading to the tightest upper bound is highly dependent on the observable norm and the frequency redundancy. We list here the different operator norms and observable norms leading to similar upper bounds on the Fourier coefficients variance but with different scaling depending on the setting.

All the operators norms used in the following proof are Schatten p-norms. For completeness, we recall that the Schatten p-norm of an operator Ω is defined as follows: $\|\Omega\|_p = [\text{Tr}(|\Omega|^p)]^{1/p} = [\sum_i \sigma_i^p]^{1/p}$ where $|\Omega| = \sqrt{\Omega^\dagger \Omega}$ and σ_i s are the singular values of Ω . We also recall that the diamond norm of a Hermiticity preserving linear map $\mathcal{S}_A$ is defined as

$$\|\mathcal{S}_A\|_\diamond = \sup_n \sum_{\Omega_{AB} \neq 0} \frac{\|(\mathcal{S}_A \otimes \mathcal{I}_B^{(n)})(\Omega_{AB})\|}{\|\Omega_{AB}\|_1} \tag{B38}$$

where $\Omega_{AB} \in \mathcal{L}(\mathcal{H}_A \otimes \mathcal{H}_B)$. The diamond norm distance $\|\mathcal{M} - \mathcal{N}\|_\diamond$ is a measure of the distinguishability of two quantum operations $\mathcal{M}$ and $\mathcal{N}$.

Diamond norm bound: We start by applying the triangular inequality on all terms in Eq.(B37) and obtain:

$$\begin{aligned}
\mathbb{E} [|c_\omega|^2] &\leq \mathbb{E}_{\text{Haar}} [|c_\omega|^2] + \frac{\|O^{\otimes 2}\|_\infty}{d(d+1)} \left\| \sum_{j_1, j'_1 \in R(\omega)} |j_1, j'_1\rangle \langle j'_1, j_1| \right\|_1 \varepsilon_\diamond \\
&\quad + K_2 \left\| \sum_{j_1, j'_1 \in R(\omega)} |j'_1, j_1\rangle \langle j_1, j'_1| \right\|_\infty \varepsilon_\diamond \\
&\quad + \|\mathcal{A}_W^{(2)}(|00\rangle\langle 00|) \otimes \mathcal{A}_W^{(2)}(O^{\otimes 2})\|_1 \left\| \sum_{\substack{i_1, i'_1 \in R(\omega) \\ j_1, j'_1 \in R(\omega)}} |j'_1, i_1\rangle \langle j_1, i'_1| \otimes |j_1, i'_1\rangle \langle j'_1, i_1| \right\|_\infty \\
&\leq \mathbb{E}_{\text{Haar}} [|c_\omega|^2] + \frac{\|O\|_\infty^2}{d(d+1)} |R(\omega)| \varepsilon_\diamond + K_2 \varepsilon_\diamond + \|O\|_1^2 \varepsilon_\diamond^2
\end{aligned} \tag{B39}$$

where in the first inequality we used hölder's inequality $|Tr[AB]| \leq \|A\|_1 \|B\|_\infty$ and the property $\|\mathcal{A}_W^{(2)}(X)\|_1 \leq \|\mathcal{A}_W^{(2)}\|_\diamond \|X\|_1$ with the shorthand $\varepsilon_\diamond := \|\mathcal{A}_W^{(2)}\|_\diamond$ and in the second inequality we used the property $\|A \otimes B\| = \|A\| \|B\|$. We also used the following elementary calculations to show that the singular values of matrices of the form $\sum_{j_1, j'_1 \in R(\omega)} |j_1, j'_1\rangle \langle j'_1, j_1|$ are zeros and ones with 1s having multiplicity $|R(\omega)|$:

$$\left(\sum_{j_1, j'_1 \in R(\omega)} |j_1, j'_1\rangle \langle j'_1, j_1| \right)^\dagger \left(\sum_{i_1, i'_1 \in R(\omega)} |i_1, i'_1\rangle \langle i'_1, i_1| \right) = \sum_{\substack{i_1, i'_1 \in R(\omega) \\ j_1, j'_1 \in R(\omega)}} |j'_1, j_1\rangle \langle j_1, j'_1| |i_1, i'_1\rangle \langle i'_1, i_1| \tag{B40}$$

$$= \sum_{j_1, j'_1 \in R(\omega)} |j'_1, j_1\rangle \langle j'_1, j_1| \tag{B41}$$

We note that the bound on the last term can be slightly modified and its tightness will depend on the observable norm and the frequency redundancy. Specifically, we can obtain:

$$\begin{aligned}
\mathbb{E} [|c_\omega|^2] &\leq \mathbb{E}_{\text{Haar}} [|c_\omega|^2] + \frac{\|O\|_\infty^2}{d(d+1)} |R(\omega)| \varepsilon_\diamond + K_2 \varepsilon_\diamond + \|O\|_1^2 \varepsilon_\diamond^2 \\
&\leq \mathbb{E}_{\text{Haar}} [|c_\omega|^2] + \frac{\|O\|_\infty^2}{d^2} |R(\omega)| \varepsilon_\diamond + \frac{\|O\|_2^2}{d^2} \varepsilon_\diamond + \|O\|_1^2 \varepsilon_\diamond^2
\end{aligned} \tag{B42}$$

Spectral norm bound:

$$\begin{aligned}
\mathbb{E} [|c_\omega|^2] &\leq \mathbb{E}_{\text{Haar}} [|c_\omega|^2] + \frac{\|O^{\otimes 2}\|_2}{d(d+1)} \left\| \sum_{j_1, j'_1 \in R(\omega)} |j_1, j'_1\rangle \langle j'_1, j_1| \right\|_2 \varepsilon_\infty \\
&\quad + K_2 \left\| \sum_{j_1, j'_1 \in R(\omega)} |j'_1, j_1\rangle \langle j_1, j'_1| \right\|_2 \varepsilon_\infty \\
&\quad + \left\| \mathcal{A}_W^{(2)}(|00\rangle\langle 00|) \otimes \mathcal{A}_W^{(2)}(O^{\otimes 2}) \right\|_2 \left\| \sum_{\substack{i_1, i'_1 \in R(\omega) \\ j_1, j'_1 \in R(\omega)}} |j'_1, i_1\rangle \langle j_1, i'_1| \otimes |j_1, i'_1\rangle \langle j'_1, i_1| \right\|_2 \\
&\leq \mathbb{E}_{\text{Haar}} [|c_\omega|^2] + \frac{\|O\|_2^2}{d(d+1)} \sqrt{|R(\omega)|} \varepsilon_\infty + K_2 \sqrt{|R(\omega)|} \varepsilon_\infty + \|O\|_2^2 |R(\omega)| \varepsilon_\infty^2
\end{aligned} \tag{B43}$$

where in the first inequality we use $|Tr[AB]| \leq \|A\|_2 \|B\|_2$ and the property $\|\mathcal{A}_W^{(2)}(X)\|_2 \leq \|\mathcal{A}_W^{(2)}\|_\infty \|X\|_2$ with the shorthand $\varepsilon_\infty := \|\mathcal{A}_W^{(2)}\|_\infty$. The above upper bound can be further simplified.

Precisely, we obtain

$$\text{Var}_\theta[c_\omega(\theta)] \leq \text{Var}_{\text{Haar}}[c_\omega(\theta)] + 2 \frac{\|O\|_2^2}{d} \sqrt{|\tilde{R}(\omega)|\varepsilon_\infty} + d^2 \|O\|_2^2 |\tilde{R}(\omega)|\varepsilon_\infty^2 \quad (\text{B44})$$

where we introduced the normalized frequency redundancy $|\tilde{R}(\omega)| := |R(\omega)|/d^2$ and used the inequality $K_2 \leq \|O\|^2/d^2$.

Monomial norm bound:

Definition 5. [Monomial definition of ε -approximate 2-design] An ansatz $U(\Theta)$ forms a monomial ε -approximate 2-design if:

$$\max_{p,q,r,s \in [d]} |(\mathcal{A}_{U(\Theta)}^{(2)})_{p,q,r,s}| \leq \frac{\varepsilon}{d^2} \quad (\text{B45})$$

where $(\mathcal{A}_{U(\Theta)}^{(2)})_{p,q,r,s}$ is a coefficient of the d^4 -dimensional matrix $\mathcal{A}_{U(\Theta)}^{(2)}$.

The monomial definition 5 of approximate 2-design gives an upper bound on the superoperator coefficients in the computational basis. Hence, we take

$$\varepsilon = \max_{\substack{i_1, i'_1, j_1, j'_1 \in [d] \\ k, k' \in [d^4]}} d^2 \left| \mathcal{A}_{\mathbb{W}}^{(2)}(|j_1, i'_1\rangle \langle j'_1, i_1|)_{k, k'} \right|$$

Consequently, we get

$$|\text{Tr} [\mathcal{A}_{\mathbb{W}}^{(2)}(|00\rangle \langle 00|) |j'_1, i_1\rangle \langle j_1, i'_1|]| \leq \frac{\varepsilon}{d^2}$$

and

$$\text{Tr} [\mathcal{A}_{\mathbb{W}}^{(2)}(|j'_1, i_1\rangle \langle j_1, i'_1|) O^{\otimes 2}] = \sum_{k, k' \in [d]} \left| \left(\mathcal{A}_{\mathbb{W}}^{(2)} |j_1, i'_1\rangle \langle j'_1, i_1| \right)_{k, k'} \right| |O_{kk'}^{\otimes 2}| \leq \frac{\varepsilon}{d^2} \sum_{k, k' \in [d]} |O_{kk'}^{\otimes 2}|$$

By applying the triangular inequality in Eq.(B36), one arrives to the upper bound obtained in Theorem 7:

$$\mathbb{E} [|c_\omega|^2] \leq \mathbb{E}_{\text{Haar}} [|c_\omega|^2] + \frac{|R(\omega)|}{d(d+1)} \varepsilon \sum_{kk'} \frac{|[O^{\otimes 2}]_{kk'}|}{d^2} + K_2 \frac{|R(\omega)|}{d^2} \varepsilon + \frac{|R(\omega)|^2 \varepsilon^2}{d^2} \sum_{kk'} \frac{|[O^{\otimes 2}]_{kk'}|}{d^2} \quad (\text{B46})$$

The term $\sum_{kk'} [O^{\otimes 2}]_{kk'}$ can be related to the observable 2-norm as follows. Indeed, by definition the observable Schatten 2-norm can be written as

$$\|O\|_2^2 = \text{Tr}[O^\dagger O] = \sum_{k, k'=1}^d |O_{k, k'}|^2. \quad (\text{B47})$$

Hence, we obtain

$$\sum_{k, k'=1}^{d^2} |[O^{\otimes 2}]_{kk'}| \leq \sqrt{d^4 \sum_{k, k'=1}^{d^2} |[O^{\otimes 2}]_{kk'}|^2} \quad (\text{B48})$$

$$= d^2 \|O^{\otimes 2}\|_2 \quad (\text{B49})$$

$$= d^2 \|O\|_2^2 \quad (\text{B50})$$

Consequently, Eq.(B46) becomes

$$\begin{aligned} \text{Var}_\theta[c_\omega(\theta)] &\leq \text{Var}_{\text{Haar}}[c_\omega(\theta)] + (\|O\|_2^2 + \frac{\|O\|_2^2}{d^2}) |\tilde{R}(\omega)|\varepsilon + d^2 \|O\|_2^2 |\tilde{R}(\omega)|^2 \varepsilon^2 \\ &\leq \text{Var}_{\text{Haar}}[c_\omega(\theta)] + 2 \|O\|_2^2 |\tilde{R}(\omega)|\varepsilon + d^2 \|O\|_2^2 |\tilde{R}(\omega)|^2 \varepsilon^2 \end{aligned} \quad (\text{B51})$$

□

B.3 Model variance with approximate 2-design trainable unitaries

We present a formal version of Corollary 2 in the following Corollary.

Corollary 3. *Consider a quantum model $f(x, \theta)$ of the form in Eq.(1) and a parametrized circuit of the form in Eq.(2) with $L = 1$ layers and fixed encoding Hamiltonians. We assume that each of the trainable layer $W^l(\theta)$, $l \in \{1, 2\}$ form independently an ε_M -approximate 2-design according to the monomial definition introduced in Definition 5. For a fixed $x \in \mathcal{X}$, the variance of the model $f(x, \theta)$ is upper bounded for any $x \in \mathbb{R}$ as*

$$\text{Var}_\theta[f(x, \theta)] \leq \frac{\|O\|_2^2}{d^2} + \|O\|_2^2 \varepsilon. \quad (\text{B52})$$

Proof. We recall that $f(x, \theta)$ for a fixed data point $x \in \mathbb{R}$ is a real valued function. Hence, $\forall x \in \mathbb{R}$, its variance is given by $\text{Var}_\theta[f(x, \theta)] = \mathbb{E}_\theta[f^2(x, \theta)] - \mathbb{E}_\theta[f(x, \theta)]^2$.

We also recall that the model f is given by

$$f(x, \theta) = \text{Tr}[W^{(1)\dagger} S^\dagger(x) W^{(2)\dagger} O W^{(2)} S(x) W^{(1)} |0\rangle\langle 0|] \quad (\text{B53})$$

where the dependence on trainable parameters θ is hidden in the trainable unitaries $W^{(1)}$ and $W^{(2)}$.

Hence the model's second moment w.r.t to the distributions over $W^{(1)}$ and $W^{(2)}$ can be expressed as

$$\mathbb{E}_\theta[f^2(x, \theta)] := \mathbb{E}_{W^{(1)} \sim \mathbb{W}, W^{(2)} \sim \mathbb{W}}[f^2(x)] \quad (\text{B54})$$

$$= \mathbb{E}_{W^{(1)} \sim \mathbb{W}, W^{(2)} \sim \mathbb{W}}[f^2(x)] - \mathbb{E}_{W^{(1)} \sim \text{Haar}, W^{(2)} \sim \mathbb{W}}[f^2(x)] + \mathbb{E}_{W^{(1)} \sim \text{Haar}, W^{(2)} \sim \mathbb{W}}[f^2(x)] \quad (\text{B55})$$

$$= \mathbb{E}_{W^{(2)} \sim \mathbb{W}} \left[\mathbb{E}_{W^{(1)} \sim \mathbb{W}}[f^2(x)] - \mathbb{E}_{W^{(1)} \sim \text{Haar}}[f^2(x)] \right] + \mathbb{E}_{W^{(1)} \sim \text{Haar}, W^{(2)} \sim \mathbb{W}}[f^2(x)] \quad (\text{B56})$$

Using the invariance property of the Haar measure, one can easily show that

$$(\mathbb{E}_{W^{(1)} \sim \text{Haar}, W^{(2)} \sim \text{Haar}}[f(x)])^2 = (\mathbb{E}_{W^{(1)} \sim \text{Haar}}[\text{Tr}[O W^{(1)} |0\rangle\langle 0| W^{(1)\dagger}]])^2 \quad (\text{B57})$$

Similarly, we have

$$\mathbb{E}_{W^{(1)} \sim \text{Haar}, W^{(2)} \sim \mathbb{W}}[f^2(x)] = \mathbb{E}_{W^{(1)} \sim \text{Haar}}[\text{Tr}[O W^{(1)} |0\rangle\langle 0| W^{(1)\dagger}]^2] \quad (\text{B58})$$

Hence, the model's variance becomes

$$\text{Var}[f(x)] = \mathbb{E}_{W^{(2)} \sim \mathbb{W}} \left[\mathbb{E}_{W^{(1)} \sim \mathbb{W}}[f^2(x)] - \mathbb{E}_{W^{(1)} \sim \text{Haar}}[f^2(x)] \right] + \text{Var}_{W^{(1)} \sim \text{Haar}}[\text{Tr}[O W^{(1)} |0\rangle\langle 0| W^{(1)\dagger}]] \quad (\text{B59})$$

$$= \mathbb{E}_{W^{(2)} \sim \mathbb{W}} \left[\text{Tr} \left[S(x)^{\dagger \otimes 2} W^{(2)\dagger \otimes 2} O^{\otimes 2} W^{(2) \otimes 2} S(x)^{\otimes 2} (\mathbb{E}_{W^{(1)} \sim \mathbb{W}} - \mathbb{E}_{W^{(1)} \sim \text{Haar}})[W^{(1) \otimes 2} |00\rangle\langle 00| W^{(1)\dagger \otimes 2}] \right] \right] \quad (\text{B60})$$

$$+ \text{Var}_{W^{(1)} \sim \text{Haar}}[\text{Tr}[O W^{(1)} |0\rangle\langle 0| W^{(1)\dagger}]] \quad (\text{B61})$$

$$= \mathbb{E}_{W^{(2)} \sim \mathbb{W}} \left[\text{Tr} \left[S(x)^{\dagger \otimes 2} W^{(2)\dagger \otimes 2} O^{\otimes 2} W^{(2) \otimes 2} S(x)^{\otimes 2} \mathcal{A}_{\mathbb{W}}(|00\rangle\langle 00|) \right] \right] \quad (\text{B62})$$

$$+ \text{Var}_{W^{(1)} \sim \text{Haar}}[\text{Tr}[O W^{(1)} |0\rangle\langle 0| W^{(1)\dagger}]] \quad (\text{B63})$$

In what follows we focus on bounding the first term in the variance expression above. Using Holder's

inequality, we obtain

$$\text{Tr} \left[S(x)^{\dagger \otimes 2} W^{(2) \dagger \otimes 2} O^{\otimes 2} W^{(2) \otimes 2} S(x)^{\otimes 2} \mathcal{A}_W(|00\rangle\langle 00|) \right] \leq \left\| S(x)^{\dagger \otimes 2} W^{(2) \dagger \otimes 2} O^{\otimes 2} W^{(2) \otimes 2} S(x)^{\otimes 2} \right\|_2 \left\| \mathcal{A}_W(|00\rangle\langle 00|) \right\|_2 \quad (\text{B64})$$

$$= \left\| O^{\otimes 2} \right\|_2 \left\| \mathcal{A}_W(|00\rangle\langle 00|) \right\|_2 \quad (\text{B65})$$

$$= \|O\|_2^2 \sqrt{\sum_{k,k'=1}^{d^2} |\mathcal{A}_W(|00\rangle\langle 00|)|_{k,k'}^2} \quad (\text{B66})$$

$$\leq \|O\|_2^2 \sqrt{\sum_{k,k'=1}^{d^2} \left(\frac{\varepsilon}{d^2} \right)^2} \quad (\text{B67})$$

$$= \|O\|_2^2 \varepsilon \quad (\text{B68})$$

Now, we go back to the term $\text{Var}_{W^{(1)} \sim \text{Haar}} [\text{Tr}[OW^{(1)} |0\rangle\langle 0| W^{(1)\dagger}]]$ which can be directly computed using Weingarten calculus [34].

$$\text{Var}_{W^{(1)} \sim \text{Haar}} [\text{Tr}[OW^{(1)} |0\rangle\langle 0| W^{(1)\dagger}]] = \mathbb{E}_{W^{(1)} \sim \text{Haar}} [\text{Tr}[OW^{(1)} |0\rangle\langle 0| W^{(1)\dagger}]^2] - \mathbb{E}_{W^{(1)} \sim \text{Haar}} [\text{Tr}[OW^{(1)} |0\rangle\langle 0| W^{(1)\dagger}]]^2 \quad (\text{B69})$$

$$= \frac{\text{Tr}[O^2] + \text{Tr}[O]^2}{d(d+1)} - \frac{\text{Tr}[O]^2}{d^2} \quad (\text{B70})$$

$$\leq \frac{\|O\|_2^2}{d^2} \quad (\text{B71})$$

which concludes the proof. $\square$

B.4 Variance under the local 2-design setting

B.4.1 Preliminaries

In this section, we present properties of Haar integration over Haar random unitaries that we will use in the main proof of Theorem 3.

Lemma 1. *Let $\mathcal{H} = \mathcal{H}_w \otimes \mathcal{H}_{\bar{w}}$ be a bipartite Hilbert space and let $\{W(\theta)\}_{\theta \in \Theta}$ be a parameterized unitary acting on $\mathcal{H}_w$ and forming a 2-design when considering a uniform distribution over the variational parameters Θ . Then, for any arbitrary linear operators $A, B, C, D : \mathcal{H} \rightarrow \mathcal{H}$ such that $A = A_w \otimes A_{\bar{w}}$ and $C = C_w \otimes C_{\bar{w}}$, we have:*

$$\begin{aligned} & \int_{\Theta} \text{Tr}_w[A(W \otimes \mathbb{1}_{\bar{w}})B(W^\dagger \otimes \mathbb{1}_{\bar{w}})] \text{Tr}_w[C(W \otimes \mathbb{1}_{\bar{w}})D(W^\dagger \otimes \mathbb{1}_{\bar{w}})] dW(\Theta) \\ &= \int_W \text{Tr}_w[A(W \otimes \mathbb{1}_{\bar{w}})B(W^\dagger \otimes \mathbb{1}_{\bar{w}})] \text{Tr}_w[C(W \otimes \mathbb{1}_{\bar{w}})D(W^\dagger \otimes \mathbb{1}_{\bar{w}})] dW \\ &= \frac{\text{Tr}_w[B] \text{Tr}_w[D]}{2^{2m} - 1} A_{\bar{w}} C_{\bar{w}} \left(\text{Tr}[A_w] \text{Tr}[C_w] - \frac{\text{Tr}[A_w C_w]}{2^m} \right) + \frac{\text{Tr}_w[B(C_{\bar{w}} \otimes \mathbb{1}_w)D]}{2^{2m} - 1} \left(\text{Tr}[A_w C_w] - \frac{\text{Tr}[A_w] \text{Tr}[C_w]}{2^m} \right) \end{aligned} \quad (\text{B72})$$

Lemma 2. *Let $\mathcal{H} = \mathcal{H}_w \otimes \mathcal{H}_{\bar{w}}$ be a bipartite Hilbert space. Then, for any arbitrary linear operators $A_w, B_w : \mathcal{H}_w \rightarrow \mathcal{H}_w, S : \mathcal{H} \rightarrow \mathcal{H}$, we have:*

$$\text{Tr}_{\bar{w}}[|p\rangle\langle q| (A_w \otimes \mathbb{1}_{\bar{w}}) S (B_w \otimes \mathbb{1}_{\bar{w}})] = |p\rangle\langle q|_w A_w \text{Tr}_{\bar{w}}[(\mathbb{1}_w \otimes |p\rangle\langle q|_{\bar{w}}) S] B_w$$

Lemma 3. Let $\mathcal{H}$ be a composite Hilbert space such that $\mathcal{H} = \mathcal{H}_{\tau_a} \otimes \mathcal{H}_{\tau_b} \otimes \mathcal{H}_z \otimes \mathcal{H}_y$ and $\mathcal{H}_x$ a subspace of $\mathcal{H}_{\tau_a} \otimes \mathcal{H}_{\tau_b} \otimes \mathcal{H}_z$. Then, for any unitary X acting on $\mathcal{H}$, W_τ acting non trivially on subspace $\mathcal{H}_\tau = \mathcal{H}_{\tau_a} \otimes \mathcal{H}_{\tau_b}$ and $Y_i = (W_\tau^\dagger \otimes \mathbb{1}_{\bar{\tau}})X_i(W_\tau \otimes \mathbb{1}_{\bar{\tau}})\forall i \in \{1, 2\}$, we get:

$$E_\tau[Tr[Tr_x[Y_1]Tr_x[Y_2]]] = \begin{cases} \frac{2^{m/2}}{2^m+1}Tr[Tr_{x,\tau}[X_1]Tr_{x,\tau}[X_2] + Tr_{x\setminus\tau}[X_1]Tr_{x\setminus\tau}[X_2]], & \text{if } x \cap \tau \in \{\tau_a, \tau_b\} \\ Tr[Tr_x[X_1]Tr_x[X_2]], & \text{otherwise} \end{cases} \quad (\text{B73})$$

Proof. For ease of notation, we denote the Hilbert space by its associated subsystem index.

$$Tr_x[Y] = \begin{cases} Tr_\tau[(W_\tau^\dagger \otimes \mathbb{1}_y)X(W_\tau \otimes \mathbb{1}_y)], & \text{if } x = \tau \\ (W_\tau^\dagger \otimes \mathbb{1}_y)Tr_x[X](W_\tau \otimes \mathbb{1}_y), & \text{if } x \cap \tau = \emptyset \\ Tr_\tau[(W_\tau^\dagger \otimes \mathbb{1}_y)Tr_{x\setminus\tau}[X](W_\tau \otimes \mathbb{1}_y)], & \text{if } \tau \subset x \\ Tr_{\tau_a}[(W_\tau^\dagger \otimes \mathbb{1}_y)Tr_{x\setminus\tau}[X](W_\tau \otimes \mathbb{1}_y)], & \text{if } x \cap \tau = \tau_a \\ Tr_{\tau_b}[(W_\tau^\dagger \otimes \mathbb{1}_y)Tr_{x\setminus\tau}[X](W_\tau \otimes \mathbb{1}_y)], & \text{if } x \cap \tau = \tau_b \end{cases}$$

For the case where $x \cap \tau = \emptyset$,

$$\begin{aligned} E_\tau[Tr[Tr_x[Y_1]Tr_x[Y_2]]] &= E_\tau[Tr[(W_\tau^\dagger \otimes \mathbb{1}_y)Tr_x[X_1]Tr_x[X_2](W_\tau \otimes \mathbb{1}_y)]] \\ &= E_\tau[Tr[Tr_x[X_1]Tr_x[X_2]]] \\ &= Tr[Tr_x[X_1]Tr_x[X_2]] \end{aligned}$$

In both cases $x = \tau$ and $\tau \subset x$, we get $Tr_x[Y] = Tr_\tau[(W_\tau^\dagger \otimes \mathbb{1}_y)Tr_z[X](W_\tau \otimes \mathbb{1}_y)]$ where $z = \emptyset, x\setminus\tau$ respectively. Hence, to compute $E_\tau[Tr[Tr_x[Y_1]Tr_x[Y_2]]]$, we apply 2 for $\omega = \tau, A = C = \mathbb{1}, B = Tr_z[X_1]$ and $D = Tr_z[X_2]$:

$$\begin{aligned} E_\tau[Tr[Tr_x[Y_1]Tr_x[Y_2]]] &= Tr[Tr_{\tau,z}[X_1]Tr_{\tau,z}[X_2]] \\ &= Tr[Tr_x[X_1]Tr_x[X_2]] \end{aligned}$$

For the last 2 cases where the intersection between x and τ is a singleton (i.e. $x \cap \tau \in \{\tau_a, \tau_b\}$), we use the *Random Tensor Network Integrator* (RTNI) package [36] and get the following:

$$E_\tau[Tr[Tr_x[Y_1]Tr_x[Y_2]]] = \frac{2^{m/2}}{2^m+1} [Tr[Tr_{x,\tau}[X_1]Tr_{x,\tau}[X_2]] + Tr[Tr_{x\setminus\tau}[X_1]Tr_{x\setminus\tau}[X_2]]]$$

□

B.4.2 Integrate over a brickwise circuit of local 2-design blocks

In this section, we provide a general method to compute the second moment superoperator of a brickwise circuit made of parameterized local 2-design blocks (no encoding part is evolved). The methods used and the expression of the second moment superoperator will be useful to prove Theorem 3 and to give a better intuition about the monomial ε -distance to a 2-design as defined in Definition 5, for this type of circuits.

We therefore consider a circuit with a brickwise structure as described in Fig.4 acting on $n = m \times A$ qubits with depth L , A blocks per layer and circular connectivity. Each block acts on m qubits and forms a 2-design on the corresponding m -qubit subsystem. We denote by $U(\theta)$ the parameterized unitary generated by this circuit and by $U_{j,l}(\theta)$ the unitary corresponding to the j^{th} block in the l^{th} layer.

For layers with an odd index l , we denote the set of subsystems on which its blocks act non-trivially by $\{s_h\}_{h=1}^A$ and for even layers, we use the notation $\{\tau_h\}_{h=1}^A$ for the corresponding subsystems. Using this notation, we can decompose each subsystem s_h into $s_h = (s_h)_1 \cup (s_h)_2$ and similarly each subsystem τ_h into $\tau_h = (\tau_h)_1 \cup (\tau_h)_2$ where each subsystem $(s_h)_i, (\tau_h)_i$ acts on $m/2$ qubits.

We note that in the following proof, we will consider that each layer l contains A_l blocks in order to account for brickwise circuits with slightly different configurations (mainly the brickwise light cone). However, to compute the second moment superoperator, we simply consider that $A_l = A \forall l \in [L]$.

We recall here the expression of a single coefficient of the second moment superoperator corresponding to the circuit unitary $U(\theta)$ indexed by multi-indices $\mathbf{qp} := (q_1, q_2, p_1, p_2) \in [2^{4n}]$ and $\mathbf{rs} := (r_1, r_2, s_1, s_2) \in [2^{4n}]$:

$$\begin{aligned} (M_{U(\theta)}^{(2)})_{\mathbf{qp}, \mathbf{rs}} &= \int_{\theta} dU(\theta) (U(\theta)^{\otimes 2} \otimes (U(\theta)^*)^{\otimes 2})_{\mathbf{qp}, \mathbf{rs}} \\ &:= \int_{\theta} U_{q_1, r_1} U_{q_2, r_2} U_{p_1, s_1}^* U_{p_2, s_2}^* dU(\theta) \end{aligned} \quad (\text{B74})$$

For convenience, we rewrite the second moment superoperator as follows:

$$\begin{aligned} (M_{U(\theta)}^{(2)})_{\mathbf{p}, \mathbf{q}, \mathbf{r}, \mathbf{s}} &= \int_{\theta} U_{q_1, r_1} U_{q_2, r_2} U_{p_1, s_1}^* U_{p_2, s_2}^* dU(\theta) \\ &= \int_{\mathbb{U}} Tr[[p_1] \langle q_1 | U Y_1 U^\dagger] Tr[[p_2] \langle q_2 | U Y_2 U^\dagger] dU \\ &:= E_{\mathbb{U}}[Tr[[p_1] \langle q_1 | U Y_1 U^\dagger] Tr[[p_2] \langle q_2 | U Y_2 U^\dagger]] \end{aligned} \quad (\text{B75})$$

where $Y_1 = |r_1\rangle\langle s_1|$ and $Y_2 = |r_2\rangle\langle s_2|$.

In order to compute the above expectation value, one needs to integrate over all the local 2-design blocks forming the circuit unitary $U(\theta)$. Under the assumption that the blocks are independent, we can perform the integration iteratively. To do so, we first start by integrating over the blocks of the first layer and then we integrate over the remaining blocks (blocks from the 2^{nd} up to the L^{th} layer).

Considering that the h^{th} block of the first layer acts non trivially on subsystem s_h (with the corresponding unitary matrix $U_{h,1}(\theta)$), we introduce the following notations where we drop θ for simplicity:

$$X_h^{pq} := Tr_{h^+}[(\mathbb{1}_{h^-} \otimes |p\rangle\langle q|_{h^+}) \tilde{U}^{(h,1)}] \quad \forall h \in \{1, \dots, A_1\} \quad (\text{B76})$$

$$f_h(x) := Tr[Tr_x[X_h^{p_1 q_1}] Tr_x[X_h^{p_2 q_2}]] \quad \forall h < A_1, \forall x \subseteq h^- \quad (\text{B77})$$

where $h^+ := \{s_{h+1} \cup s_{h+2} \cup \dots, s_{A_1}\}$, $h^- := \{s_1 \cup s_2 \cup \dots, s_{h-1} \cup s_h\}$, $\tilde{U}^{(h,1)} := U^{(h,1)} Y U^{(h,1)\dagger}$ and $U^{(h,1)}$ is the circuit unitary U after removing the first h blocks from the first layer. We notice here that X_h^{pq} acts on subsystem h^- and that it does no longer depend on the first h blocks of the first layer. We note that $X_0^{pq} := Tr[|p\rangle\langle q| U Y U^\dagger]$ and $X_A^{pq} := \tilde{U}^{(A_1,1)}$.

We give the following recursive relation between X_h^{pq} and X_{h+1}^{pq} using Lemma 2 and for ease of notation, we shortly use $h+1$ to denote subsystem s_{h+1} :

$$\begin{aligned} X_h^{pq} &= Tr_{h+1}[Tr_{(h+1)^+}[(\mathbb{1}_{h^-} \otimes |p\rangle\langle q|_{h^+})(\mathbb{1}_{h^-} \otimes U_{h+1,1} \otimes \mathbb{1}_{(h+1)^+}) \tilde{U}^{(h+1)} (\mathbb{1}_{h^-} \otimes U_{h+1,1}^\dagger \otimes \mathbb{1}_{(h+1)^+})]] \\ &= Tr_{h+1}[(\mathbb{1}_{h^-} \otimes |p\rangle\langle q|_{h+1})(\mathbb{1}_{h^-} \otimes U_{h+1,1}) X_{h+1}^{pq} (\mathbb{1}_{h^-} \otimes U_{h+1,1}^\dagger)] \end{aligned}$$

We then apply Lemma 1 and get the following $\forall h < A_1, \forall x \subseteq h^-$,

$$\begin{aligned} E_{U_{h+1,1}}[f_h(x)] &= E_{U_{h+1,1}}[Tr[Tr_x[X_h^{p_1 q_1}] Tr_x[X_h^{p_2 q_2}]]] \\ &= \alpha_{h+1} Tr[Tr_{h+1}[Tr_x[X_{h+1}^{p_1 q_1}]] Tr_{h+1}[Tr_x[X_{h+1}^{p_2 q_2}]]] + \beta_{h+1} Tr[Tr_x[X_{h+1}^{p_1 q_1}] Tr_x[X_{h+1}^{p_2 q_2}]] \\ &= \alpha_{h+1} f_{h+1}(x \cup (h+1)) + \beta_{h+1} f_{h+1}(x) \end{aligned} \quad (\text{B78})$$

with the coefficients α_{h+1} and β_{h+1} defined as follows and explicitly dependent on indices

(p_1, q_1, p_2, q_2) :

$$\begin{aligned}\alpha_{h+1}(p_1, q_1, p_2, q_2) &:= \frac{1}{2^{2m}-1} (\delta_{(p_1, q_1)_{h+1}} \delta_{(p_2, q_2)_{h+1}} - \frac{\delta_{(q_1, p_2)_{h+1}} \delta_{(p_1, q_2)_{h+1}}}{2^m}) \\ &= \frac{1}{2^{2m}-1} \left(C_{h+1}(p_1, q_1, p_2, q_2) - \frac{D_{h+1}(p_1, q_1, p_2, q_2)}{2^m} \right) \\ \beta_{h+1}(p_1, q_1, p_2, q_2) &:= \frac{1}{2^{2m}-1} (\delta_{(q_1, p_2)_{h+1}} \delta_{(p_1, q_2)_{h+1}} - \frac{\delta_{(p_1, q_1)_{h+1}} \delta_{(p_2, q_2)_{h+1}}}{2^m}) \\ &= \frac{1}{2^{2m}-1} \left(D_{h+1}(p_1, q_1, p_2, q_2) - \frac{C_{h+1}(p_1, q_1, p_2, q_2)}{2^m} \right)\end{aligned}$$

Using Eq.(B78), we integrate iteratively over the A_1 blocks of the first layer as follows:

$$\begin{aligned}& \prod_{j=1}^{A_1} E_{U_{j1}} [Tr[|p_1\rangle\langle q_1| WY_1 U^\dagger] Tr[|p_2\rangle\langle q_2| WY_2 U^\dagger]] \\ &= \prod_{j=1}^{A_1} E_{U_{j1}} [f_0(\emptyset)] \\ &= \prod_{j=2}^{A_1} E_{U_{j1}} [\alpha_1 f_1(1) + \beta_1 f_1(\emptyset)] \\ &= \prod_{j=3}^{A_1} E_{U_{j1}} [\alpha_1 (\alpha_2 f_2(12) + \beta_2 f_2(1)) + \beta_1 (\alpha_2 f_2(2) + \beta_2 f_2(\emptyset))] \\ &\vdots \\ &= \sum_{x_1 \in P_1} a_{x_1}(p_1, q_1, p_2, q_2) f_{A_1}(x_1), \quad a_{x_1}(p_1, q_1, p_2, q_2) := \prod_{h \in x_1} \alpha_h(p_1, q_1, p_2, q_2) \prod_{h \in \bar{x}_1} \beta_h(p_1, q_1, p_2, q_2)\end{aligned} \tag{B79}$$

Here P_1 is the power set of the ensemble $\{1, 2, \dots, A_1\}$, hence containing 2^{A_1} subsystems. We also recall from Eq.(B77) that:

$$f_{A_1}(x_1) := Tr[Tr_{x_1}[U^{(A_1,1)} Y_1 U^{(A_1,1)\dagger}] Tr_{x_1}[U^{(A_1,1)} Y_2 U^{(A_1,1)\dagger}]]$$

where $U^{(A_1,1)}$ is the unitary representing the circuit from the 2^{nd} up to the L^{th} layer (i.e. the A_1 blocks of the first layer are removed).

We notice that in the expression of $f_{A_1}(x_1)$, we no longer have the dependence on the indices $\mathbf{qp}$, which contribution is now contained in the coefficients a_{x_1} .

To integrate over the remaining blocks, we introduce the following notations $\forall t \in \{1, 2\}, \forall l \in \{2, \dots, L\}$ and $\forall h \in \{1, \dots, A_l\}$

$$\begin{aligned}Z_{h,l}^t &:= U^{(h,l)} Y_t U^{(h,l)\dagger} \\ g_{h,l}(x) &:= Tr[Tr_x[Z_{h,l}^1] Tr_x[Z_{h,l}^2]]\end{aligned}$$

where $U^{(h,l)}$ is the circuit unitary after removing the layers $i \in \{1, \dots, l-1\}$ and the first h blocks of the l^{th} layer. We also note that $g_{0,l}(x) = g_{A_{l-1}, l-1}(x) \forall l \in \{3, \dots, L\}$, $g_{0,2}(x) := f_{A_1}(x)$ and $g_{A_L, L}(x) := Tr[Tr_x[Y_1] Tr_x[Y_2]]$.

By applying Lemma 3, we obtain the following formula allowing to integrate recursively over the remaining blocks. We note that depending on the parity of the layer index l , we use s_h or τ_h (here we use τ_h but we get the same formula for s_h) to denote the subsystem on which the block h is acting non trivially:

$$E_{U_{h+1,l}}[g_{h,l}(x)] = \begin{cases} g_{h+1,l}(x), & \text{if } \tau_{h+1} \cap x \in \{\emptyset, \tau_{h+1}\} \\ \frac{2^{m/2}}{2^{m+1}} (g_{h+1,l}(x \cup \tau_{h+1}) + g_{h+1,l}(x \setminus \tau_{h+1})), & \text{otherwise} \end{cases} \tag{B80}$$

According to the recursive formula above, one can easily notice that for a fixed layer l , integrating $g_{0,l}(x)$ over its A_l blocks amounts to determining all the *non trivial* intersections between the current subsystem x and the layer blocks subsystems $\{\tau_h\}_{h=1}^{A_l}$ (if l is even otherwise $\{s_h\}_{h=1}^{A_l}$). In other terms, the set $I_l(x) := \{h|x \cap \tau_h \in \{(\tau_h)_1, (\tau_h)_{h+1}\} \forall h \in [A_l]\}$ will contain the blocks over which we are going to effectively integrate by adding the multiplicative factor $\frac{2^{m/2}}{2^m+1}$ and *branching* subsystem x into the two new subsystems $x \cup \tau_{h+1}$ and $x \setminus \tau_{h+1}$.

Formally, we get $\forall 2 \leq l \leq L$:

$$\prod_{j=1}^{A_l} E_{U_{ji}}[g_{0,l}(x_{l-1})] = \sum_{x_l \in P_l(x_{l-1})} \left(\frac{2^{m/2}}{2^m+1}\right)^{d_l(x_{l-1})} g_{A_l,l}(x_l) \quad (\text{B81})$$

where $d_l(x_{l-1}) := |I_l(x_{l-1})|$ and $P_l(x_{l-1})$ is the set of subsystems generated after each *branching* of x_{l-1} obtained by following the rule in Eq.(B80).

For ease of notation we drop the explicit dependence of a_{x_1} on (p_1, q_1, p_2, q_2) . By combining Eq.(B79) and Eq.(B81), we get the final expression of the second moment coefficient:

$$\begin{aligned} (M_{U(\theta)}^{(2)})_{\mathbf{p}, \mathbf{q}, \mathbf{r}, \mathbf{s}} &= \prod_{i=1}^L \prod_{j=1}^{A_i} E_{U_{ji}}[f_0(\emptyset)] \\ &= \prod_{i=3}^L \prod_{j=1}^{A_i} E_{U_{ji}} \left[\sum_{x_1 \in P_1} a_{x_1} \prod_{j=1}^{A_2} E_{U_{j2}}[f_{A_1}(x_1)] \right] \\ &= \prod_{i=3}^L \prod_{j=1}^{A_i} E_{U_{ji}} \left[\sum_{x_1 \in P_1} a_{x_1} \prod_{j=1}^{A_2} E_{U_{j2}}[g_{0,2}(x_1)] \right] \\ &= \prod_{i=3}^L \prod_{j=1}^{A_i} E_{U_{ji}} \left[\sum_{x_1 \in P_1} a_{x_1} \sum_{x_2 \in P_2(x_1)} \left(\frac{2^{m/2}}{2^m+1}\right)^{d_2(x_1)} g_{0,3}(x_2) \right] \\ &= \sum_{x_1 \in P_1} a_{x_1} \sum_{x_2 \in P_2(x_1)} \sum_{x_3 \in P_3(x_2)} \cdots \sum_{x_L \in P_L(x_{L-1})} \left(\frac{2^{m/2}}{2^m+1}\right)^{\sum_{i=2}^L d_i(x_{i-1})} g_{A_L,L}(x_L) \\ &= \sum_{x_1 \in P_1} a_{x_1} \sum_{x_2 \in P_2(x_1)} \sum_{x_3 \in P_3(x_2)} \cdots \sum_{x_L \in P_L(x_{L-1})} \left(\frac{2^{m/2}}{2^m+1}\right)^{\sum_{i=2}^L d_i(x_{i-1})} \text{Tr}[Tr_{x_L}[Y_1] Tr_{x_L}[Y_2]] \\ &= \sum_{x_1 \in P_1} a_{x_1} \sum_{x_L \in P_f(x_1)} \left(\frac{2^{m/2}}{2^m+1}\right)^{d(x_L)} \text{Tr}[Tr_{x_L}[Y_1] Tr_{x_L}[Y_2]] \\ &= \sum_{x_1 \in P_1} a_{x_1} \sum_{x_L \in P_f(x_1)} \left(\frac{2^{m/2}}{2^m+1}\right)^{d(x_L)} \delta(r_1, s_1)_{x_L} \delta(r_1, s_2)_{\overline{x_L}} \delta(r_2, s_2)_{x_L} \delta(r_2, s_1)_{\overline{x_L}} \end{aligned} \quad (\text{B82})$$

Where $P_f(x_1)$ is the set of all subsystems obtained by applying the branching rule in Eq.(B80) starting from x_1 and iterating through all the blocks in the circuits. We note as well that $d(x_L) := \sum_{i=2}^L d_i(x_{i-1})$ is the total number of branchings of x_1 leading to the final subsystem $x_L \in P_f(x_1)$.

One can remark from the above expression of the second moment superoperator coefficient that the coefficient a_{x_1} may scale in the number of blocks of the first layer $A_1 := A$ and that the sum $\sum_{x_L \in P_f(x_1)} \left(\frac{2^{m/2}}{2^m+1}\right)^{d(x_L)}$ will eventually scale in L , the circuit depth.

Moreover, we note that the expression holds for a light cone made of local 2-design blocks and acting on $n = m \times A_1$ qubits where $A_l = L - l + 1, \forall l \in [L]$. In this setting, we have exactly the same expression as in Eq.(B82), with $P_f(x_1)$ the set of subsystems obtained by applying the branching rule from Eq.B80 on the light cone blocks. Therefore, what differs between considering a full circuit made of local 2-design blocks and a light cone of the same circuit is the set of final subsystems $P_f(x_1)$ and its corresponding branching count $d(x_L) \forall x_L \in P_f(x_1)$.

However, in both cases, the construction of $P_f(x_1) \forall x_1 \in P_1$ and hence the calculation of $d(x_L)$ is quite cumbersome. Therefore, we derive the following upper bound on this sum.

To do so, we know that

$$f_A(x_1) = \text{Tr}[\text{Tr}_{x_1}[U^{(A)}Y_1U^{(A)\dagger}]\text{Tr}_{x_1}[U^{(A)}Y_2U^{(A)\dagger}]] = \sum_{x_L \in P_f(x_1)} \left(\frac{2^{m/2}}{2^m + 1}\right)^{d(x_L)} \text{Tr}[\text{Tr}_{x_L}[Y_1]\text{Tr}_{x_L}[Y_2]] \quad (\text{B83})$$

If Y_1, Y_2 are projectors (i.e. $r_i = s_i, i \in \{1, 2\}$), then using the sub multiplicative property of positive semi definite matrices, we get $f_A(x_1) \leq \text{Tr}[Y_1]\text{Tr}[Y_2] = 1$

For $Y_1 = Y_2 = |0\rangle\langle 0|$, we get $f_A(x_1) = \sum_{x \in P_f(x_1)} \left(\frac{2^{m/2}}{2^m + 1}\right)^{d(x)} \leq 1$. Therefore for arbitrary Y_1 and Y_2 , we get:

$$f_A(x_1) \leq \text{Tr}[\text{Tr}_{x_L}[Y_1]\text{Tr}_{x_L}[Y_2]] \quad (\text{B84})$$

We note that this upper bound holds for a full brickwise circuit or a lightcone over the same circuit, which can be observed in Section 5 (See Fig.14).

Indeed, as will be detailed in the next section, Eq.(B82) and Eq.(B84) will be key in obtaining the expression of the Fourier coefficients variance over a light cone, leading to the upper bound in theorem 3.

B.4.3 Proof of Theorem 3

We consider the brick-wise circuit architecture made of m -local encoding blocks and m -local 2-design parameterized blocks as illustrated in Fig.4. We recall Theorem 3 in this setting.

Theorem 8. *Consider a quantum model of the form in Eq.(1) and a parametrized circuit of the form in Eq.(2) using a brickwise architecture with $L = 1$ layers and observable $O = \hat{O}_{s_k} \otimes \mathbb{1}_{\overline{s_k}}$ acting non trivially on the m -qubit subsystem s_k . Assume that each trainable m -qubit unitary forms a local 2-design. The variance of each Fourier coefficient $c_\omega(\theta)$ for the frequencies $\omega \in \Omega$ appearing in the model Fourier decomposition in Eq.(3) is upper bounded as*

1. If $\|\hat{O}_{s_k}\|_2^2 \leq 2^m$, we have

$$\text{Var}[c_\omega] \leq \left(\frac{2^{m+1}}{2^{2m} - 1}\right)^{2L_2} |R_{E_k}(\omega)|^2. \quad (\text{B85})$$

2. If $\hat{O}_{s_k}$ is a projector of rank r , we have

$$\text{Var}[c_\omega] \leq \left(\frac{2^{m+1}}{2^{2m} - 1}\right)^{2L_2} \left(\frac{r}{2^m}\right)^2 |R_{E_k}(\omega)|^2. \quad (\text{B86})$$

Here $R_{E_k}(\omega)$ is the frequency generator obtained from the encoding blocks inside the observable backward light cone $\mathcal{L}_k$ (acting non trivially on $\mathcal{S}_{E_k}$) and L_2 is the depth of the post-encoding parameterized block.

Proof. In the following proof, using similar techniques as presented in B.4.2, we give an expression of the variance of the Fourier coefficients under the local 2-design setting described in Section 3.3. However, as the obtained expression involves cumbersome calculations, we give in Theorem 3 an upper bound on the Fourier coefficients variance that depends on the circuit depth and the frequency redundancy.

In what follows, we use the same notations from the previous section B.4.2 and we consider a circuit with a first parameterized layer $W^{(1)}$ of depth L_1 , a layer of encoding blocks $S(x)$ and a final parameterized layer $W^{(2)}$ of depth L_2 as depicted in Fig.4 acting initially on $n = m \times A$ qubits. We recall here the expression of the expectation of the modulus squared of a Fourier coefficient c_ω (See Eq.(B33) for details):

$$\begin{aligned}
\mathbb{E} [|c_\omega|^2] &= \sum_{\substack{I, I' \in R(\omega) \\ J, J' \in R(\omega)}} \mathbb{E}_{W^{(1)}} \left[\text{Tr} [|J'\rangle \langle J| W^{(1)} |0\rangle \langle 0| W^{(1)\dagger}] \text{Tr} [|I\rangle \langle I'| W^{(1)} |0\rangle \langle 0| W^{(1)\dagger}] \right] \\
&\quad \times \mathbb{E}_{W^{(2)}} \left[\text{Tr} [|J\rangle \langle J'| W^{(2)\dagger} O W^{(2)}] \text{Tr} [|I'\rangle \langle I| W^{(2)\dagger} O W^{(2)}] \right] \\
&= \sum_{\substack{I, I' \in R(\omega) \\ J, J' \in R(\omega)}} \Lambda(J, J', I, I') \Gamma(J, J', I, I')
\end{aligned} \tag{B87}$$

Here $J, J' \in R(\omega)$ are multi indices of length A , the number of blocks in a single layer and each component (j_k, j'_k) corresponds to the indices of 2 eigenvalues of the 2^m -dimensional encoding unitary acting on subsystem $s_k, \forall k \in \{1, \dots, A\}$. For ease of notation, we introduce the shorthand $\Lambda(J, J', I, I')$ for the expectation over the pre-encoding unitary $W^{(1)}$ and $\Gamma(J, J', I, I')$ for the expectation over the post-encoding unitary $W^{(2)}$ in Eq.(B87).

We then consider an m -local observable of the form $O_k = \hat{O}_{s_k} \otimes \mathbb{1}_{\overline{s_k}}$ acting non trivially on subsystem s_k and we denote by $\mathcal{L}_k$ the backward light cone associated to this observable. Precisely, $\mathcal{L}_k$ is the subcircuit containing all blocks with at least one qubit causally connected to the input qubits of $\hat{O}_{s_k}$. We also denote the subsystem on which $\mathcal{L}_k$ acts non trivially by $\mathcal{S}_{\mathcal{L}_k}$.

Due to the brick-wise circuit architecture and the locality of the observable, the quantum model is reduced to an effective model obtained by considering the restricted action of the circuit unitary on subsystem $\mathcal{S}_{\mathcal{L}_k}$:

$$f_{\mathcal{L}_k}(x) := \text{Tr} [|0\rangle \langle 0|_{\mathcal{L}_k} U(x, \theta)_{\mathcal{L}_k}^\dagger O_{\mathcal{L}_k} U(x, \theta)_{\mathcal{L}_k}] \tag{B88}$$

where $U(x, \theta)_{\mathcal{L}_k}^\dagger = W_{\mathcal{L}_k}^{(1)\dagger} S(x)_{\mathcal{L}_k}^\dagger W_{\mathcal{L}_k}^{(2)\dagger}$ is the adjoint circuit unitary restricted to the light cone acting on subsystem $\mathcal{S}_{\mathcal{L}_k}$.

Hence, Eq.(B87) can be rewritten using the following:

$$\begin{aligned}
\Lambda(J, J', I, I') &= \mathbb{E}_{W_{\mathcal{L}_k}^{(1)}} \left[\text{Tr} [|J'\rangle \langle J|_{\mathcal{L}_k} W_{\mathcal{L}_k}^{(1)} |0\rangle \langle 0|_{\mathcal{L}_k} W_{\mathcal{L}_k}^{(1)\dagger}] \text{Tr} [|I\rangle \langle I'|_{\mathcal{L}_k} W_{\mathcal{L}_k}^{(1)} |0\rangle \langle 0|_{\mathcal{L}_k} W_{\mathcal{L}_k}^{(1)\dagger}] \right] \\
\Gamma(J, J', I, I') &= \mathbb{E}_{W_{\mathcal{L}_k}^{(2)}} \left[\text{Tr} [|J\rangle \langle J'|_{\mathcal{L}_k} W_{\mathcal{L}_k}^{(2)\dagger} O_{\mathcal{L}_k} W_{\mathcal{L}_k}^{(2)}] \text{Tr} [|I'\rangle \langle I|_{\mathcal{L}_k} W_{\mathcal{L}_k}^{(2)\dagger} O_{\mathcal{L}_k} W_{\mathcal{L}_k}^{(2)}] \right]
\end{aligned}$$

Another consequence of restricting calculations to the light cone $\mathcal{L}_k$ is spectrum reduction where only encoding blocks acting on subsystem $\mathcal{S}_{\mathcal{L}_k}$ are involved. Hence, in what follows $R(\omega)$ refers to the frequency generator made only of encoding blocks acting non trivially on subsystem $\mathcal{S}_{\mathcal{L}_k}$. Let us denote by $\mathcal{S}_{E_k}$ the subsystem on which $W_{\mathcal{L}_k}^{(2)}$ (and similarly $S(x)$) acts non trivially and $\mathcal{S}_{\overline{E_k}} := \mathcal{S}_{E_k} \setminus \mathcal{S}_{E_k}$ as depicted in Fig.4. Consequently, we get

$$\begin{aligned}
\Gamma(J, J', I, I') &= \delta_{(J, J')_{\overline{E_k}}} \delta_{(I, I')_{\overline{E_k}}} \mathbb{E}_{W_{E_k}^{(2)}} \left[\text{Tr} [|J\rangle \langle J'|_{E_k} W_{E_k}^{(2)\dagger} O_{E_k} W_{E_k}^{(2)}] \text{Tr} [|I'\rangle \langle I|_{E_k} W_{E_k}^{(2)\dagger} O_{E_k} W_{E_k}^{(2)}] \right] \\
&= \delta_{(J, J')_{\overline{E_k}}} \delta_{(I, I')_{\overline{E_k}}} \Gamma_2(J, J', I, I')
\end{aligned} \tag{B89}$$

In what follows whenever $\Gamma(J, J', I, I')$ is used, it actually refers to $\Gamma(J, J', I, I')_2$. Therefore, Eq.(B87) becomes:

$$\mathbb{E} [|c_\omega|^2] = \sum_{\substack{I, I' \in R(\omega) \\ J, J' \in R(\omega)}} \delta_{(J, J')_{\overline{E_k}}} \delta_{(I, I')_{\overline{E_k}}} \Lambda(J, J', I, I') \Gamma(J, J', I, I') \tag{B90}$$

This implies that the above sum over $R(\omega)$ will effectively contain just the pairs $(J, J') \in R(\omega)$ such that $\forall s_l \in \mathcal{S}_{\overline{E_k}}, \lambda_{j_l} - \lambda_{j'_l} = 0$. Therefore the frequency generator becomes $R(\omega) := \{(J, J') \in [|2^m|]^{L_2} \times [|2^m|]^{L_2} \mid \sum_{l \in \mathcal{S}_{E_k}} (\lambda_{j_l} - \lambda_{j'_l}) = \omega\} \times \{(J, J) \in [|2^m|]^{L_1-1}\}$. In other terms, the effective spectrum

is the one generated by encoding blocks acting non trivially on subsystems of $\mathcal{S}_{E_k}$ and that is made redundant by adding null contributions from the encoding blocks outside the light cone.

In order to calculate $\Gamma(J, J', I, I')$ and $\Lambda(J, J', I, I')$, we use the same calculations from Eq.(B82) with a minor change of notations where A will take the values $L_1 + L_2 - 1$ and L_2 respectively. Therefore, we get:

$$\begin{aligned}
\Lambda(J, J', I', I) &= \sum_{x_1 \in P_1^1} a_{x_1}(J, J', I', I) f^1(x_1, |0\rangle\langle 0|_{\mathcal{L}_k}) \\
&= \sum_{x_1 \in P_1^1} a_{x_1}(J, J', I', I) \sum_{x_L \in P_f^1(x_1)} \left(\frac{2^{m/2}}{2^m + 1}\right)^{d(x_L)} \text{Tr}[Tr_{x_L}[|0\rangle\langle 0|_{\mathcal{L}_k}]^2] \\
\Gamma(J, J', I', I) &= \sum_{x_1 \in P_1^2} a_{x_1}(J, J', I', I) f^2(x_1, O_{E_k}) \\
&= \sum_{x_1 \in P_1^2} a_{x_1}(J, J', I', I) \sum_{x_L \in P_f^2(x_1)} \left(\frac{2^{m/2}}{2^m + 1}\right)^{d(x_L)} \text{Tr}[Tr_{x_L}[O_{E_k}]^2]
\end{aligned} \tag{B91}$$

where P_1^1 and P_1^2 are the power sets of $\mathcal{S}_{\mathcal{L}_k}$ and $\mathcal{S}_{E_k}$ respectively, P_f^1 and P_f^2 are the sets of final subsystems after branching over blocks of $W_{\mathcal{L}_k}^{(1)}$ and $W_{E_k}^{(2)}$ respectively, according to the construction detailed in Eq.(B82) and $f^i(x_1, H)$ is a shorthand for $\sum_{x_L \in P_f^i(x_1)} \left(\frac{2^{m/2}}{2^m + 1}\right)^{d(x_L)} \text{Tr}[Tr_{x_L}[H]^2]$, $\forall i \in \{1, 2\}$.

By substituting Eq.(B91) in Eq.(B90), we get the following expression of $\mathbb{E}[|c_\omega|^2]$:

$$\begin{aligned}
\mathbb{E}[|c_\omega|^2] &= \sum_{\substack{I, I' \in R(\omega) \\ J, J' \in R(\omega)}} \delta(J, J')_{\overline{E_k}} \delta(I, I')_{\overline{E_k}} \Lambda(J, J', I', I) \Gamma(J, J', I', I) \\
&= \sum_{\substack{I, I' \in R(\omega) \\ J, J' \in R(\omega)}} \delta(J, J')_{\overline{E_k}} \delta(I, I')_{\overline{E_k}} \left(\sum_{x_1 \in P_1^1} a_{x_1}(J, J', I', I) f^1(x_1, |0\rangle\langle 0|_{\mathcal{L}_k}) \right) \left(\sum_{y_1 \in P_1^2} a_{y_1}(J, J', I', I) f^2(y_1, O_{E_k}) \right) \\
&= \sum_{\substack{J, J' \in R_{E_k}(\omega) \\ J, J' \in R_{E_k}(\omega)}} \left(\sum_{y_1 \in P_1^2} a_{y_1} f^2(y_1, O_{E_k}) \right) \sum_{\substack{(J, J')_{\overline{E_k}} \\ (I, I')_{\overline{E_k}}}} \delta(J, J')_{\overline{E_k}} \delta(I, I')_{\overline{E_k}} \left(\sum_{x_1 \in P_1^1} a_{x_1} f^1(x_1, |0\rangle\langle 0|_{\mathcal{L}_k}) \right) \\
&= \left(\frac{1}{2^{2m} - 1} \right)^{L_1 + 2L_2 - 1} \sum_{\substack{J, J' \in R_{E_k}(\omega) \\ J, J' \in R_{E_k}(\omega)}} \left(\sum_{y_1 \in P_1^2} \tilde{a}_{y_1} f^2(y_1, O_{E_k}) \right) \sum_{x_1 \in P_1^1} f^1(x_1, |0\rangle\langle 0|_{\mathcal{L}_k}) \sum_{\substack{(J, J')_{\overline{E_k}} \\ (I, I')_{\overline{E_k}}}} \delta(J, J')_{\overline{E_k}} \delta(I, I')_{\overline{E_k}}
\end{aligned} \tag{B92}$$

where $a_{x_1} := \left(\frac{1}{2^{2m} - 1}\right)^{L_1 + L_2 - 1} \tilde{a}_{x_1}$ and $a_{y_1} := \left(\frac{1}{2^{2m} - 1}\right)^{L_2} \tilde{a}_{y_1}$.

Summing over indices in $\overline{E_k}$ gives the following:

$$\begin{aligned}
\sum_{\substack{(JJ')_{\overline{E_k}} \\ (I,I')_{\overline{E_k}}}} \delta(J, J')_{\overline{E_k}} \delta(I, I')_{\overline{E_k}} \tilde{a}_{x_1} &= \sum_{J_{\overline{E_k}}, I_{\overline{E_k}}} \tilde{a}_{x_1} \delta(J, J')_{\overline{E_k}} \delta(I, I')_{\overline{E_k}} \\
&= \sum_{J_{\overline{E_k}}, I_{\overline{E_k}}} \prod_{l \in x_1} (C_l - \frac{D_l}{2^m}) \prod_{l \in \overline{x_1}} (D_l - \frac{C_l}{2^m}) (\delta(J, J')_{\overline{E_k}} \delta(I, I')_{\overline{E_k}}) \\
&= \prod_{l \in x_1 \cap E_k} (C_l - \frac{D_l}{2^m}) \prod_{l \in \overline{x_1} \cap E_k} (D_l - \frac{C_l}{2^m}) \sum_{J_{\overline{E_k}}, I_{\overline{E_k}}} \prod_{l \in x_1 \cap \overline{E_k}} (1 - \frac{\delta_{j_l, i_l}}{2^m}) \prod_{l \in \overline{x_1} \cap \overline{E_k}} (\delta_{j_l, i_l} - \frac{1}{2^m}) \\
&= \prod_{l \in x_1 \cap E_k} (C_l - \frac{D_l}{2^m}) \prod_{l \in \overline{x_1} \cap E_k} (D_l - \frac{C_l}{2^m}) \prod_{l \in x_1 \cap \overline{E_k}} \sum_{j_l, i_l} (1 - \frac{\delta_{j_l, i_l}}{2^m}) \prod_{l \in \overline{x_1} \cap \overline{E_k}} \sum_{j_l, i_l} (\delta_{j_l, i_l} - \frac{1}{2^m}) \\
&= \prod_{l \in x_1 \cap E_k} (C_l - \frac{D_l}{2^m}) \prod_{l \in \overline{x_1} \cap E_k} (D_l - \frac{C_l}{2^m}) \prod_{l \in x_1 \cap \overline{E_k}} (2^{2m} - 1) \prod_{l \in \overline{x_1} \cap \overline{E_k}} (2^m - 2^m) \\
&= \prod_{l \in x_1 \cap E_k} (C_l - \frac{D_l}{2^m}) \prod_{l \in \overline{x_1} \cap E_k} (D_l - \frac{C_l}{2^m}) (2^{2m} - 1)^{|x_1 \cap \overline{E_k}|} 0^{|\overline{x_1} \cap \overline{E_k}|} \\
&= \begin{cases} \prod_{l \in x_1 \cap E_k} (C_l - \frac{D_l}{2^m}) \prod_{l \in \overline{x_1} \cap E_k} (D_l - \frac{C_l}{2^m}) (2^{2m} - 1)^{|\overline{E_k}|}, & \text{if } \overline{E_k} \subset x_1 \\ 0, & \text{otherwise} \end{cases} \\
&= \begin{cases} \prod_{l \in x_1 \cap E_k} (C_l - \frac{D_l}{2^m}) \prod_{l \in \overline{x_1} \cap E_k} (D_l - \frac{C_l}{2^m}) (2^{2m} - 1)^{L_1 - 1}, & \text{if } \overline{E_k} \subset x_1 \\ 0, & \text{otherwise} \end{cases}
\end{aligned}$$

Thus, by substituting this above sum in Eq.(B92), we obtain:

$$\begin{aligned}
\mathbb{E} [|c_\omega|^2] &= (\frac{1}{2^{2m} - 1})^{2L_2} \sum_{\substack{J, J' \in R_{E_k}(\omega) \\ I, I' \in R_{E_k}(\omega)}} \sum_{y_1 \in P_1^2} \tilde{a}_{y_1} f^2(y_1, O_{E_k}) \sum_{\substack{x_1 \in P_1^1 \\ \overline{E_k} \subset x_1}} \prod_{l \in x_1 \cap E_k} (C_l - \frac{D_l}{2^m}) \prod_{l \in \overline{x_1} \cap E_k} (D_l - \frac{C_l}{2^m}) f^1(x_1, |0\rangle\langle 0|_{\mathcal{L}_k}) \\
&= (\frac{1}{2^{2m} - 1})^{2L_2} \sum_{\substack{J, J' \in R_{E_k}(\omega) \\ I, I' \in R_{E_k}(\omega)}} \sum_{y_1 \in P_1^2} \tilde{a}_{y_1} f^2(y_1, O_{E_k}) \sum_{x_1 \in P_1^2} \tilde{a}_{x_1} f^1(x_1 \cup \overline{E_k}, |0\rangle\langle 0|_{\mathcal{L}_k})
\end{aligned} \tag{B93}$$

In the above expression of $\mathbb{E} [|c_\omega|^2]$, we have $|\tilde{a}_x| \leq 1 \ \forall x \in P_1^2$ and from Eq.(B84) we have:

$$\begin{aligned}
f^1(x_1 \cup \overline{E_k}, |0\rangle\langle 0|_{\mathcal{L}_k}) &= \sum_{x \in P_f^1(x_1 \cup \overline{E_k})} (\frac{2^{m/2}}{2^m + 1})^{d(x)} Tr[Tr_{x_L} [|0\rangle\langle 0|_{\mathcal{L}_k}]^2] \\
&\leq Tr[Tr_{x_L} [|0\rangle\langle 0|_{\mathcal{L}_k}]^2] = 1 \\
f^2(y_1, O_{E_k}) &= \sum_{x \in P_f^2(y_1)} (\frac{2^{m/2}}{2^m + 1})^{d(x)} Tr[Tr_{x_L} [O_{E_k}]^2] \\
&\leq Tr[Tr_{x_L} [O_{E_k}]^2] \\
&\leq 2^{2mL_2}
\end{aligned}$$

where we use in the last inequality the hypothesis that $Tr[\hat{O}_{s_k}^2] \leq 2^m$.

Finally, by combining all the previous steps, we retrieve the upper bound in Theorem 3:

$$\mathbb{E} [|c_\omega|^2] \leq \left(\frac{2^{m+1}}{2^{2m} - 1} \right)^{2L_2} |R_{E_k}(\omega)|^2$$

□

B.4.4 Fourier coefficients variance in a light cone forming a 2-design

In this section, we consider the same settings and notations from the previous section B.4.3 but we assume that $W_{\mathcal{L}_k}^{(1)}$ and $W_{E_k}^{(2)}$ form each a 2-design on subsystems $\mathcal{S}_{\mathcal{L}_k}$ and $\mathcal{S}_{E_k}$ respectively. We follow the same steps as in the previous proof until we get to Eq.(B90). Then, we compute $\Gamma(J, J', I, I')$ and $\Lambda(J, J', I, I')$ under the 2-design assumption. Hence, by applying Weingarten calculus expression of the second moment [34], we obtain:

$$\begin{aligned}\Lambda(J, J', I, I') &= \frac{1}{2^{2m(L_1+L_2-1)} - 1} \left(1 - \frac{1}{2^{m(L_1+L_2-1)}} \right) [\delta_{J,J'}\delta_{I,I'} + \delta_{J,I}\delta_{J',I'}] \\ \Gamma(J, J', I, I') &= \frac{1}{2^{2mL_2} - 1} \left[\delta_{(J,J')_{E_k}} \delta_{(I,I')_{E_k}} \left(\text{Tr}[O_{E_k}]^2 - \frac{\text{Tr}[O_{E_k}^2]}{2^{mL_2}} \right) + \delta_{(J,I)_{E_k}} \delta_{(J',I')_{E_k}} \left(\text{Tr}[O_{E_k}^2] - \frac{\text{Tr}[O_{E_k}]^2}{2^{mL_2}} \right) \right]\end{aligned}\quad (\text{B94})$$

Therefore, we substitute $\Gamma(J, J', I, I')$ and $\Lambda(J, J', I, I')$ in Eq.(B90) using the new expressions in Eq.(B94) and we get $\forall \omega \in \Omega \setminus \{0\}$:

$$\begin{aligned}\mathbb{E}[|c_\omega|^2] &= \frac{1}{2^{m(L_1+L_2-1)}(2^{m(L_1+L_2-1)} + 1)(2^{2mL_2} - 1)} \sum_{\substack{I, I' \in R(\omega) \\ J, J' \in R(\omega)}} \delta_{(J,J')_{\overline{E_k}}} \delta_{(I,I')_{\overline{E_k}}} \delta_{J,I} \delta_{J',I'} \left(\text{Tr}[O_{E_k}^2] - \frac{\text{Tr}[O_{E_k}]^2}{2^{mL_2}} \right) \\ &= \frac{\left(\text{Tr}[O_{E_k}^2] - \frac{\text{Tr}[O_{E_k}]^2}{2^{mL_2}} \right)}{2^{m(L_1+L_2-1)}(2^{m(L_1+L_2-1)} + 1)(2^{2mL_2} - 1)} \sum_{\substack{(I,I')_{\overline{E_k}} \\ (J,J')_{\overline{E_k}}}} \delta_{(J,J')_{\overline{E_k}}} \delta_{(I,I')_{\overline{E_k}}} \delta_{(J,I)_{\overline{E_k}}} \delta_{(J',I')_{\overline{E_k}}} \\ &\quad \times \sum_{\substack{I, I' \in R_{E_k}(\omega) \\ J, J' \in R_{E_k}(\omega)}} \delta_{(J,I)_{E_k}} \delta_{(J',I')_{E_k}} \\ &= \frac{2^{m(L_2-1)} \left(\text{Tr}[\hat{O}_k^2] - \frac{\text{Tr}[\hat{O}_k]^2}{2^m} \right)}{2^{m(L_1+L_2-1)}(2^{m(L_1+L_2-1)} + 1)(2^{2mL_2} - 1)} 2^{m(L_1-1)} |R_{E_k}(\omega)| \\ &= \frac{1}{2^m(2^{m(L_1+L_2-1)} + 1)(2^{2mL_2} - 1)} \left(\text{Tr}[\hat{O}_k^2] - \frac{\text{Tr}[\hat{O}_k]^2}{2^m} \right) |R_{E_k}(\omega)|\end{aligned}\quad (\text{B95})$$

B.5 Proof of theorem 4

We recall Theorem 4.

Theorem 9 (Fourier Norm Bound). *Consider a quantum model $f(x, \theta)$ of the form in Eq.(1) using an observable O and a parametrized circuit of the form in Eq.(2) with $L \geq 1$ layers. Also assume that the encoding Hamiltonians are fixed, giving rise to a spectrum Ω . Then,*

$$\forall x \in \mathbb{R}^d, \forall \theta \in \Theta, |f(x, \theta)|^2 \leq \|O\|_\infty^2 \quad (\text{B96})$$

$$\forall \theta \in \Theta, \sum_{\omega \in \Omega} |c_\omega(\theta)|^2 \leq \|O\|_\infty^2 \quad (\text{B97})$$

Proof. The first point can be proven by remarking that $\langle \psi | O | \psi \rangle$ can be maximized by taking $|\psi\rangle$ as the eigenvector associated to the largest eigenvalue of O .

For the second point, one can write by considering the half spectrum Ω_+

$$|f(x)|^2 = \left| \sum_{\omega \in \Omega_+} c_\omega e^{-i\omega^\top x} + c_\omega^* e^{i\omega^\top x} \right|^2 \quad (\text{B98})$$

$$= \sum_{\omega \in \Omega_+} (c_\omega e^{-i\omega^\top x} + c_\omega^* e^{i\omega^\top x})(c_\omega^* e^{i\omega^\top x} + c_\omega e^{-i\omega^\top x}) \quad (\text{B99})$$

$$+ 2 \sum_{\omega_1 \neq \omega_2} (c_{\omega_1} e^{-i\omega_1^\top x} + c_{\omega_1}^* e^{i\omega_1^\top x})(c_{\omega_2} e^{-i\omega_2^\top x} + c_{\omega_2}^* e^{i\omega_2^\top x}) \quad (\text{B100})$$

$$= \sum_{\omega \in \Omega_+} 2|c_\omega|^2 + c_\omega^2 e^{-2i\omega^\top x} + c_\omega^{2*} e^{2i\omega^\top x} \quad (\text{B101})$$

$$+ 2 \sum_{\omega_1 \neq \omega_2} c_{\omega_1} c_{\omega_2}^* e^{-i(\omega_1 - \omega_2)^\top x} + c_{\omega_1} c_{\omega_2} e^{-i(\omega_1 + \omega_2)^\top x} + c_{\omega_1}^* c_{\omega_2}^* e^{i(\omega_1 + \omega_2)^\top x} + c_{\omega_1}^* c_{\omega_2} e^{i(\omega_1 - \omega_2)^\top x} \quad (\text{B102})$$

$$= 2 \sum_{\omega \in \Omega_+} |c_\omega|^2 + g(x) \quad (\text{B103})$$

We finish the proof by finding x_0 such that $g(x_0) = 0$. This can be done with the lemma 3. We first reduce the problem to one variable and we introduce $h(t) = \sum_{\omega \in \Omega} a_\omega \cos(\omega_1 t + \omega_{2:d}^\top x'_0 + \phi_\omega)$ where x'_0 is an arbitrary vector of size $d - 1$. $\square$

Lemma 3. *Let $h(t) = \sum_{\omega \in \Omega} a_\omega \cos(\omega t + \phi_\omega)$, with Ω being a discrete subset of $\mathbb{R}^d$, and $\phi_\omega \in \mathbb{R}$. Then it exists t_0 such that $g(t_0) = 0$.*

Proof. Let us suppose that $h(t)$ is of constant sign, we can assume it is positive on $\mathbb{R}$. If h is not of constant sign, by continuity it means that it exists t_0 such that $h(t_0) = 0$ and the proof is finished. We will show that $\forall \varepsilon > 0$ it exists T such that $|\int_0^T h(t) dt| < \varepsilon$. Then it means that h is equal to 0 over all the interval $[0, T]$, which proves the result.

We will now prove that $\forall \varepsilon > 0$, there exists a real T and integers q_ω s such that $\forall \omega \in \Omega, |2\pi T - q_\omega 2\pi/\omega| \leq \varepsilon$ which proves that $|\int_0^{2\pi T} h(t) dt| < \varepsilon$ since for each frequency we integrate over an integer number of periods.

We will show how to construct T and q_ω s for three values ω_1, ω_2 and ω_3 .

Let $R_2 = \{k \frac{\omega_2}{\omega_1} - \lfloor k \frac{\omega_2}{\omega_1} \rfloor, k \in \mathbb{N}, k \in [0, N^2]\}$

If one divides the interval $[0, 1]$ into N equal subintervals, then there are at least $N + 1$ elements of R_2 that are in the same subinterval. Let $\{k_1, \dots, k_{N+1}\}$ the integers corresponding to these elements.

Then let $R_3 = \{k \frac{\omega_3}{\omega_1} - \lfloor k \frac{\omega_3}{\omega_1} \rfloor, k \in \mathbb{N}, k \in \{k_1, \dots, k_{N+1}\}\}$. If one divides again the interval $[0, 1]$ into N equal subintervals, then there are at least 2 elements of R_3 that are in the same subinterval.

Then there exists k, k' such that

$$\begin{aligned} \left| k \frac{\omega_3}{\omega_1} - \lfloor k \frac{\omega_3}{\omega_1} \rfloor - (k' \frac{\omega_3}{\omega_1} - \lfloor k' \frac{\omega_3}{\omega_1} \rfloor) \right| &\leq \frac{1}{N} \\ \left| (k - k') \frac{\omega_3}{\omega_1} - (\lfloor k \frac{\omega_3}{\omega_1} \rfloor - \lfloor k' \frac{\omega_3}{\omega_1} \rfloor) \right| &\leq \frac{1}{N} \\ \left| (k - k') \frac{1}{\omega_1} - \frac{1}{\omega_3} (\lfloor k \frac{\omega_3}{\omega_1} \rfloor - \lfloor k' \frac{\omega_3}{\omega_1} \rfloor) \right| &\leq \frac{1}{N\omega_3} \end{aligned}$$

And since k and k' are in $\{k_1, \dots, k_{N+1}\}$, we also have

$$\left| (k - k') \frac{1}{\omega_1} - \frac{1}{\omega_2} (\lfloor k \frac{\omega_2}{\omega_1} \rfloor - \lfloor k' \frac{\omega_2}{\omega_1} \rfloor) \right| \leq \frac{1}{N\omega_2}$$

By taking $T = (k - k') \frac{1}{\omega_1}$ and $1/N = \varepsilon$, we proved the result for 3 numbers. One can apply the same construction for $|\Omega|$ number of frequencies by taking $N^{|\Omega|}$ integers at the beginning. $\square$

C Additional Numerics

C.1 Fourier coefficients variance in the 2-design setting for a reuploading model

We consider the same settings of Theorem 6 where we take a reuploading circuit with $L = 2$ circuit layers acting on $n = 4$ qubits. In Fig.15, we see that the simulated variance of the Fourier coefficients match the theoretical variance values predicted by Theorem 6 for the Pauli and exponential encoding strategies. The explicit dependence on the redundancies is harder to visualize as was done in Fig.7 because the variance expression from Theorem 6 include the partial redundancies. Nonetheless, we clearly observe in the case of Pauli encoding that we have a gaussian distribution (negative frequencies are not plotted) with half of the frequencies suppressed. For the exponential encoding strategy, we see a gaussian distribution but with a higher variance than the Pauli case which is due to redundancies caused by the reuploading scheme and the fact that we are using the same encoding layer twice as explained in Appendix A.2.




Figure 15: For $n = 4$ qubits, two circuit layers ($L = 2$) with the same encoding gates in each of the two encoding layers, five repetitions of the strongly entangling ansatz per trainable layer (see Appendix D); relation between the variance of each Fourier coefficient $\text{Var}[c_\omega]$ (blue dots) and its corresponding partial redundancies given by theorem 6 (green bars). Values are given for two different encoding strategies, a) Pauli encoding and b) exponential encoding.

C.2 Fourier coefficients variance in the approximate 2-design setting for a reuploading model

In Section 3.2, we showed in Theorem 7 an upper bound on $\text{Var}(c_\omega)$ in the case of a model with a single circuit layer ($L = 1$), being a second degree polynomial in $|R(\omega)|$. It was left as an open question to demonstrate a similar bound for $L > 1$. Fig.16 shows a simulation in the case of $L = 2$ circuit layers. We tried to fit a second degree polynomial in order to express the relation between $\text{Var}(c_\omega)$ and $|R(\omega)|$ directly. Even though one cannot simulate an upper bound but just a direct correlation, seeing this second-degree polynomial fitting well could indicate that the same kind of bound could be expected.

D Ansätze for trainable layers

This section serves as a reference to the different circuit architectures that we considered in Section 5, explaining how a particular ansatz structure can be scaled in the number of qubits.




Figure 16: Fitting Fourier coefficients variance $\text{Var}[c_\omega]$ to a second degree polynomial in the redundancies $|R(\omega)|$ for four different circuit architecture (connectivity). As the locality of the trainable block decreases V , ϵ gets bigger and the trainable layers farther away from being a 2-design. R^2 assesses the goodness of the fit, the closer to one the better.




Figure 17: *Strongly Entangling Ansatz* [37]: The Ansatz depth varies with the number of qubits. To have all possible entanglers between qubits, one needs to have the number of blocks ℓ to be $\ell \geq n - 1$, with n the number of qubits and r is the range of the control gates given by $r = \ell \bmod(n)$.




Figure 18: *Simplified Two Design Ansatz* [27]: An initial layer of Pauli-Y rotations and controlled-Z entanglers. Note that this ansatz has the same approximately the same amount of entangler gates as for the basic entangling ansatz layer, but double the number of parameters. It is composed of an initial layer and then the periodic layer is the one that will be repeated.