Progress toward practical quantum cryptanalysis by variational quantum cloning

Brian Coyle,[1] Mina Doosti,[1] Elham Kashefi,[1,2] and Niraj Kumar[1]
[1]*School of Informatics, 10 Crichton Street, Edinburgh EH8 9AB, United Kingdom,*
[2]*CNRS, LIP6, Sorbonne Université, 4 place Jussieu, 75005 Paris, France*

(Received 27 May 2021; accepted 11 March 2022; published 11 April 2022)

Cryptanalysis of quantum cryptographic systems generally involves finding optimal adversarial attack strategies on the underlying protocols. The core principle of modeling quantum attacks often reduces to the ability of the adversary to clone unknown quantum states and to extract thereby meaningful secret information. Explicit optimal attack strategies typically require high computational resources due to large circuit depths or, in many cases, are unknown. Here we introduce variational quantum cloning (VarQlone), a cryptanalysis algorithm based on quantum machine learning, which allows an adversary to obtain optimal approximate cloning strategies with short depth quantum circuits, trained using hybrid classical-quantum techniques. The algorithm contains operationally meaningful cost functions with theoretical guarantees, quantum circuit structure learning and gradient-descent-based optimization. Our approach enables the end-to-end discovery of hardware-efficient quantum circuits to clone specific families of quantum states, which we demonstrate in an implementation on the Rigetti Aspen quantum hardware. We connect these results to quantum cryptographic primitives and derive explicit attacks facilitated by VarQlone. We expect that quantum machine learning will serve as a resource for improving attacks on current and future quantum cryptographic protocols.

DOI: 10.1103/PhysRevA.105.042604

I. INTRODUCTION

Quantum cryptography is one of the prominent application areas for quantum technologies. Early proposals for quantum cryptographic protocols appeared as early as the 1980s with conceptual ideas including quantum money [1] and quantum key distribution [2]. Since then, new protocols are being developed at a staggering rate, exploiting quantum phenomena such as entanglement and nonlocality of quantum correlations for security proofs [3–6]. On the experimental side, rapid progress is being made as well, with the first violation of loophole-free Bell inequalities [7] demonstrated in 2015 and the satellite Micius implementing quantum protocols over long distances, including quantum key distribution (QKD) [8–10]. For an overview of recent advances in quantum cryptography, see the topical review [11].

At the heart of the security of many quantum protocols, is a fundamental pillar of quantum mechanics: the *no-cloning* theorem. This theorem states that the *cloning* of arbitrary quantum information perfectly and deterministically is forbidden by the laws of quantum mechanics [12]. Its connection to cryptography is easily seen in "prepare and measure" QKD protocols. If an adversary is capable of intercepting and making perfect copies of quantum states sent between two parties communicating using some secret key (encoded in quantum information) they can, in principle, learn the underlying secret information. The fact that the adversary is unable to do so under a foundational quantum mechanical principle leads to many potential advantages in using quantum communication protocols including, in many cases, realization of information-theoretic security guarantees.

However, the original formulation of the no-cloning theorem is not sufficient to fully analyze the security of protocols as it merely states that perfect and deterministic cloning is forbidden. In a remarkable discovery [13] it was established that cloning becomes possible, to some extent, if one is willing to relax the two assumptions in the no-cloning theorem. As a consequence, an adversary can partially learn hidden information. Specifically, removing the requirement of generating "perfect" clones gives *approximate* cloning [13], and relaxing determinism gives *probabilistic* cloning [14,15]. Both of these subfields of quantum information have a rich history and have been widely studied. For comprehensive reviews see Refs. [16,17].

So far, we used here the example of cloning as an attack by an adversary on quantum protocols. It turns out that in some cases these types of attacks are actually optimal [16]. In cases where they are not, cloning provides a means to at least provide lower bounds on the strategies of an adversary [18]. However, actually *implementing* such cloning-based attacks might be nontrivial in practice. For example, building general cloning circuits can require quantum circuits of large depth, which are not easily implementable on noisy intermediate scale quantum (NISQ) [19] devices (universal quantum computers with on the order of 50–200 physical, noisy qubits, which lack the capabilities of quantum error correction [20,21]). The effect of decoherence and errors in these devices puts the production of high-quality clones out of reach of an adversary with NISQ resources. Worse still, in many cases the optimal transformations to perform cloning might not be known, in particular for families of states whose optimal cloning fidelities are difficult to extract analytically.

On the other hand, there has been much interest in implementing quantum cloning and cryptographic attacks on protocols via specific and tailored experiments (for example, [22–26]) but these may not be easily reconfigurable or generalizable to other scenarios. In summary, finding and constructing quantum cloning circuits to prepare high-fidelity clones on NISQ hardware is challenging.

The present proposal adds a ingredient towards the circumvention of these issues. We give an algorithm—"variational quantum cloning" (VarQlone)—which uses quantum machine learning [27–30] (QML) techniques to *learn* to clone quantum states in an end-to-end manner. VarQlone is made possible by recent advances and techniques in the field of *variational* quantum algorithms (VQAs) [31–35]. VQAs are intentionally tailored to be useful on NISQ devices, which are too small and noisy to implement "coherent" algorithms with speedups, for example, in factoring large prime numbers [36]. However, such devices are capable of performing tasks which cannot be simulated by any classical device in reasonable time [37–39], making the search for dedicated applications a topic of likely practical relevance. VQAs have been proposed or used for various applications, including quantum chemistry [40] and combinatorial optimization [41].

The core quantum component is typically a parameterized quantum circuit (PQC) [42] and when VQAs are applied to machine-learning problems, they have come to be seen as quantum neural networks [43,44]. This is because they can achieve many of the same tasks as classical neural networks, [45,46] and can outperform them in certain cases [47–49]. Furthermore, machine learning techniques, both quantum [50–56] and classical [57–61] has proven useful in *discovering* interesting variations of, and providing insights into quantum algorithms and subroutines. This line of study even extends to the foundations of quantum mechanics [62].

VarQlone is different from other variational algorithms in that it can be viewed as a first step into a different area of application, *variational quantum cryptanalysis*. Specifically, in using QML techniques to learn to clone quantum states, VarQlone can discover unique ways to attack quantum protocols, in particular those whose underlying security reduces to quantum cloning. Furthermore, in developing such techniques more generally we can draw on the relationship between classical machine learning and deep learning, with classical cryptography [63–66].

For concreteness in this work, we focus on cloning families of quantum states used in two distinct families of quantum cryptographic protocols: those used in QKD protocols and those used in quantum coin-flipping protocols. For the former, we use the canonical BB-84 protocol [2] as an example, which can be attacked by cloning *phase-covariant* states. For quantum coin-flipping protocols so-called *fixed-overlap* states are typically used. Furthermore, the security of coin-flipping protocols are typically much less well studied than QKD protocols, so we also explicitly give analytic (cloning-based) attacks on two examples of such protocols, that of Mayers *et al.* [67] and that of Aharonov *et al.* [68] We then use VarQlone to construct these attacks explicitly. As part of our algorithm we define suitable cost functions, prove theoretical guarantees for them (including notions of faithfulness [51] and "barren plateaus" [69,70]), and use state-of-the-art techniques such as quantum architecture search and variable structure Ansätze [71]. To illustrate this, Fig. 1 shows a cartoon of how VarQlone can be used in a cryptographic scenario.

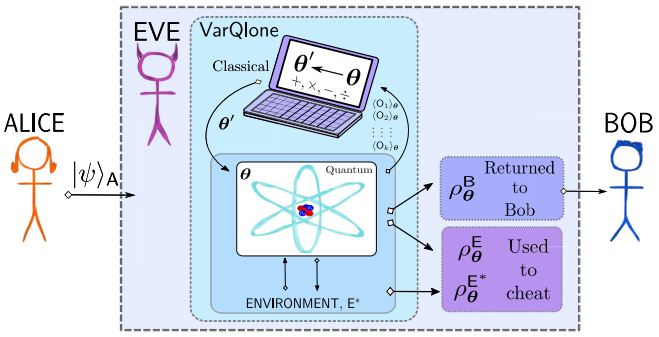

FIG. 1. Cartoon overview of VarQlone in a cryptographic attack. Here an adversary Eve, E, implements a $1 \to 2$ cloning attack on states used in a quantum protocol (for example, QKD) between Alice and Bob. Eve intercepts the states sent by Alice $|\psi\rangle_A$ and may interact with an ancillary "environment," E^*. This interaction is trained (an optimal parameter setting $\boldsymbol{\theta}$ is found) by Eve to optimally produce clones, $\rho^B_{\boldsymbol{\theta}}, \rho^E_{\boldsymbol{\theta}}$. In order to attack the protocol, Eve will return $\rho^B_{\boldsymbol{\theta}}$ to Bob and use the rest (her clone, $\rho^E_{\boldsymbol{\theta}}$ plus the remaining environment state, $\rho^{E^*}_{\boldsymbol{\theta}}$) to cheat. The training procedure consists of using a classical computer to optimize the quantum parameters, via a cost function. The cost is a function of k observables, O_k, measured from the output states, which are designed to extract fidelities of the states to compare against the ideal state.

Finally, to underline the practical potential of our approach we implemented it on the Rigetti Aspen quantum computer and show how VarQlone can learn to clone states with a higher fidelity on this device than previously known "analytic" quantum circuits, highlighting the flexibility of our approach. Furthermore, the nature of VarQlone allows us to improve cloning fidelities generically, on quantum computers available through the cloud [72], without significant tweaking and custom built experimental hardware.

II. QUANTUM CLONING

Approximate quantum cloning allows circumventing the no-cloning theorem, but there are still a number of subtleties that have to be addressed when building quantum cloning machines (QCMs), i.e., the unitaries U. Figure 1 illustrates so-called $1 \to 2$ cloning, but we can generalize to $M \to N$ cloning. Here M copies of the input state to be cloned is given, and N output "clones" are requested. Next, we have to choose the *comparison metric* which is used to compare the clones outputted from the QCM, relative to the ideal input states. It is common to use the *fidelity* [73], between quantum states ρ, σ:

$$F(\rho,\sigma) = \left(\mathrm{Tr}\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}}\right)^2. \tag{1}$$

We will require two specifications of the fidelity. The first is the *local* fidelity, $F^j_{\mathsf{L}} := F_{\mathsf{L}}(\sigma_j, \rho_A)$, which compares the ideal input state, $\rho_A := |\psi\rangle\langle\psi|_A$ (of which we may have M identical copies), to the output clones, $\sigma_j := \mathrm{Tr}_{\bar{j}}(U^\dagger \rho_A^{\otimes M} \otimes \rho_{\mathrm{R}} \otimes \rho_{\mathrm{aux}} U)$, $j \in \{1, \ldots, N\}$, i.e., the reduced states of the

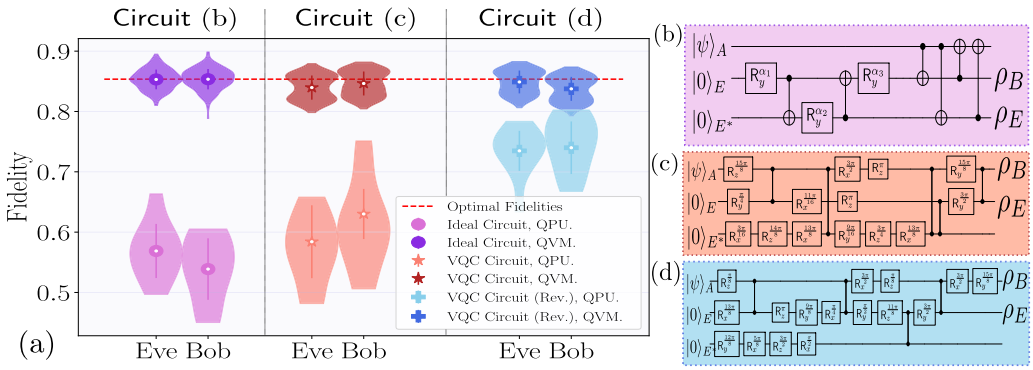

FIG. 2. Variational Quantum Cloning implemented on phase-covariant states using three qubits of the Rigetti Aspen-8 chip (QPU), plus simulated results (QVM). Violin plots in (a) show the cloning fidelities, for Bob and Eve, found using each of the circuits shown in (b)–(d), respectively. Shown in red is the maximal possible fidelity for this problem. (b) The ideal circuit with clones appearing in registers 2 and 3. (c) The structure-learned circuit for the same scenario, using one less entangling gate. (d) The effect of allowing clones to appear in registers 1 and 2. In the latter case, only four (nearest-neighbor) entangling gates are used, demonstrating a significant boost in performance on the QPU.

QCM output. $\mathrm{Tr}_{\bar{j}}(\cdot)$ denotes the partial trace over all subsystems except j, ρ_{aux} is an auxiliary subsystem which may be needed for cloning and ρ_{R} is a quantum register of size $N-M$ which will receive the clones. The second is the *global* fidelity compares the entire output state of the QCM to a product state of input copies, $F_{\mathsf{G}}(\mathrm{Tr}_{\mathrm{aux}}[U^\dagger \rho_A^{\otimes M} \otimes \rho_{\mathrm{R}} \otimes \rho_{\mathrm{aux}} U], \rho_A^{\otimes N})$.

Next, we have *symmetry*, which only applies to local cloners, and requires that all output clones must be of the same quality:

$$F_{\mathsf{L}}^{j} = F_{\mathsf{L}}^{k}, \quad \forall j,k \in \{1,\ldots,N\}. \tag{2}$$

Finally, we have *universality*, which refers to the set of states, $\mathcal{S}$, the cloner can optimally clone. A universal cloning machine can clone all states equally well, whereas a *state-dependent* cloner restricts the family of states to be some subset of the Hilbert space. All of these are practically relevant dimensions when studying the relationship between cloning and quantum cryptography.

III. VARIATIONAL QUANTUM CRYPTANALYSIS

In this work we propose the merging of quantum machine learning with quantum cryptography and cryptanalysis. Our attack vector is given by learning to clone the quantum states used in two different families of protocols with VarQlone. Let us begin with QKD protocols by focusing on BB84. [2] In this protocol one party, say, Alice, sends single-qubit states in two orthogonal bases (for instance, the eigenstates of the Pauli X and Pauli Y matrices, $|\pm\rangle$ and $|\pm i\rangle$) to a second party, Bob, via a quantum channel that is susceptible to an eavesdropping adversary, Eve. Eve's goal is to extract the secret information sent between Alice and Bob, encoded in the states. It turns out that the optimal "individual" (or, incoherent) attack [16] on this protocol by Eve is given by cloning so called *phase-covariant* [74] states of the form

$$|\psi_{xy}(\eta)\rangle = \frac{1}{\sqrt{2}}\big(|0\rangle + e^{i\eta}|1\rangle\big). \tag{3}$$

For these states, Figs. 2(c) and 2(d) illustrate some candidate cloning circuits learned by VarQlone, compared to the optimal "analytic" circuit given in Refs. [17,75,76] [Fig. 2(b)]. For this family of states, Eve can construct a cloning machine with fidelity $F_{\mathsf{L,opt}}^{\mathrm{PC,E}} \approx 0.85$. First, circuit (c) is a circuit learned in the same circumstances as circuit (b). We see that it can achieve a higher cloning fidelity on the Rigetti Aspen-8 chip, due to the fewer number of gates required. Circuit (d) improves performance on chip with a further reduction in gate numbers, by changing in which registers the clones need to appear. As a final comment, we note that the circuit (b), while providing optimal cloning fidelities for equatorial states, is not strictly the most economical version one could use for this particular task. Specifically, one can perform a variation of phase-covariant cloning, known as *mirror* phase-covariant cloning [77], using only four entangling gates, comparable to circuit (d). We discuss this in Appendix F. Nevertheless, we chose circuit (b) to compare against to demonstrate the flexibility and reconfigurable nature of VarQlone. Now, let us now analyze the performance of one of these VarQlone-learned circuits [circuit (c) specifically] in an attack on BB84. The security criterion for the optimal individual attack indicates that if Eve has as much information about the states as Bob, then a secret key can no longer be extracted. As such, the protocol requires a critical "error rate," D_{crit}, above which Alice and Bob will detect malfeasance, and abort the protocol. The key rate which can be extracted here can be generally computed in terms of the Holevo quantity, χ:

$$R = I(A{:}B) - \min\{\chi(A{:}E_Q), \chi(B{:}E_Q)\}, \tag{4}$$

where $I(A{:}B)$ denotes the mutual information (MI) between Alice and Bob and the index Q denotes that Eve may employ general quantum strategies. The Holevo quantity is given by von Neumann entropy:

$$\chi(Q{:}E) := S(\rho_E) - \tfrac{1}{2}S\big(\rho_E^0\big) - \tfrac{1}{2}S\big(\rho_E^1\big). \tag{5}$$

In order to calculate the critical error rate, we follow the analysis in Ref. [16]. Intuitively, Eq. (4) formalizes the above statement that no key can be extracted ($R=0$) when Alice and Bob's MI is equal to the minimum value between Alice and Eve and Bob and Eve. In Eq. (5), ρ_E denotes the mixed state of Eve over all of the combinations of Alice's choice of input, and ρ_E^0 and ρ_E^1 denotes the states of Eve for the random variables that encode 0 and 1 in the protocol, respectively.

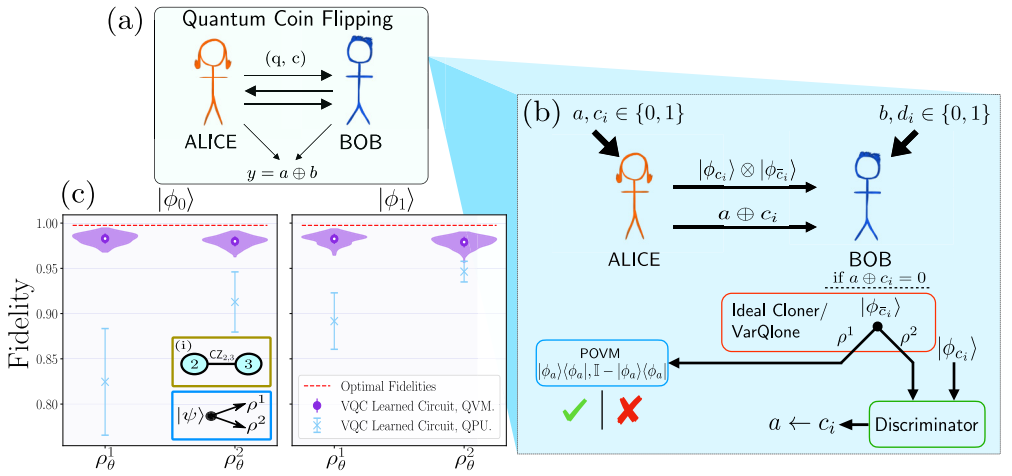

FIG. 3. Overview of cloning-based attack on the protocol of Mayers *et al.* [67], plus corresponding numerical results for VarQlone. (a) Cartoon of coin-flipping protocols, Alice and Bob send quantum (q) and/or classical (c) information to agree on a final "coin-flip" bit, y. (b) The relevant part of the protocol of Mayers *et al.*, $\mathcal{P}_1$, plus a cloning-based attack on Bob's side. He builds a cloning machine using VarQlone to produce two clones of Alice's sent states, one of which he returns, and the other is used to guess Alice's input bit, a. (c) Fidelities of each output clone, ρ^j_θ achieved using VarQlone when $(1 \to 2)$ cloning the family of states used in, $\mathcal{P}_1$. In the left (right) panel, $|\phi_0\rangle$ ($|\phi_1\rangle$) is. Both simulated (QVM, purple circles) and on Rigetti hardware (QPU, blue crosses) are shown. For the QVM (QPU) results, 256 (five) samples of each state are used to generate statistics. Violin plots show complete distribution of outcomes and error bars show the means and standard deviations. Inset (i) shows the two qubits of the Aspen-8 chip which were used, with the allowed connectivity of a CZ between them. Note an ancilla was also allowed, but VarQlone chose not to use it in this example. The corresponding learned circuit is given in Appendix F.

For circuit (c) we compute the error rate based on the BB84 protocol run in the x-y Pauli basis, and find $D_{\text{crit}} = 15.8\%$. By contrast, the optimal error rate for the ideal incoherent attack is $D^{\text{incoh}}_{\text{crit}} = 1 - F^{\text{PC,E}}_{\text{L,opt}} \approx 14.6\%$ [16]. As such, VarQlone is able to learn a close-to-optimal attack on BB84 but which achieves a higher fidelity on quantum hardware. For further discussion and details of the calculation, see Appendix C. Finally, we can see that circuit (d) in Fig. 2 achieves a higher still fidelity on hardware, but it does so by not using the ancilla to reduce the circuit depth. As such, it is a better circuit for purely performing cloning than either circuit (b) or (c), but does not provide an optimal attack for BB84 [16]. This is because the ancilla-free cloning machine only allows Eve to effectively perform a "one-sided" attack, i.e., she can guess Alice's BB84-encoded bit well, but without the ancilla she has limited information about Bob's bit.

Next, we move to quantum coin-flipping protocols, for which our primary example is the protocol of Mayers *et al.* [67] (which we denote $\mathcal{P}_1$ for brevity). The goal of such protocols is for two (mutually distrusting) parties (Alice and Bob) to agree on a random "coin flip" bit, y; a high-level overview of such protocols can be seen in Fig. 3(a). As a result of our analysis on $\mathcal{P}_1$, we provide a cloning-based attack to completely break it analytically (in contrast to the above BB84 protocol, whose security is maintained as long as the critical error rate is observed). We then use VarQlone to construct a candidate circuit to implement the cloning part of the attack. The relevant parts of the protocol $\mathcal{P}_1$ can be seen in Fig. 3(b). At a high level, Alice will choose bits, a, c_i and sends them to Bob encoded in the states in Eq. (6). Bob's attack consists of running a cloning attack and performing a state discrimination, with the purpose of guessing Alice's input bit to the coin, a. For further details of the protocol specifics, and our attack, see Appendix D. Also shown in Fig. 3(c), are the results of VarQlone when cloning the states of $\mathcal{P}_1$, both simulated and on the Rigetti Aspen-8 chip.

Specifically, the states used in $\mathcal{P}_1$ are of the form

$$|\phi_x\rangle = \cos\phi|0\rangle + (-1)^x \sin\phi|1\rangle, \quad x \in \{0,1\} \tag{6}$$

with ϕ chosen to be equal to $\pi/18$. These two states are so-called *fixed-overlap* states, which are defined by their overlap $s := \langle\phi_0|\phi_1\rangle = \cos(\pi/9)$. This family of states was one of the original scenarios studied in the realm of approximate cloning [78], but are difficult to tackle analytically [16]. Figure 3 shows the results of cloning these states, using a circuit learned by VarQlone. We also show in this figure a high-level overview of a quantum coin-flipping protocol and further details of the protocol, $\mathcal{P}_1$.

The key quantity in quantum coin-flipping protocols between two parties, Alice and Bob, is the *bias* on the coin that can be achieved by a cheating party (we assume this is Bob in this work). For example, a bias of $\epsilon = 0.1$ indicates the coin has 60% probability of favoring Bob's preferred outcome, in contrast to 50% which would be achieved by a fair (unbiased) coin. We begin by first deriving the following theorem (proof is given in Appendix D 2) which gives a bias assuming a *perfect* state-dependent cloning attack on the states in Eq. (6).

Theorem 1. [Ideal Cloning Attack Bias on $\mathcal{P}_1$]

Bob can achieve a bias of $\epsilon \approx 0.27$ using a state-dependent cloning attack on the protocol, $\mathcal{P}_1$, with a single copy of Alice's state.

The above theorem is proven using the Holevo-Helstrom [79,80] bound. From this, we then have the following corollary, by plugging in the VarQlone-learned circuit for cloning the states of $\mathcal{P}_1$ to build the full attack:

Corollary 1. [VarQlone Attack Bias on $\mathcal{P}_1$]
Bob can achieve a bias of $\epsilon \approx 0.29$ using a state-dependent VarQlone attack on the protocol, $\mathcal{P}_1$, with a single copy of Alice's state.

However, this applies only for a *single* round of the protocol, whereas $\mathcal{P}_1$, which was originally described [67] to contain k rounds. If this attack is repeated on all k rounds, we show in Appendix D 2 how this completely breaks the protocol (in other words, the bias in the quantum coin can be made to approach $1/2$ as $k \to \infty$).

As a final note on coin-flipping protocols, in Appendix D we discuss a second quantum coin flipping protocol (that of Aharonov *et al.* [68], denoted $\mathcal{P}_2$), and derive similar cloning-based attacks on it. This protocol uses the two states in Eq. (6) plus their orthogonal counterparts, for an alternative choice of the angle ϕ. In the next section, we use these states to demonstrate $1 \to 3$ and $2 \to 4$ cloning, and to highlight some features of the VarQlone algorithm.

IV. THE ALGORITHM: VARIATIONAL QUANTUM CLONING

Now we outline details of our variational quantum cloning algorithm, which generated the cloning-based attacks in the previous section. To reiterate, our motivation is to find short-depth circuits to clone a given family of states, and also use this toolkit to investigate state families where the optimal figure of merit is unknown.

Figure 1 illustrates how VarQlone is used in a cryptographic scenario; however, the core variational part is common to many variational algorithms. In particular, a variational method uses a parameterized state, denoted by $\rho_{\boldsymbol{\theta}}$, typically prepared by some short-depth parameterized unitary on some initial state, $\rho_{\boldsymbol{\theta}} := U(\boldsymbol{\theta})|0\rangle\langle 0|U^\dagger(\boldsymbol{\theta})$. The parameters are then optimized by minimizing (or maximizing) a *cost function*, typically a function of k observable measurements on $\rho_{\boldsymbol{\theta}}$, O_k. This resembles a classical neural network, and indeed techniques and ideas from classical machine learning can be borrowed and adapted.

We propose to use primarily local cost functions of the following functional form:

$$C_{\mathsf{loc}}^{M\to N}(\boldsymbol{\theta}) := \mathop{\mathbb{E}}_{|\psi\rangle\in\mathcal{S}} f(\mathsf{O}_{\mathsf{L}}^{\psi}, \rho_{\boldsymbol{\theta}}, M, N). \tag{7}$$

Here we use two specific realizations of the function, f, to generate the results, but in the Appendixes we provide alternatives, such as a *global* cost function, and a specific form the local cost Eq. (7) which enforces asymmetry in cloning.

Choosing f as follows:

$$f_{\mathsf{sq}} := \sum_{i=1}^{N}\left[1 - F_{\mathsf{L}}^{i}(\boldsymbol{\theta})\right]^2 + \sum_{i<j}^{N}\left[F_{\mathsf{L}}^{i}(\boldsymbol{\theta}) - F_{\mathsf{L}}^{j}(\boldsymbol{\theta})\right]^2 \tag{8}$$

results in what for brevity we refer to as the *squared* cost function, a generalization of the cost proposed in Ref. [81]. Here $F_{\mathsf{L}}^{j}(\boldsymbol{\theta}) := F_{\mathsf{L}}(|\psi\rangle\langle\psi|, \rho_{\boldsymbol{\theta}}^{j})$ is the local fidelity of the parameterized state relative to output clone j. This is generated using the observable $\mathsf{O}_{\mathsf{sq}}^{\psi} = |\psi\rangle\langle\psi|$ for the specific instance of state to be cloned from the set, $|\psi\rangle \in \mathcal{S}$. As such, we define $C_{\mathsf{sq}}^{M\to N}(\boldsymbol{\theta}) := \mathbb{E}_{|\psi\rangle\in\mathcal{S}}[f_{\mathsf{sq}}]$. By choosing f to be a *linear* function of the fidelities, we get a local cost more familiar from variational-algorithm literature [51,70,82,83]. We discuss the tradeoff between these different functional forms in the Appendixes, which turns out to be important for our specific application.

The cost functions we propose can be differentiated using the parameter shift rule [45,84], and we explicitly give their gradients in the Appendixes. For all the results described here, we use the gradient-descent-based Adam [85] optimizer, with our cost functions. This is in contrast to Ref. [81] which exclusively used gradient-free optimization approaches. As a theoretical guarantee on these cost functions, we prove notions of *faithfulness*, i.e., that small values of the cost functions imply near-optimal solutions (meaning the quantum states of the output clones are sufficiently close to optimality). As an example, we can prove the squared cost function above is ϵ-weakly faithful with respect to the Fubini-Study metric [86,87], D_{FS} (definitions are given in Appendix B 3):

Theorem 2. The squared cost function is ϵ-weakly faithful with respect to D_{FS}. If the cost is ϵ-close to its minimum, i.e.,

$$C_{\mathsf{sq}}(\boldsymbol{\theta}) - C_{\mathsf{sq}}^{\mathrm{opt}} \leqslant \epsilon, \tag{9}$$

where $C_{\mathsf{sq}}^{\mathrm{opt}}$ is the optimal theoretical cost, then

$$D_{\mathrm{FS}}\left(\rho_{\boldsymbol{\theta}}^{\psi,j}, \rho_{\mathrm{opt}}^{\psi,j}\right) \leqslant \frac{\mathcal{N}\epsilon}{2(1-F_{\mathrm{opt}})\sin(F_{\mathrm{opt}})} := f_1(\epsilon), \quad \forall|\psi\rangle \in \mathcal{S}, \forall j \in [N]. \tag{10}$$

We also remark that typically *global* cost functions are usually more favorable from the point of view of *operational meaning*. For example, in variational compilation [51], this cost function compares the closeness of two global unitaries. In this respect, local cost functions are usually used as a proxy to optimize a global cost function, meaning optimization with respect to the local function typically provides insights into the convergence of desired global properties. In contrast to many previous applications, by the nature of quantum cloning, VarQlone allows the local cost functions to have immediate operational meaning. Furthermore, in our use cases, there is a more subtle relationship between local and global optimization than in other applications, as for cloning in both cases the results can lead to different outcomes. We also revisit this discussion in Appendix B 3.

A key element in variational algorithms is the choice of Ansatz that is used in the PQC. The primary Ansatz we choose is one with a *variable* structure. This allows us to learn cloning circuits in an end-to-end manner. The idea is to optimize over both the continuous parameters of a quantum circuit, but also the gates within the circuit itself, which come from a discrete set.

The goal is to solve the following optimization problem [88]:

$$(\boldsymbol{\theta}^*, \boldsymbol{g}^*) = \arg\min_{\boldsymbol{\theta}, \boldsymbol{g}\in\mathcal{G}} C(\boldsymbol{\theta}, \boldsymbol{g}). \tag{11}$$

Such variable-structure Ansatz approaches can be broadly dubbed *quantum architecture search* (QAS) [89] to draw parallels with neural architecture search [9,90] (NAS) in classical ML. Approaches to QAS have appeared in many forms [71,88,91–94]. In this work, $\mathcal{G}$ is a gateset *pool*, from which a particular sequence, $\boldsymbol{g}$ is chosen. As a summary, to solve this

problem, we iterate over $\boldsymbol{g}$, swap out gates, and reoptimize the parameters, $\boldsymbol{\theta}$, until a minimum of the cost, $C(\boldsymbol{\theta}^*, \boldsymbol{g}^*)$ is found. This is a combination of a discrete and continuous optimization problem, where the discrete parameters are the indices of the gates in $\boldsymbol{g}$ (i.e., the circuit structure), and the continuous parameters are $\boldsymbol{\theta}$. Each time the circuit structure is changed (a subset of gates are altered), the continuous parameters are reoptimized, as in Ref. [71]. Variations of this approach have been proposed in Refs. [88,95] which could be easily incorporated, and we leave such investigation to future work. For the results shown in Fig. 2 [for $1 \to 2$ cloning phase-covariant states, Eq. (3)], we use the following three qubit gate pool:

$$\mathcal{G}_{\text{PC}} := \{\mathrm{R}_z^2(\theta), \mathrm{R}_z^3(\theta), \mathrm{R}_z^4(\theta), \mathrm{R}_x^2(\theta), \mathrm{R}_x^3(\theta), \mathrm{R}_x^4(\theta), \mathrm{R}_y^2(\theta), \mathrm{R}_y^3(\theta), \mathrm{R}_y^4(\theta), \mathrm{CZ}_{2,3}, \mathrm{CZ}_{3,4}, \mathrm{CZ}_{2,4}\}. \quad (12)$$

In order to achieve the results in Fig. 3, to attack protocol $\mathcal{P}_1$ using $1 \to 2$ state-dependent cloning, we use the following pool:

$$\mathcal{G}_{\mathcal{P}_1^{1\to2}} := \{\mathrm{R}_j^i(\theta), \mathrm{CZ}_{2,3}, \mathrm{CZ}_{3,4}\}, \quad \forall i \in \{2,3,4\}, \forall j \in \{x,y,z\}, \quad (13)$$

where R_j^i indicates the jth Pauli rotation acting on the ith qubit and CZ is the controlled-Z gate. In both cases, we use the qubits indexed 2, 3, and 4 in an Aspen-8 sublattice. Note that in the latter case, we allow only a linear, nearest-neighbor (NN) connectivity, which removes the need for inserting SWAP gates by the quantum compiler. In Appendix E, we include a more detailed discussion of the specifics of the algorithm via supplementary numerical results. These considerations include the cost functions, Ansätze, barren plateaus, sample complexity of the algorithm along with the VarQlone-based analysis of the protocols mentioned above. As a final demonstration, we extend to the more general $M \to N$ cloning [96,97], where M copies of the input state are transformed into $N > M$ output clones. Specifically, we test $1 \to 3$ and $2 \to 4$ cloning of the states used in the coin-flipping protocol of Aharonov *et al.* [68] mentioned above and the results are shown in Fig. 4, where again we are able to find high-fidelity results for these two problems. Figure 4 also shows the effect of circuit connectivity [nearest-neighbor (NN) versus fully connected (FC)] allowed in the gate pool for both of these cases. The best VarQlone-learned quantum circuits for $1 \to 2$ cloning of the states used in [67], $\mathcal{P}_1$, and $1 \to 2$, $1 \to 3$ and $2 \to 4$ cloning of the states of [68], $\mathcal{P}_2$, are given explicitly in Appendix F.

V. DISCUSSION

Quantum cloning is one of the most important ingredients not just as a tool in quantum cryptanalysis, but also with roots in foundational questions of quantum mechanics. However, given the amount of attention this field has received, a fundamental question remained elusive: how do we construct efficient, flexible, and noise-tolerant circuits to actually perform approximate or probabilistic cloning? This question is especially pertinent in the current NISQ era, where search for useful applications on small scale noisy quantum devices remains at the forefront. In this work, we attempt to answer this question by proposing variational quantum cloning

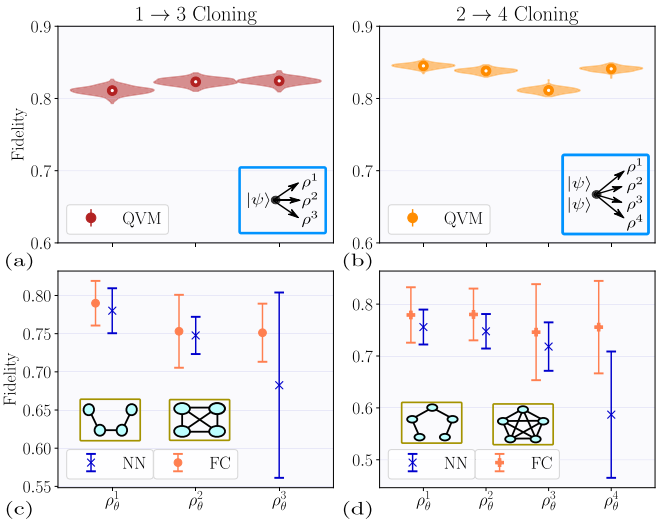

FIG. 4. Clone fidelities for optimal circuits learned by VarQlone for (a) $1 \to 3$ and (b) $2 \to 4$ cloning of the states used in the coin-flipping protocol of [68] . Mean and standard deviations of 256 samples are shown (violin plots show full distribution of fidelities), where the fidelities are computed using tomography only on the Rigetti QVM. In both cases, VarQlone is able to achieve average fidelities >80%. (c)–(d) The mean and standard deviation of the optimal fidelities found by VarQlone over 15 independent runs (15 random initial structures, $\boldsymbol{g}$) for a nearest neighbor (NN, purple) vs (d) fully connected (FC, pink) entanglement connectivity allowed in the variable structure Ansatz for $1 \to 3$ and $2 \to 4$ cloning of $\mathcal{P}_2$ states. Insets of (c)–(d) show corresponding allowed CZ gates in each example.

(VarQlone), a cloning device that utilizes the capability of short-depth quantum circuits and the power of classical computation to learn the ability to clone a state (or set of states) using the techniques of variational algorithms. This brings into view an alternative domain of performing realistic implementation of attacks on quantum cryptographic systems. We note, however, that in order to fully implement realistic and practical attacks, one must consider all aspects of the protocol environment, including, for example, the input and output mechanisms to the quantum cloner. Incorporating VarQlone into the full analyses of experimental implementation of quantum protocols, for example, [25,26], is a fruitful avenue for future work.

In conclusion, we remark that our work opens frontiers of analyzing quantum cryptographic schemes using quantum machine learning. In particular, this is applicable to secure communication schemes which are becoming increasingly relevant in quantum internet era.

ACKNOWLEDGMENTS

We thank Atul Mantri for useful comments on the manuscript. This work was supported by the Engineering and Physical Sciences Research Council (Grant No. EP/L01503X/1), EPSRC Centre for Doctoral Training in Pervasive Parallelism at the University of Edinburgh, School of Informatics, Entrapping Machines, (Grant No. FA9550-17-1-0055), and the H2020-FETOPEN Grant PHOQUSING (Grant No. 899544). We also thank Rigetti Computing for the use

of their quantum compute resources. Views expressed in this paper are those of the authors and do not reflect the views or policies of Rigetti Computing.

B.C. and M.D. devised the concept of the work; B.C. developed the numerical results; N.K., M.D., and B.C. proved the theoretical aspects of the algorithm, and M.D. developed and proved the protocol attacks. E.K. supervised the work. All authors contributed to the manuscript writing. The authors declare no conflict of interests.

APPENDIX A: CLONING WITH MULTIPLE INPUT STATES

For further background, in this section we provide some relevant equations for the cloning fidelities in some scenarios. In the main text, we discussed phase-covariant cloning, and its relevance for attacking the BB84 protocol. However, the earliest result in approximate cloning was in fact for *universal* cloning [13]. For example, if Eve was eavesdropping on a protocol using arbitrary single-qubit states, the best local fidelity she and Bob can jointly receive is $F^{\mathrm{U,B}}_{\mathsf{L},\mathrm{opt}} = F^{\mathrm{U,E}}_{\mathsf{L},\mathrm{opt}} = 5/6 < F^{\mathrm{PC,E}}_{\mathsf{L},\mathrm{opt}}$. This is still higher, however, than trivial quantum cloning strategies [16].

Now, as mentioned in the main text, we can provide multiple (M) copies of an states to the cloner and request N output approximate clones. This is referred to as $M \to N$ cloning [96,97]. Generalizing the universal cloning fidelity $(F^{\mathrm{U},j}_{\mathsf{L},\mathrm{opt}}) := F^{\mathrm{U},j}_{\mathsf{L},\mathrm{opt}}(1,2)$ to the $M \to N$ scenario, the optimal local fidelity becomes

$$F^{\mathrm{U},j}_{\mathsf{L},\mathrm{opt}}(M,N) = F^{\mathrm{U}}_{\mathsf{L},\mathrm{opt}}(M,N) = \frac{M}{N} + \frac{(N-M)(M+1)}{N(M+2)}. \tag{A1}$$

In the limit $M \to \infty$, an optimal cloning machine becomes equivalent to an quantum state estimation machine [16] for universal cloning. In the context of cryptography, $M \to N$ cloning can be modeled as having N adversaries, $E_1 \ldots, E_N$ who receive M copies of the state to be cloned. $N - M$ ancilla qubits are used to assist, so the initial state is $|\psi_A\rangle^{\otimes M} \otimes |0\rangle^{\otimes N-M}$.

Examining $M \to N$ cloning in the case of fixed-overlap states reveals an interesting feature of QCMs, which has relevance for the cost functions we define in Appendix B. For these states, the optimal *global* fidelity of cloning the two states in Eq. (6) is given by

$$F^{\mathrm{FO}}_{\mathsf{G},\mathrm{opt}}(M,N) = \tfrac{1}{2}(1 + s^{M+N} + \sqrt{1-s^{2M}}\sqrt{1-s^{2N}}). \tag{A2}$$

Interestingly, it can be shown that the state-dependent quantum cloning machine (SDQCM) which achieves this optimal *global* fidelity, does not actually saturate the optimal *local* fidelity. Instead, computing the local fidelity for the globally optimized SDQCM gives [98]

$$F^{\mathrm{FO},j}_{\mathsf{L},*}(M,N) = \frac{1}{4}\left[\frac{1+s^M}{1+s^N}(1+s^2+2s^N) + \frac{1-s^M}{1-s^N}(1+s^2-2s^N) + 2\frac{1-s^{2M}}{1-s^{2N}}(1-s^2)\right], \quad \forall j. \tag{A3}$$

In contrast, computing the optimal local fidelity for this scenario [78] (for $1 \to 2$ cloning) is

$$F^{\mathrm{FO},j}_{\mathsf{L},\mathrm{opt}} = \frac{1}{2} + \frac{\sqrt{2}}{32s}(1+s)(3-3s+\sqrt{1-2s+9s^2})\sqrt{-1+2s+3s^2+(1-s)\sqrt{1-2s+9s^2}}, \ \forall j. \tag{A4}$$

It can be shown that the *minimum* value for this expression is achieved when $s = \frac{1}{2}$ and gives $F^{\mathrm{FO},j}_{\mathsf{L},\mathrm{opt}} \approx 0.987$, which is much better than the symmetric phase-covariant cloner.

Comparing Eq. (A4) and Eq. (A3) reveals that $F^{\mathrm{FO},j}_{\mathsf{L},*}(1,2)$ is actually a *lower* bound for the optimal local fidelity, $F^{\mathrm{FO},j}_{\mathsf{L},\mathrm{opt}}$ in Eq. (A4). This point is crucially relevant in our designs for a variational cloning algorithm and affects our ability to prove faithfulness arguments.

APPENDIX B: VARIATIONAL QUANTUM CLONING: COST FUNCTIONS, GRADIENTS, AND GUARANTEES

In this Appendix, we elaborate on the details of our variational quantum cloning algorithm discussed in the main text. Before doing so, we provide a gentle introduction to the motivations, methods and mind set for developing variational algorithms.

As mentioned in the main text, the primary reason for the development of variational quantum algorithms (VQAs) [31–35] is the availability of NISQ quantum computers (through the cloud [72,99,100]). The small size of these devices and current noise rates put a speedup in, for example, factoring large prime numbers [36] out of reach, and "coherent" algorithms more generally.

VQAs have been proposed for a range of applications from learning Grover's algorithm [50] and compiling quantum circuits [51–53] to solving linear systems of equations [54–56] and even extending to the foundations of quantum mechanics [62], among others [82,101–104]. Deeper fundamental questions about the computational complexity [31,32], trainability [69,70,105–109], and noise resilience [83,110,111] of VQAs have also been considered. While all of the above are tightly related, each application and problem domain presents its own unique challenges, for example, requiring domain-specific knowledge, efficiency, interpretability of solution, etc. The tangential relationship of these algorithms to machine learning techniques also opens the door to the wealth of information and techniques available in that field [112,113]. A parallel and related line of research has focused on purely classical machine learning techniques (for example, reinforcement learning) to discover alternative quantum experiments [57–60] and quantum communication protocols [61].

The variational approach has been useful in quantum information and has been applied successfully to learn quantum algorithms. More interestingly, given the flexibility of the

method, it can learn alternate versions of quantum primitives, or even *improved* versions in some cases to achieve a particular task. For example, the work of Ref. [71] found alternative methods to compute quantum state overlap and Ref. [114] is able to learn circuits which are better suited to a given hardware.

In light of the above, let us now turn to the specifics of VarQlone. An overview of the relevant parts is given in Fig. 7. In the following sections, we highlight some important ingredients of the algorithm, including the cost functions we use, and the derivations of their corresponding gradients, and theoretical guarantees on them. The one key ingredient missing from this section is the choice of Ansatz we use in the algorithm, but we will revisit this later when discussing the numerical results in greater detail.

1. Cost functions

In the main text, we introduced a functional form for our *local* cost functions, and discussed one specific instance of it. Here we elaborate on the alternative choices of cost functions one can choose, including a global cost, a cost for asymmetric cloning and an alternative local symmetric cost function.

We begin by stating the functions, and then discussing the various ingredients and their relative advantages. The first cost (introduced in the main text) we refer to as the "squared local cost" or just "squared cost" for brevity:

$$C_{\mathsf{sq}}^{M\to N}(\boldsymbol{\theta}) := \mathop{\mathbb{E}}_{|\psi\rangle\in\mathcal{S}} \left\{ \sum_{i=1}^{N} \left[1 - F_{\mathsf{L}}^{i}(\boldsymbol{\theta})\right]^2 + \sum_{i<j}^{N} \left[F_{\mathsf{L}}^{i}(\boldsymbol{\theta}) - F_{\mathsf{L}}^{j}(\boldsymbol{\theta})\right]^2 \right\}. \tag{B1}$$

The second local cost we call the *linear* local cost or "local cost" again for brevity, given by

$$C_{\mathsf{L}}^{M\to N}(\boldsymbol{\theta}) := \mathop{\mathbb{E}}_{|\psi\rangle\in\mathcal{S}} \left[C_{\mathsf{L}}^{\psi}(\boldsymbol{\theta})\right] := \mathop{\mathbb{E}}_{|\psi\rangle\in\mathcal{S}} \left[\mathrm{Tr}\left(\mathsf{O}_{\mathsf{L}}^{\psi}\rho_{\boldsymbol{\theta}}\right)\right],$$
$$\mathsf{O}_{\mathsf{L}}^{\psi} := \mathbb{1} - \frac{1}{N}\sum_{j=1}^{N} |\psi\rangle\langle\psi|_j \otimes \mathbb{1}_{\bar{j}}, \tag{B2}$$

where $|\psi\rangle \in \mathcal{S}$ is the family of states to be cloned.

Now, these first two cost functions are related only in that they are both functions of *local* observables, i.e., the local fidelities. The third and final cost is fundamentally different to the other two, in that it instead uses global observables, and as such, we refer to it as the "global cost":

$$C_{\mathsf{G}}^{M\to N}(\boldsymbol{\theta}) := \mathop{\mathbb{E}}_{|\psi\rangle\in\mathcal{S}} \left[\mathrm{Tr}\left(\mathsf{O}_{\mathsf{G}}^{\psi}\rho_{\boldsymbol{\theta}}\right)\right], \qquad \mathsf{O}_{\mathsf{G}}^{\psi} := \mathbb{1} - |\psi\rangle\langle\psi|^{\otimes N}. \tag{B3}$$

The second local cost and our global cost functions are adapted from the literature on variational algorithms [51,70,82,83]. For compactness, we will drop the superscript $M \to N$ when the meaning is clear from context.

Now, we motivate our choices for the above cost functions. For Eq. (B1), if we restrict to the special case of $1 \to 2$ cloning (i.e., we have only two output parties, $j \in \{B, E\}$), and remove the expectation value over states, we recover the cost function used in Ref. [81]. A useful feature of this cost is that symmetry is explicitly enforced by the difference term $[F_i(\boldsymbol{\theta}) - F_j(\boldsymbol{\theta})]^2$.

In contrast, the local and global cost functions are inspired by other variational algorithm literature [51,70,82,83] where their properties have been extensively studied, particularly in relation to the phenomenon of "barren plateaus" [69,70]. It has been demonstrated that hardware-efficient Ansätze are untrainable (with either differentiable or nondifferentiable methods) using a global cost function similar to Eq. (B3), since they have exponentially vanishing gradients [84]. In contrast, local cost functions [Eq. (B2) and Eq. (B1)] are shown to be efficiently trainable with $O(\log N)$ depth hardware-efficient Ansätze [70]. We explicitly prove this property also for our local cost [Eq. (B2)] in Appendix E 1 b.

We also remark that typically global cost functions are usually more favorable from the point of view of *operational meaning*, for example, in variational compilation [51], this cost function compares the closeness of two global unitaries. In this respect, local cost functions are usually used as a proxy to optimize a global cost function.

In our case, the nature of quantum cloning allows VarQlone local cost functions to have immediate operational meaning, illustrated through the following example [using the local cost, Eq. (B2)] for $1 \to 2$ cloning:

$$C_{\mathsf{L}}^{\psi}(\boldsymbol{\theta}) = \mathrm{Tr}\left[\left(\mathbb{1} - \frac{1}{2}\sum_{j=1}^{2} |\psi\rangle\langle\psi|_j \otimes \mathbb{1}_{\bar{j}}\right)\rho_{\boldsymbol{\theta}}\right] \Rightarrow C_{\mathsf{L}}(\boldsymbol{\theta}) = 1 - \frac{1}{2}\mathbb{E}\left[F_{\mathsf{L}}\left(|\psi\rangle\langle\psi|, \rho_{\boldsymbol{\theta}}^{1}\right) + F_{\mathsf{L}}\left(|\psi\rangle\langle\psi|, \rho_{\boldsymbol{\theta}}^{2}\right)\right],$$

where $\mathbb{E}[F_{\mathsf{L}}]$ is the average fidelity [16] over the possible input states. The final expression of $C_{\mathsf{L}}(\boldsymbol{\theta})$ in the above equation follows from the expression of fidelity when one of the states is pure. Similarly, the global cost function relates to the global fidelity of the output state with respect to the input state(s).

2. Cost function gradients

In this work, we opt for a gradient-descent-based optimization approach, for which we require the efficient computation of gradients. Here we derive the analytic gradients for our cost functions. We use the local cost function, Eq. (B1) as an explicit example and the derivations for the other cost functions follow straightforwardly. As a reminder, the squared cost is given by

$$C_{\mathsf{sq}}^{M\to N}(\boldsymbol{\theta}) := \mathop{\mathbb{E}}_{|\psi\rangle\in\mathcal{S}} \left\{ \sum_{i=1}^{N} \left[1 - F_{\mathsf{L}}^{i}(\boldsymbol{\theta})\right]^2 + \sum_{i<j}^{N} \left[F_{\mathsf{L}}^{i}(\boldsymbol{\theta}) - F_{\mathsf{L}}^{j}(\boldsymbol{\theta})\right]^2 \right\}, \tag{B4}$$

where the expectation is taken over the uniform distribution. For example, in the phase-covariant cloner of the states Eq. (3), the parameters, η are sampled uniformly from the interval, $[0, 2\pi)$.

Now, the derivative of Eq. (B4), with respect to a single parameter, θ_l, is given by

$$\frac{\partial C_{\mathsf{sq}}(\boldsymbol{\theta})}{\partial \theta_l} = 2 \mathop{\mathbb{E}}_{|\psi\rangle\in\mathcal{S}} \left\{ \sum_{i=1}^{N} \left[1 - F_{\mathsf{L}}^{i}(\boldsymbol{\theta})\right] \left[-\frac{\partial F_{\mathsf{L}}^{i}(\boldsymbol{\theta})}{\partial \theta_l} \right] + \sum_{i<j}^{N} \left(F_{\mathsf{L}}^{i}(\boldsymbol{\theta}) - F_{\mathsf{L}}^{j}(\boldsymbol{\theta})\right) \left[\frac{\partial F_{\mathsf{L}}^{i}(\boldsymbol{\theta})}{\partial \theta_l} - \frac{\partial F_{\mathsf{L}}^{j}(\boldsymbol{\theta})}{\partial \theta_l} \right] \right\}.$$

We can rewrite the expression for the fidelity of the jth clone as

$$F_{\mathsf{L}}^{j}(\boldsymbol{\theta}) = \langle\psi|\rho_j(\boldsymbol{\theta})|\psi\rangle = \mathrm{Tr}[|\psi\rangle\langle\psi|\rho_j] = \mathrm{Tr}[|\psi\rangle\langle\psi|\mathrm{Tr}_{\bar{j}}(U(\boldsymbol{\theta})\rho_{\mathrm{init}}U(\boldsymbol{\theta})^{\dagger})]. \tag{B5}$$

Using the linearity of the trace, the derivative of the fidelities with respect to the parameters, θ_l, can be computed:

$$\frac{\partial F_{\mathsf{L}}^{j}(\boldsymbol{\theta})}{\partial \theta_l} = \mathrm{Tr}\left[|\psi\rangle\langle\psi|\mathrm{Tr}_{\bar{j}}\left(\frac{\partial U(\boldsymbol{\theta})\rho_{\mathrm{init}}U(\boldsymbol{\theta})^{\dagger}}{\partial \theta_l} \right) \right]. \tag{B6}$$

Now if we assume that each $U(\boldsymbol{\theta}) := U(\theta_d)U(\theta_{d-1})\cdots U(\theta_1)$ is composed of unitary gates of the form: $U(\theta_l) = \exp(-i\theta_l \Sigma_l)$, where $\Sigma_l^2 = \mathbb{1}$[1] (for example, a tensor product of Pauli operators), then from Refs. [45,84], we get

$$\frac{\partial U(\boldsymbol{\theta})\rho_{\mathrm{init}}U(\boldsymbol{\theta})^{\dagger}}{\partial \theta_l} = U^{l+\frac{\pi}{2}}(\boldsymbol{\theta})\rho_{\mathrm{init}}[U(\boldsymbol{\theta})^{l+\frac{\pi}{2}}]^{\dagger} - U^{l-\frac{\pi}{2}}(\boldsymbol{\theta})\rho_{\mathrm{init}}[U(\boldsymbol{\theta})^{l-\frac{\pi}{2}}]^{\dagger}, \tag{B7}$$

where the notation $U^{l\pm\frac{\pi}{2}}$ indicates the lth parameter has been shifted by $\pm\frac{\pi}{2}$, i.e., $U^{l\pm\frac{\pi}{2}} := U(\theta_d)U(\theta_{d-1})\cdots U(\theta_l \pm \pi/2)\cdots U(\theta_1)$. Now,

$$\begin{aligned}
\frac{\partial F_{\mathsf{L}}^{j}(\boldsymbol{\theta})}{\partial \theta_l} &= \mathrm{Tr}\left(|\psi\rangle\langle\psi|\mathrm{Tr}_{\bar{j}}\left\{ U^{l+\frac{\pi}{2}}(\boldsymbol{\theta})\rho_{\mathrm{init}}\left[U(\boldsymbol{\theta})^{l+\frac{\pi}{2}}\right]^{\dagger} \right\} \right) - \mathrm{Tr}\left(|\psi\rangle\langle\psi|\mathrm{Tr}_{\bar{j}}\left\{ U^{l-\frac{\pi}{2}}(\boldsymbol{\theta})\rho_{\mathrm{init}}\left[U(\boldsymbol{\theta})^{l-\frac{\pi}{2}}\right]^{\dagger} \right\} \right) \\
\Rightarrow \frac{\partial F_{\mathsf{L}}^{j}(\boldsymbol{\theta})}{\partial \theta_l} &= \mathrm{Tr}\left[|\psi\rangle\langle\psi|\rho_j^{l+\frac{\pi}{2}}(\boldsymbol{\theta}) \right] - \mathrm{Tr}\left[|\psi\rangle\langle\psi|\rho_j^{l-\frac{\pi}{2}}(\boldsymbol{\theta}) \right] = F_{\mathsf{L}}^{(j,l+\frac{\pi}{2})}(\boldsymbol{\theta}) - F_{\mathsf{L}}^{(j,l-\frac{\pi}{2})}(\boldsymbol{\theta}),
\end{aligned}$$

where we define $F_j^{(l\pm\frac{\pi}{2})}(\boldsymbol{\theta}) := \langle\psi|\rho_j^{l\pm\frac{\pi}{2}}(\boldsymbol{\theta})|\psi\rangle$ the fidelity of the jth clone, when prepared using a unitary whose lth parameter is shifted by $\pm\frac{\pi}{2}$, with respect to a target input state, $|\psi\rangle$.

Plugging this into the above expression, we get

$$\frac{\partial C_{\mathsf{sq}}(\boldsymbol{\theta})}{\partial \theta_l} = \mathop{2\mathbb{E}}_{|\psi\rangle\in\mathcal{S}} \left[\sum_{i<j}^{N} \left(F_{\mathsf{L}}^{i} - F_{\mathsf{L}}^{j}\right)\left[F_{\mathsf{L}}^{(i,l+\frac{\pi}{2})} - F_{\mathsf{L}}^{(i,l-\frac{\pi}{2})} - F_{\mathsf{L}}^{(j,l+\frac{\pi}{2})} + F_{\mathsf{L}}^{(j,l-\frac{\pi}{2})}\right] - \sum_{i=1}^{N} \left(1 - F_{\mathsf{L}}^{i}\right)\left[F_{\mathsf{L}}^{(i,l+\frac{\pi}{2})} - F_{\mathsf{L}}^{(i,l-\frac{\pi}{2})}\right] \right]. \tag{B8}$$

Using the same method, we can also derive the gradient of the local cost, Eq. (B2), with N output clones as

$$\frac{\partial C_{\mathsf{L}}(\boldsymbol{\theta})}{\partial \theta_l} = \mathbb{E}\left(\sum_{i=1}^{N} \left[F_{\mathsf{L}}^{i,l-\pi/2} - F_{\mathsf{L}}^{i,l+\pi/2}\right] \right). \tag{B9}$$

Finally, similar techniques result in the analytical expression of the gradient of the global cost function:

$$\frac{\partial C_{\mathsf{G}}(\boldsymbol{\theta})}{\partial \theta_l} = \mathbb{E}\left(F_{\mathsf{G}}^{l-\pi/2} - F_{\mathsf{G}}^{l+\pi/2}\right), \tag{B10}$$

where $F_{\mathsf{G}} := F(|\psi\rangle\langle\psi|^{\otimes N}, \rho_{\boldsymbol{\theta}})$ is the global fidelity between the parameterized output state and an N-fold tensor product of input states to be cloned.

3. Cost function guarantees

We would like to have theoretical guarantees about the above cost functions in order to use them. One particular desirable feature is *faithfulness* [51,54], meaning achieving the cost minimum indicates a solution to the problem in question.

Unfortunately, due to the hard limits on approximate quantum cloning, the above costs cannot have a minimum at 0, but instead at some finite value (say, $C_{\mathsf{L}}^{\mathrm{opt}}$ for the local cost). If one has knowledge of the optimal cloning fidelities for the problem at hand, then normalized cost functions with a minimum at zero can be defined. Otherwise, one must take the lowest value found to be the approximation of the cost minimum.

Despite this, we can still derive certain theoretical guarantees about them. Specifically, we consider notions of *strong* and *weak* faithfulness, relative to the error in our solution. Our goal is to provide statements about the *generalization performance* of the cost functions, by considering how close the states we output from our cloning machine are to those which would be outputted from the "optimal" cloner, relative to some metrics. In the following, we denote $\rho_{\mathrm{opt}}^{\psi,j}$ ($\rho_{\boldsymbol{\theta}}^{\psi,j}$) to be the optimal (VarQlone learned) reduced state for qubit j, for a particular input state, $|\psi\rangle$. If the superscript j is not present, we mean the global state of all clones.

[1]From Ref. [84], we actually need to assume only that Σ_l has at most two unique eigenvalues.

Definition 1 (Strong Faithfulness). A cloning cost function, C, is strongly faithful if

$$C(\boldsymbol{\theta}) = C^{\text{opt}} \Rightarrow \rho_{\boldsymbol{\theta}}^{\psi} = \rho_{\text{opt}}^{\psi}, \quad \forall |\psi\rangle \in \mathcal{S}, \tag{B11}$$

where C^{opt} is the minimum value achievable for the cost, C, according to quantum mechanics, and $\mathcal{S}$ is the given set of states to be cloned.

Definition 2 (ϵ-Weak Local Faithfulness). A local cloning cost function, C_{L}, is ϵ-weakly faithful if

$$\left|C_{\mathsf{L}}(\boldsymbol{\theta}) - C_{\mathsf{L}}^{\text{opt}}\right| \leqslant \epsilon \Rightarrow D\left(\rho_{\boldsymbol{\theta}}^{\psi,j}, \rho_{\text{opt}}^{\psi,j}\right) \leqslant f(\epsilon), \quad \forall |\psi\rangle \in \mathcal{S}, \forall j, \tag{B12}$$

where $D(\cdot,\cdot)$ is a chosen metric in the Hilbert space between the two states and f is a polynomial function.

Definition 3 (ϵ-Weak Global Faithfulness). A global cloning cost function, C_{G}, is ϵ-weakly faithful if

$$\left|C_{\mathsf{G}}(\boldsymbol{\theta}) - C_{\mathsf{G}}^{\text{opt}}\right| \leqslant \epsilon \Rightarrow D\left(\rho_{\boldsymbol{\theta}}^{\psi}, \rho_{\text{opt}}^{\psi}\right) \leqslant f(\epsilon), \quad \forall |\psi\rangle \in \mathcal{S}. \tag{B13}$$

One could also define local and global versions of strong faithfulness, but this is less interesting so we do not focus on it here. Let us begin by examining the squared local cost function. For this case, we will provide the most extensive analysis, and faithfulness proofs for the other cost functions can be derived using similar methods.

a. Squared cost function

We first can write the squared cost function as

$$C_{\mathsf{sq}}^{M\to N}(\boldsymbol{\theta}) = \frac{1}{\mathcal{N}} \int_{\mathcal{S}} \left[\sum_{j=1}^{N} [1 - F_i(\boldsymbol{\theta})]^2 + \sum_{i<j}^{N} [F_i(\boldsymbol{\theta}) - F_j(\boldsymbol{\theta})]^2 \right] d\psi, \tag{B14}$$

where the expectation of a fidelity F_i over the states in distribution $\mathcal{S}$ is defined as $\mathbb{E}[F_i] = \frac{1}{\mathcal{N}} \int_{\mathcal{S}} F_i d\psi$, with the normalization condition being $\mathcal{N} = \int_{\mathcal{S}} d\psi$. For qubit states, if the normalization is over the entire Bloch sphere in $SU(2)$, then $\mathcal{N} = 4\pi$. For notation simplicity, we herein denote the $C_{\mathsf{sq}}^{M\to N}(\boldsymbol{\theta})$ as $C_{\mathsf{sq}}(\boldsymbol{\theta})$.

We begin with a proof of the how the cost function is *strongly* faithful.

1. Strong Faithfulness:

Theorem 3. The squared local cost function is locally strongly faithful, i.e.,

$$C_{\mathsf{sq}}(\boldsymbol{\theta}) = C_{\mathsf{sq}}^{\text{opt}} \Rightarrow \rho_{\boldsymbol{\theta}}^{\psi,j} = \rho_{\text{opt}}^{\psi,j} \quad \forall |\psi\rangle \in \mathcal{S}, \forall j \in [N]. \tag{B15}$$

Proof. The cost function $C_{\text{sq}}(\boldsymbol{\theta})$ achieves a minimum at the joint maximum of $\mathbb{E}[F_i(\boldsymbol{\theta})]$ for all $i \in [N]$. In symmetric $M \to N$ cloning, the expectation value of all the N output fidelities peak at $F_i = F_{\text{opt}}$ for all input states $|\psi\rangle$. This corresponds to a unique optimal joint state $\rho_{\text{opt}}^{\psi,j} = U_{\text{opt}} |\psi^{\otimes M}, 0^{\otimes N-M}\rangle\langle\psi^{\otimes M}, 0^{\otimes N-M}| U_{\text{opt}}^{\dagger}$ for each $|\psi\rangle \in S$, where U_{opt} is the unitary producing the the optimal state. Since the joint optimal state and the corresponding fidelities are unique for all input states in the distribution, we conclude that the cost function achieves a minimum under precisely the unique condition, i.e., $\mathbb{E}[F_j(\boldsymbol{\theta})] = F_{\text{opt}}$ for all $j \in [N]$. This condition implies that

$$\rho_{\boldsymbol{\theta}}^{\psi,j} = \rho_{\text{opt}}^{\psi,j}, \quad \forall |\psi\rangle \in S, \forall j \in [N]. \tag{B16}$$

We note that since F_{opt} is the same for all the reduced states $j \in [N]$, this implies that the optimal reduced states are all the same for a given $|\psi\rangle \in \mathcal{S}$. Thus Eq. (B16) provides the necessary guarantee that minimizing the cost function results in the corresponding circuit output being equal to the optimal cloned state for all the inputs. ■

2. Weak Faithfulness:

Computing the exact fidelities of the output states requires an infinite number of copies. In reality, we run the iteration only a finite number of times and thus our cost function can only reach the optimal cost up to some precision. This is also relevant when running the circuit on devices in the NISQ era which would inherently introduce noise in the system. Thus, we can hope only to minimize the the cost function up to within some precision of the optimal cost.

Formally, we state this as the following lemma:

Lemma 1. Suppose the cost function is ϵ-close to the optimal cost in symmetric cloning

$$C_{\text{sq}}(\boldsymbol{\theta}) - C_{\text{sq}}^{\text{opt}} \leqslant \epsilon, \tag{B17}$$

then we have

$$\text{Tr}\left[\left(\rho_{\text{opt}}^{\psi,j} - \rho_{\boldsymbol{\theta}}^{\psi,j}\right)|\psi\rangle\langle\psi|\right] \leqslant \frac{\mathcal{N}\epsilon}{2(1 - F_{\text{opt}})}, \quad \forall |\psi\rangle \in \mathcal{S}, \forall j \in [N]. \tag{B18}$$

Proof. In $M \to N$ symmetric cloning, the optimal cost function value is achieved when each output clone achieves the fidelity F_{opt}. Thus, using Eq. (B1), the optimal cost function value is given by

$$C_{\text{sq}}^{\text{opt}} = N(1 - F_{\text{opt}})^2. \tag{B19}$$

The optimal cost function corresponds to all output clones having the same fidelity. Therefore, as we begin to minimize the cost $C_{\text{sq}}(\boldsymbol{\theta})$, all the output clones start to produce states with approximately same fidelity. This is explicitly enforced by taking the limit $\epsilon \to 0$, in which case the difference terms of Eq. (B1) vanish. Thus, the cost function explicitly enforces the symmetry property. Let us assume $\epsilon \to 0$, and consider the quantity $C_{\text{sq}}(\boldsymbol{\theta}) - C_{\text{sq}}^{\text{opt}}$:

$$\begin{aligned} &C_{\text{sq}}(\boldsymbol{\theta}) - C_{\text{sq}}^{\text{opt}} \\ &= \frac{1}{\mathcal{N}} \int_{\mathcal{S}} \left[\sum_{i}^{N} [1 - F_i(\boldsymbol{\theta})]^2 + \sum_{i<j}^{N} [F_i(\boldsymbol{\theta}) - F_j(\boldsymbol{\theta})]^2 \right] \\ &\quad \times d\psi - N(1 - F_{\text{opt}})^2 \\ &\underset{\epsilon\to 0}{\approx} \frac{1}{\mathcal{N}} \int_{\mathcal{S}} \left[\sum_{j}^{N} [1 - F_j(\boldsymbol{\theta})]^2 - N(1 - F_{\text{opt}})^2 \right] d\psi \\ &\approx \frac{1}{\mathcal{N}} \int_{\mathcal{S}} \left[\sum_{j}^{N} [F_{\text{opt}} - F_j(\boldsymbol{\theta})][2 - F_{\text{opt}} - F_j(\boldsymbol{\theta})] \right] d\psi \end{aligned}$$

$$\geqslant \frac{2(1-F_{\text{opt}})}{\mathcal{N}} \int_{\mathcal{S}} \left[\sum_j^N [F_{\text{opt}} - F_j(\boldsymbol{\theta})] \right] d\psi$$

$$= \frac{2(1-F_{\text{opt}})}{\mathcal{N}} \left[\sum_j^N \int_{\mathcal{S}} \text{Tr}\left[\left(\rho_{\text{opt}}^{\psi,j} - \rho_{\boldsymbol{\theta}}^{\psi,j} \right) |\psi\rangle\langle\psi| \right] d\psi \right]. \tag{B20}$$

The second line follows since F_{opt} is the same for each input state, $|\psi\rangle$. Utilizing the inequality in Eq. (B17) and Eq. (B20), we obtain

$$\sum_j^N \int_{\mathcal{S}} \text{Tr}\left[\left(\rho_{\text{opt}}^{\psi,j} - \rho_{\boldsymbol{\theta}}^{\psi,j} \right) |\psi\rangle\langle\psi| \right] d\psi \leqslant \frac{\mathcal{N}\epsilon}{2(1-F_{\text{opt}})}$$

$$\Rightarrow \text{Tr}\left[\left(\rho_{\text{opt}}^{\psi,j} - \rho_{\boldsymbol{\theta}}^{\psi,j} \right) |\psi\rangle\langle\psi| \right] \leqslant \frac{\mathcal{N}\epsilon}{2(1-F_{\text{opt}})}, \quad \forall |\psi\rangle \in \mathcal{S}, \forall j \in [N]. \tag{B21}$$

■

The above inequality allows us to quantify the closeness of the state produced by VarQlone and the unique optimal clone for any $|\psi\rangle \in \mathcal{S}$. We quantify this closeness of the states in the two popular distance measures in quantum information, the Fubini-Study (or Bures angle) distance [115] and the trace distance between the two quantum states.

Using the above lemma, we can prove the following two theorems for the squared local cost function:

Theorem 4. The squared cost function as defined Eq. (B1), is ϵ-weakly faithful with respect to the Fubini-distanc measure D_{FS}. In other words, if the squared cost function, Eq. (B1), is ϵ-close to its minimum, i.e.,

$$C_{\text{sq}}(\boldsymbol{\theta}) - C_{\text{sq}}^{\text{opt}} \leqslant \epsilon, \tag{B22}$$

where $C_{\text{sq}}^{\text{opt}} := \min_{\boldsymbol{\theta}} \sum_i^N [1 - F_i(\boldsymbol{\theta})]^2 + \sum_{i<j}^N [F_i(\boldsymbol{\theta}) - F_j(\boldsymbol{\theta})]^2 = N(1-F_{\text{opt}})^2$ is the optimal theoretical cost using fidelities produced by the ideal *symmetric* cloning machine, then the following fact holds:

$$D_{\text{FS}}\left(\rho_{\boldsymbol{\theta}}^{\psi,j}, \rho_{\text{opt}}^{\psi,j} \right) \leqslant \frac{\mathcal{N}}{2(1-F_{\text{opt}})\sin(F_{\text{opt}})} \epsilon := f_1(\epsilon), \quad \forall |\psi\rangle \in \mathcal{S}, \forall j \in [N]. \tag{B23}$$

Proof. To prove Theorem 4, we revisit and rewrite the Fubini-Study distance as [115]

$$D_{\text{FS}}(\rho, \sigma) = \arccos \sqrt{F(\rho, \sigma)} = \arccos \langle \phi | \tau \rangle, \tag{B24}$$

where $|\phi\rangle$ and $|\tau\rangle$ are the purifications of ρ and σ, respectively, which maximize the overlap. We note that $D_{\text{FS}}(\rho, \sigma)$ lies between $[0, \pi/2]$, with the value $\pi/2$ corresponding to the unique solution of $\rho = \sigma$. Since this distance is a metric, it obeys the triangle's inequality, i.e., for any three states ρ, σ, and δ,

$$D_{\text{FS}}(\rho, \sigma) \leqslant D_{\text{FS}}(\rho, \delta) + D_{\text{FS}}(\sigma, \delta). \tag{B25}$$

Rewriting the result of Lemma 1 in terms of fidelity for each $|\psi\rangle \in \mathcal{S}$ and correspondingly in terms of Fubini-Study distance using Eq. (B24) is

$$F\left(\rho_{\text{opt}}^{\psi,j}, |\psi\rangle \right) - F\left(\rho_{\boldsymbol{\theta}}^{\psi,j}, |\psi\rangle \right) \leqslant \epsilon' \Rightarrow \cos^2\left[D_{\text{FS}}\left(\rho_{\text{opt}}^{\psi,j}, |\psi\rangle \right) \right] - \cos^2\left[D_{\text{FS}}\left(\rho_{\boldsymbol{\theta}}^{\psi,j}, |\psi\rangle \right) \right] \leqslant \epsilon', \tag{B26}$$

where $\epsilon' = \mathcal{N}\epsilon/2(1-F_{\text{opt}})$. Let us denote $D_{\pm}^{\psi} = D_{\text{FS}}(\rho_{\text{opt}}^{\psi,j}, |\psi\rangle) \pm D_{\text{FS}}(\rho_{\boldsymbol{\theta}}^{\psi,j}, |\psi\rangle)$ This inequality in Eq. (B26) can be further rewritten as

$$\cos\left[D_{\text{FS}}\left(\rho_{\text{opt}}^{\psi,j}, |\psi\rangle \right) \right] - \cos\left[D_{\text{FS}}\left(\rho_{\boldsymbol{\theta}}^{\psi,j}, |\psi\rangle \right) \right] \leqslant \frac{\epsilon'}{\cos\left[D_{\text{FS}}\left(\rho_{\text{opt}}^{\psi,j}, |\psi\rangle \right) \right] + \cos\left[D_{\text{FS}}\left(\rho_{\boldsymbol{\theta}}^{\psi,j}, |\psi\rangle \right) \right]},$$

$$\cos\left[D_{\text{FS}}\left(\rho_{\text{opt}}^{\psi,j}, |\psi\rangle \right) \right] - \cos\left[D_{\text{FS}}\left(\rho_{\boldsymbol{\theta}}^{\psi,j}, |\psi\rangle \right) \right] \lessapprox \frac{\epsilon'}{2\cos\left[D_{\text{FS}}\left(\rho_{\text{opt}}^{\psi,j}, |\psi\rangle \right) \right]},$$

$$2\sin\left(\frac{D_+^{\psi}}{2} \right) \sin\left(\frac{D_-^{\psi}}{2} \right) \leqslant \frac{\epsilon'}{2\cos\left[D_{\text{FS}}\left(\rho_{\text{opt}}^{\psi,j}, |\psi\rangle \right) \right]}$$

$$\Rightarrow D_-^{\psi} \leqslant \frac{\epsilon'}{\sin\left[D_{\text{FS}}\left(\rho_{\text{opt}}^{\psi,j}, |\psi\rangle \right) \right]} = \frac{\mathcal{N}\epsilon}{2(1-F_{\text{opt}})\sin(F_{\text{opt}})}, \tag{B27}$$

where we have used the approximations that in the limit $\epsilon \to 0$, $D_{\text{FS}}(\rho_{\text{opt}}^{\psi,j}, |\psi\rangle) \approx D_{\text{FS}}(\rho_{\boldsymbol{\theta}}^{\psi,j}, |\psi\rangle)$ and the trigonometric identities $\cos(x-y) = 2\sin(\frac{x+y}{2})\sin(\frac{x-y}{2})$, and $\sin 2x = 2\sin x \cos x$.

Further, using the Fubini-Study metric triangle's inequality on the states $\{\rho_{\text{opt}}^{\psi,j}, \rho_{\boldsymbol{\theta}}^{\psi,j}, |\psi\rangle\}$ results in

$$D_{\text{FS}}\left(\rho_{\boldsymbol{\theta}}^{\psi,j}, |\psi\rangle \right) \leqslant D_{\text{FS}}\left(\rho_{\text{opt}}^{\psi,j}, |\psi\rangle \right) + D_{\text{FS}}\left(\rho_{\boldsymbol{\theta}}^{\psi,j}, \rho_{\text{opt}}^{\psi,j} \right). \tag{B28}$$

Combining the above inequality and Eq. (B27) results in

$$D_{\mathrm{FS}}\left(\rho_{\boldsymbol{\theta}}^{\psi,j}, \rho_{\mathrm{opt}}^{\psi,j}\right) \leqslant \frac{\mathcal{N}}{2(1-F_{\mathrm{opt}})\sin(F_{\mathrm{opt}})}\epsilon, \quad \forall |\psi\rangle \in \mathcal{S}. \tag{B29}$$

This bounds the closeness of the trained output state and the optimal output state as a function of ϵ. ■

As our second result, we prove Theorem 5 which provides a similar result, but relative to the trace distance. Contrary to the Fubini-Study distance, this result holds true only when the input states are qubits. The trace distance is a desirable bound to have since it is a strong notion of distance between quantum states, generalizing the total variation distance between classical probability distributions [115].

Theorem 5. The squared cost function, Eq. (B1), is ϵ-weakly faithful with respect to the trace distance D_{Tr}:

$$D_{\mathrm{Tr}}\left(\rho_{\mathrm{opt}}^{\psi,j}, \rho_{\boldsymbol{\theta}}^{\psi,j}\right) \leqslant g_1(\epsilon), \quad \forall j \in [N], \tag{B30}$$

where

$$g_1(\epsilon) \approx \frac{1}{2}\sqrt{4F_{\mathrm{opt}}(1-F_{\mathrm{opt}}) + \epsilon\frac{\mathcal{N}(1-2F_{\mathrm{opt}})}{2(1-F_{\mathrm{opt}})}}. \tag{B31}$$

Proof. First, we note that $F_{\mathrm{opt}} = \langle\psi|\rho_{\mathrm{opt}}^{\psi,j}|\psi\rangle$ is the same value for all input states $|\psi\rangle \in \mathcal{S}$. We apply the change of basis from $|\psi\rangle \to |0\rangle$ by applying the unitary $V|\psi\rangle = |0\rangle$. Then the effective change on the state $\rho_{\mathrm{opt}}^{\psi,j}$ to have a fidelity F_{opt} with the state $|0\rangle$ is, $\rho_{\mathrm{opt}}^{\psi,j} \to V\rho_{\mathrm{opt}}^{\psi,j}V^\dagger$.

We can write the state $V\rho_{\mathrm{opt}}^{\psi,j}V^\dagger$ as

$$V\rho_{\mathrm{opt}}^{\psi,j}V^\dagger = \begin{pmatrix} F_{\mathrm{opt}} & a^* \\ a & 1-F_{\mathrm{opt}} \end{pmatrix}, \tag{B32}$$

where we use the usual properties of a density matrix and $a \in \mathbb{C}$. The upper bound condition in Eq. (B48) states that $\langle\psi|\rho_{\mathrm{opt}}^{\psi,j} - \rho_{\boldsymbol{\theta}}^{\psi,j}|\psi\rangle = \langle 0|V(\rho_{\mathrm{opt}}^{\psi,j} - \rho_{\boldsymbol{\theta}}^{\psi,j})V^\dagger|0\rangle = \mathcal{N}\epsilon/2(1-F_{\mathrm{opt}}) = \epsilon'$ then becomes

$$V\rho_{\boldsymbol{\theta}}^{\psi,j}V^\dagger = \begin{pmatrix} F_{\mathrm{opt}}+\epsilon' & b^* \\ b & 1-(F_{\mathrm{opt}}+\epsilon') \end{pmatrix} \tag{B33}$$

for some $b \in \mathbb{C}$. The condition that $V\rho_{\mathrm{opt}}^{\psi,j}V^\dagger, V\rho_{\boldsymbol{\theta}}^{\psi,j}V^\dagger \geqslant 0$, i.e., they are positive, implies that

$$\begin{aligned} |a|^2 &\leqslant F_{\mathrm{opt}}(1-F_{\mathrm{opt}}) \equiv r_{F_{\mathrm{oct}}}^2, \\ |b|^2 &\leqslant (F_{\mathrm{opt}}+\epsilon')[1-(F_{\mathrm{opt}}+\epsilon')] \equiv r_{F_{\mathrm{opt}}+\epsilon'}. \end{aligned} \tag{B34}$$

The trace distance between two general qubit states is related to the positive eigenvalue of the difference of the two qubit states. Consider the eigenvalues of $V\rho_{\mathrm{opt}}^{\psi,j}V^\dagger - V\rho_{\boldsymbol{\theta}}^{\psi,j}V^\dagger$. The two eigenvalues of this matrix is $\lambda_\pm = \pm\sqrt{\epsilon'^2 + |a-b|^2}$. From this, the trace distance between the two states is

$$\begin{aligned} D_{\mathrm{Tr}}\left(V\rho_{\mathrm{opt}}^{\psi,j}V^\dagger, V\rho_{\boldsymbol{\theta}}^{\psi,j}V^\dagger\right) &= \tfrac{1}{2}\left|\left|V\rho_{\mathrm{opt}}^{\psi,j}V^\dagger - V\rho_{\boldsymbol{\theta}}^{\psi,j}V^\dagger\right|\right| \\ &= \tfrac{1}{2}|\lambda_+| = \tfrac{1}{2}\sqrt{\epsilon'^2 + |a-b|^2}. \end{aligned} \tag{B35}$$

We note that the trace distance is unitary invariant. Thus,

$$\begin{aligned} D_{\mathrm{Tr}}\left(\rho_{\mathrm{opt}}^{\psi,j}, \rho_{\boldsymbol{\theta}}^{\psi,j}\right) &= D_{\mathrm{Tr}}\left(V\rho_{\mathrm{opt}}^{\psi,j}V^\dagger, V\rho_{\boldsymbol{\theta}}^{\psi,j}V^\dagger\right) \\ &= \frac{1}{2}\sqrt{\epsilon'^2 + |a-b|^2} \\ &\leqslant \frac{1}{2}\sqrt{\epsilon'^2 + \left(r_{F_{\mathrm{oct}}} + r_{F_{\mathrm{opt}}+\epsilon'}\right)^2} \\ &\approx \frac{1}{2}\sqrt{4F_{\mathrm{opt}}(1-F_{\mathrm{opt}}) + \epsilon'(1-2F_{\mathrm{opt}})} \\ &= \frac{1}{2}\sqrt{4F_{\mathrm{opt}}(1-F_{\mathrm{opt}}) + \epsilon\frac{\mathcal{N}(1-2F_{\mathrm{opt}})}{2(1-F_{\mathrm{opt}})}}, \end{aligned} \tag{B36}$$

where we have used the inequality $|a-b|^2 \leqslant ||a|+|b||^2$ for all $a, b \in \mathbb{C}$. ■

b. Local cost function

Next, we prove analogous results for the local cost function, defined for $M \to N$ cloning to include the distribution $\mathcal{S}$ over the input states is

$$\begin{aligned} C_{\mathsf{L}}(\boldsymbol{\theta}) &:= \mathbb{E}\left[1 - \frac{1}{N}\left(\sum_{j=1}^{N} F_j(\boldsymbol{\theta})\right)\right] \\ &= 1 - \frac{1}{N\mathcal{N}}\int_{\mathcal{S}}\sum_{j=1}^{N} F_j(\boldsymbol{\theta})\, d\psi, \end{aligned} \tag{B37}$$

where $\mathcal{N} = \int_{\mathcal{S}} d\psi$ is the normalization condition. As above, we can show this cost function also exhibits strong faithfulness:

1. Strong Faithfulness:

Theorem 6. The squared local cost function is locally strongly faithful:

$$C_{\mathsf{L}}(\boldsymbol{\theta}) = C_{\mathsf{L}}^{\mathrm{opt}} \Rightarrow \rho_{\boldsymbol{\theta}}^{\psi,j} = \rho_{\mathrm{opt}}^{\psi,j}, \quad \forall |\psi\rangle \in \mathcal{S}, \forall j \in [N]. \tag{B38}$$

Proof. Similar the faithfulness arguments of the squared cost function, one can immediately see that the cost function $C_{\mathsf{L}}(\boldsymbol{\theta})$ achieves a unique minimum at the joint maximum of $\mathbb{E}[F_j(\boldsymbol{\theta})]$ for all $j \in [N]$. Thus, the minimum of $C_{\mathsf{L}}(\boldsymbol{\theta})$ corresponds to the unique optimal joint state with its unique local reduced states $\rho_{\mathrm{opt}}^{\psi,j}$ for each $j \in [N]$ for each input state $|\psi\rangle \in \mathcal{S}$. Thus the cost function achieves a minimum under precisely the unique condition, i.e., the output state is equal to the optimal clone state. ■

2. Weak Faithfulness:

Now we can also prove analogous versions of weak faithfulness. Many of the steps in the proof follow similarly to the squared cost derivations above, so we omit them for brevity where possible. As above, we first have the following lemma:

Lemma 2. Suppose the cost function is ϵ-close to the optimal cost in symmetric cloning

$$C_{\mathsf{L}}(\boldsymbol{\theta}) - C_{\mathsf{L}}^{\mathrm{opt}} \leqslant \epsilon, \tag{B39}$$

where we assume $\lim_{\epsilon\to 0}|\mathbb{E}[F_i(\boldsymbol{\theta})]-\mathbb{E}[F_j(\boldsymbol{\theta})]|\to 0, \forall i, j$, and therefore $C_{\text{opt}} := 1 - F_{\text{opt}}$. Then,

$$\text{Tr}\big[\big(\rho_{\text{opt}}^{\psi,j}-\rho_{\boldsymbol{\theta}}^{\psi,j}\big)|\psi\rangle\langle\psi|\big]\leqslant\mathcal{N}\epsilon,\quad \forall|\psi\rangle\in\mathcal{S},\forall j\in[N]. \tag{B40}$$

Now we can prove the following theorem:

Theorem 7. The local cost function, Eq. (B2), is ϵ-weakly faithful with respect to D_{FS}:

$$C_{\text{L}}(\boldsymbol{\theta})-C_{\text{L}}^{\text{opt}}\leqslant\epsilon, \tag{B41}$$

where $C_{\text{L}}^{\text{opt}} := 1 - F_{\text{opt}}$ then the following fact holds:

$$D_{\text{FS}}\big(\rho_{\boldsymbol{\theta}}^{\psi,j},\rho_{\text{opt}}^{\psi,j}\big)\leqslant\frac{\mathcal{N}\epsilon}{\sin(F_{\text{opt}})}=:f_2(\epsilon),\quad\forall|\psi\rangle\in\mathcal{S},\forall j\in[N]. \tag{B42}$$

Proof. We rewrite the Eq. (B40) in terms of the Fubini-Study distance,

$$F\big(\rho_{\text{opt}}^{\psi,j},|\psi\rangle\big)-F\big(\rho_{\boldsymbol{\theta}}^{\psi,j},|\psi\rangle\big)\leqslant\mathcal{N}\epsilon\Rightarrow\cos^2\big[D_{\text{FS}}\big(\rho_{\text{opt}}^{\psi,j},|\psi\rangle\big)\big]-\cos^2\big[D_{\text{FS}}\big(\rho_{\boldsymbol{\theta}}^{\psi,j},|\psi\rangle\big)\big]\leqslant\mathcal{N}\epsilon. \tag{B43}$$

Following the derivation in the squared cost function section, we obtain the Fubini-Study closeness as

$$D_{\text{FS}}\big(\rho_{\boldsymbol{\theta}}^{\psi,j},\rho_{\text{opt}}^{\psi,j}\big)\leqslant\frac{\mathcal{N}\epsilon}{\sin(F_{\text{opt}})},\quad\forall|\psi\rangle\in\mathcal{S},\forall j\in[N]. \tag{B44}$$

■

Finally, we have Theorem 8 relating to the trace distance. The proof follows identically to Theorem 5 so we just state the result:

Theorem 8. The local cost function, Eq. (B2), is ϵ-weakly faithful with respect to D_{Tr} on qubits:

$$D_{\text{Tr}}\big(\rho_{\text{opt}}^{\psi,j},\rho_{\boldsymbol{\theta}}^{\psi,j}\big)\leqslant\tfrac{1}{2}\sqrt{4F_{\text{opt}}(1-F_{\text{opt}})+\mathcal{N}\epsilon(1-2F_{\text{opt}})}=:g_2(\epsilon),\quad\forall j\in[N]. \tag{B45}$$

c. Global cost function

Finally, we show in the next theorems that the global cost function exhibits similar notions of faithfulness:

Theorem 9. The global cost function is globally strongly faithful, i.e.,

$$C_{\text{G}}(\boldsymbol{\theta})=C_{\text{G}}^{\text{opt}}\Rightarrow\rho_{\boldsymbol{\theta}}^{\psi}=\rho_{\text{opt}}^{\psi},\quad\forall|\psi\rangle\in\mathcal{S}. \tag{B46}$$

Proof. The global cost function $C_{\text{G}}(\boldsymbol{\theta})$ achieves the minimum value $C_{\text{G}}^{\text{opt}}$ at a unique point corresponding to $\mathbb{E}[F_{\text{G}}(\boldsymbol{\theta})]=F_{\text{G}}^{\text{opt}}$, where $F_{\text{G}}^{\text{opt}}$ corresponds to the fidelity term for $C_{\text{G}}^{\text{opt}}$. This corresponds to the unique global clone state ρ_{opt}^{ψ}. Thus the cost function, achieves a unique minimum under precisely the unique condition, i.e., the output global state is equal to the optimal clone state for all inputs in the distribution. ■

Now we provide statements of weak faithfulness, something that is much more relevant in the practical implementation of the cloning scheme using global optimization.

Lemma 3. Suppose the cost function is ϵ-close to the optimal cost in symmetric cloning

$$C_{\text{G}}(\boldsymbol{\theta})-C_{\text{G}}^{\text{opt}}\leqslant\epsilon, \tag{B47}$$

where $C_{\text{G}}^{\text{opt}} := 1 - F_{\text{G}}^{\text{opt}}$. Then

$$\text{Tr}\big[\big(\rho_{\text{opt}}^{\psi}-\rho_{\boldsymbol{\theta}}^{\psi}\big)|\psi\rangle^{\otimes 2}\langle\psi|^{\otimes 2}\big]\leqslant\mathcal{N}\epsilon,\quad\forall|\psi\rangle\in\mathcal{S}. \tag{B48}$$

Proof. The proof of Lemma 3 follows identically to Lemma 2 but with the exception that $C_{\text{G}}(\boldsymbol{\theta})-C_{\text{G}}^{\text{opt}}=\mathbb{E}[F_{\text{G}}^{\text{opt}}-F_{\text{G}}(\boldsymbol{\theta})]$. ■

The proof of Lemma 2 follows identically to Lemma 1, but with the exception that we can write $C_{\text{L}}(\boldsymbol{\theta})-C_{\text{L}}^{\text{opt}}=\mathbb{E}(F_{\text{opt}}-F(\boldsymbol{\theta}))$ in the symmetric case, assuming $F_i(\boldsymbol{\theta})\approx F_j(\boldsymbol{\theta}), \forall i\neq j\in[N]$.

Finally, we have the following theorem relating to weak faithfulness of the global cost function:

Theorem 10. Suppose the cost function is ϵ-close to the optimal cost in symmetric cloning

$$C_{\text{G}}(\boldsymbol{\theta})-C_{\text{G}}^{\text{opt}}\leqslant\epsilon, \tag{B49}$$

where $C_{\text{G}}^{\text{opt}} := 1 - F_{\text{G}}^{\text{opt}}$. Then

$$D_{\text{FS}}\big(\rho_{\boldsymbol{\theta}}^{\psi},\rho_{\text{opt}}^{\psi}\big)\leqslant\frac{\mathcal{N}\epsilon}{\sin\big(F_{\text{G}}^{\text{opt}}\big)}=:f_4(\epsilon),\quad\forall|\psi\rangle\in\mathcal{S} \tag{B50}$$

and

$$\begin{aligned}&D_{\text{Tr}}\big(\rho_{\text{opt}}^{\psi},\rho_{\boldsymbol{\theta}}^{\psi}\big)\\&\leqslant\tfrac{1}{2}\sqrt{4F_{\text{G}}^{\text{opt}}\big(1-F_{\text{G}}^{\text{opt}}\big)+\mathcal{N}\epsilon\big(1-2F_{\text{G}}^{\text{opt}}\big)}\\&=:g_4(\epsilon),\quad\forall|\psi\rangle\in\mathcal{S}.\end{aligned} \tag{B51}$$

Proof. The proof of Theorem 10 follows along the same lines as the proof of closeness of the Fubini-Study distance for standard local cost function as provided in Theorem 7 and the closeness of trace distance as provided in Theorem 8. ■

d. Global versus local faithfulness

This section explores the relationship between local and global cost function optimization for different cloners (universal, phase-covariant, etc.). In particular, we address the question of whether optimizing a cloner with a local or a global cost function also achieves an optimal solution relative to the other cost (operational meaning). If the answer is affir-

mative, we can use whichever cost exhibits the most desirable qualities and be confident they will achieve the same results. If not, we must be more careful as the choice may not lead to the optimal behavior we desire and so will be application dependent.

We note that this relationship only manifests in *symmetric* cloning, since there is no possibility to enforce asymmetry in the global cost function. As we will see in Eq. (B70), the only way to enforce asymmetry is by constructing a cost function which optimizes with respect to the local asymmetric optimal fidelities.

The tradeoff between local and global faithfulness turns out to be subtle when dealing with cloning problems, and is in contrast to similar studies in analogous variational algorithm literature. To begin, we have the following theorem:

Theorem 11. For the general case of $M \to N$ cloning, the global cost function $C_{\mathsf{G}}(\boldsymbol{\theta})$ and the local cost function $C_{\mathsf{L}}(\boldsymbol{\theta})$ satisfy the inequality

$$C_{\mathsf{L}}(\boldsymbol{\theta}) \leqslant C_{\mathsf{G}}(\boldsymbol{\theta}) \leqslant N C_{\mathsf{L}}(\boldsymbol{\theta}). \tag{B52}$$

Proof. We first prove the first part of the inequality:

$$\begin{aligned} C_{\mathsf{G}}(\boldsymbol{\theta}) - C_{\mathsf{L}}(\boldsymbol{\theta}) &= \frac{1}{\mathcal{N}} \int_{\mathcal{S}} \mathrm{Tr}\big[\big(\mathsf{O}_{\mathsf{G}}^{\psi} - \mathsf{O}_{\mathsf{L}}^{\psi}\big)\rho_{\boldsymbol{\theta}}^{\psi}\big]\, d\psi \\ &= \frac{1}{\mathcal{N}N} \int_{\mathcal{S}} \mathrm{Tr}\left[\left(\sum_{j=1}^{N} (|\psi\rangle\langle\psi|_j \otimes \mathbb{1}_{\bar{j}} - |\psi\rangle\langle\psi|_1 \otimes \cdots |\psi\rangle\langle\psi|_N)\right)\rho_{\boldsymbol{\theta}}^{\psi}\right] \geqslant 0 \\ &\Rightarrow C_{\mathsf{G}}(\boldsymbol{\theta}) \geqslant C_{\mathsf{L}}(\boldsymbol{\theta}), \end{aligned} \tag{B53}$$

where $\mathsf{O}_{\mathsf{L}}^{\psi}$ is defined in Eq. (B2), and the inequality in the second line holds because

$$\sum_{j=1}^{N} \big(|\psi\rangle\langle\psi|_j \otimes \mathbb{1}_{\bar{j}} - |\psi\rangle\langle\psi|_1 \otimes \cdots |\psi\rangle\langle\psi|_N\big) = \sum_{j=1}^{N} |\psi\rangle\langle\psi|_j \otimes (\mathbb{1}_{\bar{j}} - |\psi\rangle\langle\psi|_{\bar{j}}) \geqslant 0, \quad \forall |\psi\rangle \in \mathcal{S}. \tag{B54}$$

For the second part of the inequality, we consider the operator $N\mathsf{O}_{\mathsf{L}}^{\psi} - \mathsf{O}_{\mathsf{G}}^{\psi}$,

$$\begin{aligned} N\mathsf{O}_{\mathsf{L}}^{\psi} - \mathsf{O}_{\mathsf{G}}^{\psi} &= (N-1)\mathbb{1} - \sum_{j=1}^{N} \big(|\psi\rangle\langle\psi|_j \otimes \mathbb{1}_{\bar{j}}\big) + |\psi\rangle\langle\psi|_1 \otimes \cdots |\psi\rangle\langle\psi|_N \\ &= \sum_{j=1}^{N-1} (\mathbb{1}_j \otimes \mathbb{1}_{\bar{j}} - |\psi\rangle\langle\psi|_j \otimes \mathbb{1}_{\bar{j}}) - |\psi\rangle\langle\psi|_N \otimes \mathbb{1}_{\bar{N}} + |\psi\rangle\langle\psi|_1 \otimes \cdots |\psi\rangle\langle\psi|_N \\ &= \sum_{j=1}^{N-1} [(\mathbb{1} - |\psi\rangle\langle\psi|)_j \otimes \mathbb{1}_{\bar{j}}] - \bigotimes_{j=1}^{N-1} (\mathbb{1} - |\psi\rangle\langle\psi|)_j \otimes |\psi\rangle\langle\psi|_N \\ &= (\mathbb{1} - |\psi\rangle\langle\psi|)_1 \otimes \left(\mathbb{1}_{\bar{1}} - \bigotimes_{j=2}^{N-1} (\mathbb{1} - |\psi\rangle\langle\psi|_j) \otimes |\psi\rangle\langle\psi|_N\right) + \sum_{j=2}^{N-1} [(\mathbb{1} - |\psi\rangle\langle\psi|)_j \otimes \mathbb{1}_{\bar{j}}] \\ &\geqslant 0, \end{aligned} \tag{B55}$$

where the second last line is positive because each individual operator is positive for all $|\psi\rangle \in \mathcal{S}$. ∎

A similar inequality was proven in the work of, for example, Ref. [54]. We, however, note that the inequality proven in Theorem 11 (unlike in Ref. [54]) does not allow us make statements about the similarity of individual clones from the closeness of the global cost function and vice versa. This can be seen as follows:

$$\begin{aligned} C_{\mathsf{G}}(\boldsymbol{\theta}) - C_{\mathsf{G}}^{\mathrm{opt}} \leqslant \epsilon &\Rightarrow C_{\mathsf{L}}(\boldsymbol{\theta}) - C_{\mathsf{L}}^{\mathrm{opt}} \leqslant \epsilon - (C_{\mathsf{G}}(\boldsymbol{\theta}) - C_{\mathsf{L}}(\boldsymbol{\theta})) + \big(C_{\mathsf{L}}^{\mathrm{opt}} - C_{\mathsf{G}}^{\mathrm{opt}}\big) \\ &\Rightarrow C_{\mathsf{L}}(\boldsymbol{\theta}) - C_{\mathsf{L}}^{\mathrm{opt}} \leqslant \epsilon + \big(C_{\mathsf{L}}^{\mathrm{opt}} - C_{\mathsf{G}}^{\mathrm{opt}}\big) \\ &\not\Rightarrow C_{\mathsf{L}}(\boldsymbol{\theta}) - C_{\mathsf{L}}^{\mathrm{opt}} \leqslant \epsilon. \end{aligned} \tag{B56}$$

Here we have used the result of Theorem 11 that $C_{\mathsf{G}}(\boldsymbol{\theta}) \geqslant C_{\mathsf{L}}(\boldsymbol{\theta})$, and we note that $C_{\mathsf{L}}^{\mathrm{opt}} - C_{\mathsf{G}}^{\mathrm{opt}} \neq 0$ for all the $M \to N$ cloning. In particular, for $1 \to 2$ cloning, $C_{\mathsf{L}}^{\mathrm{opt}} = 5/6$, while $C_{\mathsf{G}}^{\mathrm{opt}} = 2/3$. This is due to the nonvanishing property of these cost functions, even at the theoretical optimal, and highlights the subtlety of the case in hand.

While we are unable to leverage generic inequalities for our purpose based on the cost functions, we can make statements in *specific* cases. In other words, by restricting the cloning problem to a specific input set of states, we can guarantee that optimizing *globally* will be sufficient to also optimize *local* figures of merit.

In particular, in the following we establish this strong and weak faithfulness guarantees for the *special cases* of universal and phase-covariant cloning by analyzing problem-specific features.

Theorem 12. The *global* cost function is *locally* strongly faithful for a universal symmetric cloner, i.e.:

$$C_{\mathsf{G}}(\boldsymbol{\theta}) = C_{\mathsf{G}}^{\text{opt}} \iff \rho_{\boldsymbol{\theta}}^{\psi,j} = \rho_{\text{opt}}^{\psi,j}, \quad \forall |\psi\rangle \in \mathcal{H}, \forall j \in \{1,\ldots,N\}. \tag{B57}$$

Proof. In the symmetric universal case, $C_{\mathsf{L}}^{\text{opt}}$ has a unique minimum when each local fidelity saturates:

$$F_{\mathsf{L}}^{\text{opt}} = \frac{M(N+2)+N-M}{N(M+2)}, \tag{B58}$$

achieved by local reduced states, $\{\rho_{\text{opt}}^{\psi,j}\}_{j=1}^{N}$. Now it has been shown that the optimal global fidelity F_{G} that can be reached [16,116] is

$$F_{\mathsf{G}}^{\text{opt}} = \frac{N!(M+1)!}{M!(N+1)!}, \tag{B59}$$

which also is the corresponding unique minimum value for $C_{\mathsf{G}}^{\text{opt}}$, achieved by some global state ρ_{opt}^{ψ}.

Finally, it was proven in Refs. [117,118] that the cloner which achieves one of these bounds is unique and also saturates the other, and therefore must also achieve the unique minimum of both global and local cost functions, $C_{\mathsf{G}}^{\text{opt}}$ and $C_{\mathsf{L}}^{\text{opt}}$. Hence, the local states which optimize $C_{\mathsf{L}}^{\text{opt}}$ must be the reduced density matrices of the global state which optimizes $C_{\mathsf{G}}^{\text{opt}}$ and so

$$\rho_{\text{opt}}^{\psi,j} := \text{Tr}_{\bar{j}}\left(\rho_{\text{opt}}^{\psi}\right), \quad \forall j. \tag{B60}$$

Thus for a universal cloner, the cost function with respect to both local and global fidelities will converge to the same minimum. ■

Now, before proving an analogous statement in the case of phase-covariant cloning, we first need the following lemma (we return to the notation of B, E, and E^* for clarity):

Lemma 4. For any $1 \to 2$ phase-covariant cloning machine which takes states $|0\rangle_B \otimes |\psi\rangle_E$ and an ancillary qubit $|A\rangle_{E^*}$ as input, where $|\psi\rangle := \frac{1}{\sqrt{2}}(|0\rangle + e^{i\theta}|1\rangle)$, and outputs a three-qubit state $|\Psi_{\text{BEE}^*}\rangle$ in the following form:

$$\begin{aligned}|\Psi_{\text{BEE}^*}\rangle = \tfrac{1}{2}[(|0,0\rangle + e^{i\phi}(\sin\eta|0,1\rangle + \cos\eta|1,0\rangle))|0\rangle_{E^*} \\ + (e^{i\phi}|1,1\rangle + (\cos\eta|0,1\rangle + \sin\eta|1,0\rangle))|1\rangle_{E^*}],\end{aligned} \tag{B61}$$

the global and local fidelities are simultaneously maximized at $\eta = \frac{\pi}{4}$ where $0 \leqslant \eta \leqslant \frac{\pi}{2}$ is the "shrinking factor."

Proof. To prove this, we follow the formalism that was adopted by Cerf *et al.* [119]. This uses the fact that a symmetric phase-covariant cloner induces a mapping of the following form [16]:

$$\begin{aligned}
&|0\rangle|0\rangle|0\rangle \to |0\rangle|0\rangle|0\rangle, \\
&|1\rangle|0\rangle|0\rangle \to (\sin\eta|0\rangle|1\rangle + \cos\eta|1\rangle|0\rangle)|0\rangle, \\
&|0\rangle|1\rangle|1\rangle \to (\cos\eta|0\rangle|1\rangle + \sin\eta|1\rangle|0\rangle)|1\rangle, \\
&|1\rangle|1\rangle|1\rangle \to |1\rangle|1\rangle|1\rangle.
\end{aligned} \tag{B62}$$

Next, we calculate the global state by tracing out the ancillary state to get $\rho_{\mathsf{G}}^{\text{opt}}$:

$$\rho_{\mathsf{G}}^{\text{opt}} = \text{Tr}_{E^*}(|\Psi_{\text{BEE}^*}\rangle\langle\Psi_{\text{BEE}^*}|) = |\Phi_1\rangle\langle\Phi_1| + |\Phi_2\rangle\langle\Phi_2|, \tag{B63}$$

where $|\Phi_1\rangle := \frac{1}{2}[|0,0\rangle + e^{i\phi}(\sin\eta|0,1\rangle + \cos\eta|1,0\rangle)]$ and $|\Phi_2\rangle := \frac{1}{2}[e^{i\phi}|1,1\rangle + (\cos\eta|0,1\rangle + \sin\eta|1,0\rangle)]$. Hence the global fidelity can be found as

$$\begin{aligned}F_{\mathsf{G}}^{\text{opt}} &= \text{Tr}\left(|\psi\rangle\langle\psi|^{\otimes 2}\rho_{\mathsf{G}}^{\text{opt}}\right) = |\langle\psi^{\otimes 2}|\Phi_1\rangle|^2 + |\langle\psi^{\otimes 2}|\Phi_2\rangle|^2 \\ &= \tfrac{1}{8}(1+\sin\eta+\cos\eta)^2.\end{aligned} \tag{B64}$$

Now, optimizing $F_{\mathsf{G}}^{\text{opt}}$ with respect η, we see that $F_{\mathsf{G}}^{\text{opt}}$ has only one extremum between $[0, \frac{\pi}{2}]$ specifically at $\eta = \frac{\pi}{4}$. We can also see that the local fidelity is also achieved for the same η and is equal to

$$F_{\mathsf{L}}^{\text{opt}} = \frac{1}{2}\left(1 + \frac{\sqrt{2}}{2}\right), \tag{B65}$$

which is the upper bound for local fidelity of the phase-covariant cloner. ■

With Lemma 4 established, we can next prove:

Theorem 13. The global cost function is locally strongly faithful for phase-covariant symmetric cloner, i.e.,

$$C_{\mathsf{G}}(\boldsymbol{\theta}) = C_{\mathsf{G}}^{\text{opt}} \iff \rho_{\boldsymbol{\theta}}^{\psi,j} = \rho_{\text{opt}}^{\psi,j}, \quad \forall |\psi\rangle \in \mathcal{S}, \forall j \in \{B,E\}, \tag{B66}$$

where $\mathcal{S}$ is the distribution corresponding to phase-covariant cloning.

Proof. Now, we have in Lemma 4 that the global and local fidelities of a phase-covariant cloner are both achieved with a cloning transformation of the form in Eq. (B62). Applying this transformation unitary to $|\psi\rangle|\Phi^+\rangle_{BE}$ (where $|\Phi^+\rangle_{BE}$ is a Bell state) leads to Cerf's formalism for cloning. Furthermore, we can observe that due to the symmetry of the problem, this transformation is unique (up to global phases) and so any optimal cloner must achieve it.

Furthermore, one can check that the ideal circuit in Fig. 2(b) does indeed produce an output in the form of Eq. (B61) once the preparation angles have been set for phase-covariant cloning. By a similar argument to the above, we can see that a variational cloning machine which achieves an optimal cost function value, i.e., $C_{\mathsf{G}}(\boldsymbol{\theta}) = C_{\mathsf{G}}^{\text{opt}}$ much also saturate the optimal cloning fidelities. Furthermore, by the uniqueness of the above transformation [Eq. (B62)] we also have that the local states of VarQlone are the same as the optimal transformation, which completes the proof. ■

4. Asymmetric cloning

As discussed in the main text, for certain applications, we require asymmetric cloning in the output states, i.e., in the $1 \to 2$ cloning, the optimal reduced states of Bob and Eve do not necessarily have the same fidelities with respect to the input states. We note that the cost functions proposed for symmetric cloning does not work for the asymmetric case because the symmetric cost functions are a monotonic function of Bob's and Eve's output state fidelities with respect to the input states, thus they always converge to the optimal fidelity values which are same for for Bob and Eve. This section provides a construction for asymmetric cost function with a desired output fidelity in one of the clones.

a. Optimal asymmetric fidelities

From Ref. [16], any universal $1 \to 2$ cloning circuit producing outputs clones for Bob and Eve must satisfy the no-cloning inequality:

$$\sqrt{\left(1-F_{\mathsf{L}}^{p,B}\right)\left(1-F_{\mathsf{L}}^{q,E}\right)} \geqslant \tfrac{1}{2} - \left(1-F_{\mathsf{L}}^{p,B}\right) - \left(1-F_{\mathsf{L}}^{q,E}\right), \tag{B67}$$

where the output clones of Bob and Eve are denote by $F_{\mathsf{L}}^{p,B}$ and $F_{\mathsf{L}}^{q,E}$ for the desired parameterizations p and q.

It can be easily verified that the fidelities that saturate the above inequality are

$$F_{\mathsf{L}}^{p,B} = 1 - \frac{p^2}{2}, \quad F_{\mathsf{L}}^{q,E} = 1 - \frac{q^2}{2}, \quad p, q \in [0,1], \tag{B68}$$

with p, q satisfy $p^2 + q^2 + pq = 1$. This implies that Eve is free to choose a desired fidelity for either clone, by varying the parameter, p. For example, suppose Eve wishes to send a clone to Bob with a particular fidelity $F_B^p = 1 - p^2/2$, then from Eq. (B67) her clone would have a corresponding fidelity:

$$F_E^p = 1 - \tfrac{1}{4}(2 - p^2 - p\sqrt{4-3p^2}). \tag{B69}$$

b. Asymmetric cost functions

From the inequality presented in the previous section, we can derive an *asymmetric* cost function for $1 \to 2$ cloning. Note that it can be generalized to arbitrary $M \to N$ cloning. This cost function for a particular input state family, $\mathcal{S}$ is then

$$\begin{aligned} &C_{\mathsf{L},\mathrm{asym}}(\boldsymbol{\theta}) \\ &:= \mathbb{E}\left[F_{\mathsf{L}}^{p,E} - F_{\mathsf{L}}^{E}(\boldsymbol{\theta})\right]^2 + \mathbb{E}\left[F_{\mathsf{L}}^{p,E} - F_{\mathsf{L}}^{E}(\boldsymbol{\theta})\right]^2 \\ &= \frac{1}{\mathcal{N}} \int_{\mathcal{S}} \left(\left[F_{\mathsf{L}}^{p,B} - F_{\mathsf{L}}^{B}(\boldsymbol{\theta})\right]^2 + \left[F_{\mathsf{L}}^{p,E} - F_{\mathsf{L}}^{E}(\boldsymbol{\theta})\right]^2\right) d\psi \end{aligned} \tag{B70}$$

with $F_{\mathsf{L}}^{p,j}$, $j \in \{B, E\}$ defined according to the conditions in Eq. (B69). We note Eve could also choose a specific fidelity for her clone, parameterized by q, $F_{\mathsf{L}}^{q,E} = 1 - q^2/2$, which would in turn determine $F_{\mathsf{L}}^{q,B}$ as above.

c. Asymmetric faithfulness

1. Strong Faithfulness:

Theorem 14. The asymmetric $1 \to 2$ local cost function is strongly faithful:

$$C_{\mathsf{L},\mathrm{asym}}(\boldsymbol{\theta}) = C_{\mathsf{L},\mathrm{asym}}^{\mathrm{opt}}(\boldsymbol{\theta}) \Rightarrow \rho_{\boldsymbol{\theta}}^{\psi,i} = \rho_{\mathrm{opt}}^{\psi,i}, \quad \forall |\psi\rangle \in \mathcal{S}, \forall i \in \{B, E\}. \tag{B71}$$

Proof. The cost function $C_{\mathsf{L},\mathrm{asym}}(\boldsymbol{\theta})$ achieves the minimum value of zero, uniquely when $F_{\mathsf{L}}^{B}(\boldsymbol{\theta}) = F_{\mathsf{L}}^{p,B}$ and $F_{\mathsf{L}}^{E}(\boldsymbol{\theta}) = F_{\mathsf{L}}^{p,E}$ for all input states $|\psi\rangle \in \mathcal{S}$. This corresponds to the unique reduced states $\rho_{\mathrm{opt}}^{\psi,B}$ and $\rho_{\mathrm{opt}}^{\psi,E}$ for Bob and Eve. Thus the cost function, achieves a unique minimum of zero precisely when the output reduced state for Bob and Eve is equal to the optimal clones for all inputs in $\mathcal{S}$. ∎

2. Weak Faithfulness:

Returning again to ϵ-weak faithfulness, we get similar results as in the symmetric case above:

Theorem 15. The asymmetric cost function, Eq. (B70), is ϵ-weakly faithful with respect to D_{FS}:

$$C_{\mathsf{L},\mathrm{asym}}(\boldsymbol{\theta}) - C_{\mathsf{L},\mathrm{asym}}^{\mathrm{opt}} \leqslant \epsilon, \tag{B72}$$

where $C_{\mathsf{L},\mathrm{asym}}^{\mathrm{opt}} = 0$. Then the following fact holds for Bob's and Eve's reduced states:

$$D_{\mathrm{FS}}\left(\rho_{\boldsymbol{\theta}}^{\psi,B}, \rho_{\mathrm{opt}}^{\psi,B}\right) \leqslant \frac{\sqrt{\mathcal{N}\epsilon}}{\sin(1-p^2/2)}, \quad D_{\mathrm{FS}}\left(\rho_{\boldsymbol{\theta}}^{\psi,E}, \rho_{\mathrm{opt}}^{\psi,E}\right) \leqslant \frac{\sqrt{\mathcal{N}\epsilon}}{\sin(1-q^2/2)}. \tag{B73}$$

Furthermore, we also have the following trace distance bounds:

$$D_{\mathrm{Tr}}\left(\rho^{\psi,B}, \rho_{\boldsymbol{\theta}}^{\psi,B}\right) \leqslant \frac{1}{2}\sqrt{p^2(2-p^2) - \sqrt{\mathcal{N}\epsilon}(1-p^2)}, \quad D_{\mathrm{Tr}}\left(\rho^{\psi,E}, \rho_{\boldsymbol{\theta}}^{\psi,E}\right) \leqslant \frac{1}{2}\sqrt{q^2(2-q^2) - \sqrt{\mathcal{N}\epsilon}(1-q^2)}. \tag{B74}$$

Proof. First, we derive a similar result to Lemma 2 and Lemma 1. By expanding the term $|C_{\mathsf{L},\mathrm{asym}}(\boldsymbol{\theta}) - C_{\mathsf{L},\mathrm{opt}}|$ in terms of the corresponding output states, we obtain

$$\begin{aligned} |C_{\mathsf{L},\mathrm{asym}}(\boldsymbol{\theta}) - C_{\mathsf{L},\mathrm{opt}}| &= \left| \frac{1}{\mathcal{N}} \int_{\mathcal{S}} \left(\left[F_{\mathsf{L}}^{p,B} - F_{\mathsf{L}}^{B}(\boldsymbol{\theta})\right]^2 + \left[F_{\mathsf{L}}^{p,E} - F_{\mathsf{L}}^{E}(\boldsymbol{\theta})\right]^2\right) d\psi \right| \\ &= \frac{1}{\mathcal{N}} \int_{\mathcal{S}} \left[\left(\mathrm{Tr}\left[\left(\rho^{\psi,B} - \rho_{\boldsymbol{\theta}}^{\psi,B}\right)|\psi\rangle\langle\psi|\right]\right)^2 + \left(\mathrm{Tr}\left[\left(\rho^{\psi,E} - \rho_{\boldsymbol{\theta}}^{\psi,E}\right)|\psi\rangle\langle\psi|\right]\right)^2\right] d\psi. \end{aligned} \tag{B75}$$

Using the inequalities Eq. (B72) and Eq. (B75), we get

$$\frac{1}{\mathcal{N}}\int_{\mathcal{S}}\left(\mathrm{Tr}\left[\left(\rho_{\mathrm{opt}}^{\psi,j}-\rho_{\boldsymbol{\theta}}^{\psi,j}\right)|\psi\rangle\langle\psi|\right]\right)^2 d\psi \leqslant \epsilon \Rightarrow \mathrm{Tr}\left[\left(\rho_{\mathrm{opt}}^{\psi,j}-\rho_{\boldsymbol{\theta}}^{j}\right)|\psi\rangle\langle\psi|\right] \leqslant \sqrt{\mathcal{N}\epsilon}, \tag{B76}$$

where $j \in \{B, E\}$. Thus the above inequality holds true for the output clone states corresponding to both Bob and Eve.

Next, to derive Eq. (B73) we rewrite the Eq. (B76) in terms of the Fubini-Study metric,

$$F\left(\rho_{\mathrm{opt}}^{\psi,j}, |\psi\rangle\right) - F\left(\rho_{\boldsymbol{\theta}}^{\psi,j}, |\psi\rangle\right) \leqslant \sqrt{\mathcal{N}\epsilon} \Rightarrow \cos^2\left[D_{\mathrm{FS}}\left(\rho_{\mathrm{opt}}^{\psi,j}, |\psi\rangle\right)\right] - \cos^2\left[D_{\mathrm{FS}}\left(\rho_{\boldsymbol{\theta}}^{\psi,j}, |\psi\rangle\right)\right] \leqslant \sqrt{\mathcal{N}\epsilon}. \tag{B77}$$

Following the derivation in the squared symmetric cost function section, we obtain the Fubini-Study closeness as

$$D_{\mathrm{FS}}\left(\rho_{\boldsymbol{\theta}}^{\psi,j}, \rho_{\mathrm{opt}}^{\psi,j}\right) \leqslant \frac{\sqrt{\mathcal{N}\epsilon}}{\sin\left(F_{\mathsf{L}}^{r,j}\right)}, \quad \forall |\psi\rangle \in \mathcal{S}, \tag{B78}$$

where $F_{\mathsf{L}}^{r,j}$ is the optimal cloning fidelity corresponding to $j \in \{B, E\}$ with $r \in \{p, q\}$. Finally, plugging in the optimal asymmetric fidelities, $F_{\mathsf{L}}^{p,B} = 1 - p^2/2$, and similarly for $F_{\mathsf{L}}^{q,E}$ we arrive at

$$D_{\mathrm{FS}}\left(\rho_{\boldsymbol{\theta}}^{\psi,B}, \rho_{\mathrm{opt}}^{\psi,B}\right) \leqslant \frac{\sqrt{\mathcal{N}\epsilon}}{\sin(1-p^2/2)}, \quad D_{\mathrm{FS}}\left(\rho_{\boldsymbol{\theta}}^{\psi,E}, \rho_{\mathrm{opt}}^{\psi,E}\right) \leqslant \frac{\sqrt{\mathcal{N}\epsilon}}{\sin(1-q^2/2)}. \tag{B79}$$

Finally, to prove Eq. (B74), we follow the trace distance derivation bounds as in previous sections and obtain

$$D_{\mathrm{Tr}}\left(\rho^{\psi,j}, \rho_{\boldsymbol{\theta}}^{\psi,j}\right) \leqslant \frac{1}{2}\sqrt{4F_{\mathsf{L}}^{r,j}\left(1-F_{\mathsf{L}}^{r,j}\right) + \sqrt{\mathcal{N}\epsilon}\left(1-2F_{\mathsf{L}}^{r,j}\right)}. \tag{B80}$$

Again, plugging in the optimal fidelities for Bob and Eve completes the proof. ■

5. Sample complexity of the algorithm

VarQlone requires classical minimization of one of the cost functions $C(\boldsymbol{\theta}) := \{C_{\mathrm{sq}}(\boldsymbol{\theta}), C_{\mathsf{L}}(\boldsymbol{\theta}), C_{\mathrm{asym}}(\boldsymbol{\theta}), C_{\mathsf{G}}(\boldsymbol{\theta})\}$ to achieve the optimal cost value. In order to do so, we must be able to efficiently evaluate the cost function of choice. In our case, this can be achieved by a method to compute the fidelity between quantum states:

(1) Prepare an initial circuit with some values for $\boldsymbol{\theta}$. Input a state $|\psi\rangle \in \mathcal{S}$ into the circuit and compute the cost function value $C^{\psi}(\boldsymbol{\theta})$ using the SWAP test [120] to estimate overlap between the output clone and the input state. Let L denote the number of copies of the state $|\psi\rangle$ used to estimate $C^{\psi}(\boldsymbol{\theta})$.

(2) Compute an estimator for the true cost $C(\boldsymbol{\theta}) = \mathbb{E}_{|\psi\rangle\in\mathcal{S}}[C^{\psi}(\boldsymbol{\theta})]$ using K different states sampled from $\mathcal{S}$.

Estimating the overlap in the above steps is sufficient for our purposes, since this coincides with the fidelity when at least one of the states is a pure state:

$$F(|\psi\rangle\langle\psi|, \rho) = \langle\psi|\rho|\psi\rangle = \mathrm{Tr}(|\psi\rangle\langle\psi|\rho). \tag{B81}$$

Since VQAs are heuristic algorithms, there are no guarantees on the number of training iterations over $\boldsymbol{\theta}$ to converge to C^{opt}. However, one can at least provide guarantees on the number of samples required to estimate the cost, for a particular instance of the parameters. Since this is a necessary subroutine in the algorithm, it must be efficient.

Theorem 16. The number of samples $L \times K$ required to estimate the cost function $C(\boldsymbol{\theta})$ up to ϵ'-additive error with a success probability δ is

$$L \times K = O\left(\frac{1}{\epsilon'^2}\log\frac{2}{\delta}\right), \tag{B82}$$

where K is the number of distinct states $|\psi\rangle$ sampled uniformly at random from the distribution $\mathcal{S}$, and L is the number of copies of each input state.

Proof. We provide the proof for the cost function $C_{\mathsf{G}}(\boldsymbol{\theta})$. However, this proof extends in a straightforward manner to other cost functions. As a reminder, the global cost function is defined as

$$C_{\mathsf{G}}^{\psi}(\boldsymbol{\theta}) = 1 - \langle\phi|\rho_{\boldsymbol{\theta}}^{\psi}|\phi\rangle \;\Rightarrow\; C_{\mathsf{G}}(\boldsymbol{\theta}) = 1 - \frac{1}{\mathcal{N}}\int_{\mathcal{S}}\langle\phi|\rho_{\boldsymbol{\theta}}^{\psi}|\phi\rangle\, d\psi. \tag{B83}$$

The estimation of $C_{\mathsf{G}}^{\psi}(\boldsymbol{\theta})$ requires the computation of the overlap $\langle\phi|\rho_{\boldsymbol{\theta}}^{\psi}|\phi\rangle$. Let us denote $|\phi\rangle := |\psi\rangle^{\otimes N}$ to be an N-fold tensor product of the input state. The SWAP test proposed by Buhrman *et al.* [120] is an algorithm to compute this overlap, and the circuit is given in Fig. 5.

This test inputs the states $|\phi\rangle$ and $\rho_{\boldsymbol{\theta}}^{\psi}$ with an additional ancilla qubit $|0\rangle$, and measures the ancilla in the end in the computational basis. The probability of obtaining an outcome "1" in the measurement is

$$\mathbb{P}['1'] = p^{\psi} = \tfrac{1}{2}\left(1 - \langle\phi|\rho_{\boldsymbol{\theta}}^{\psi}|\phi\rangle\right). \tag{B84}$$

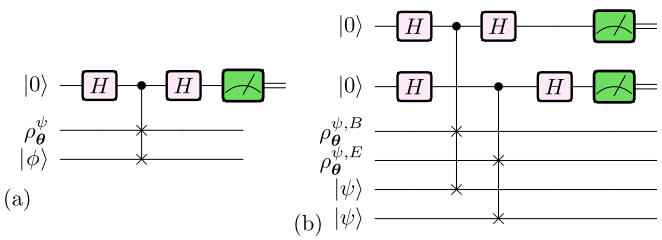

FIG. 5. SWAP test circuit illustrated for $1 \to 2$ cloning. In (a) for example, we compare the global state $\rho_{\boldsymbol{\theta}}^{\psi}$, with the state $|\phi\rangle$, where $|\phi\rangle := |\psi\rangle \otimes |\psi\rangle$ is the product state of two copies of ψ. (b) Local SWAP test with the reduced state of Bob and Eve separately. One ancilla is required for each fidelity to be computed. The generalization for N input states in $M \to N$ cloning is straightforward.

From this, we see that $C_{\mathsf{G}}^{\psi}(\boldsymbol{\theta}) = 2p^{\psi}$. To estimate this cost function value, we run the SWAP test L times and estimate the averaged number of "1" outcomes. Let us the define the estimator for cost function with L samples to be

$$\widehat{C}_{\mathsf{G},\mathrm{avg}}^{\psi}(\boldsymbol{\theta}) = \frac{1}{L}\sum_{i=1}^{L}\widehat{C}_{\mathsf{G},i}^{\psi}(\boldsymbol{\theta}), \tag{B85}$$

where $\widehat{C}_{\mathsf{G},i}^{\psi}(\boldsymbol{\theta})$ is equal to 2 if the SWAP test outcome at the ith run is "1" and is 0 otherwise. From this we can see that the expected value of $\widehat{C}_{\mathsf{G},\mathrm{avg}}^{\psi}(\boldsymbol{\theta})$ is

$$\mathbb{E}\big[\widehat{C}_{\mathsf{G},\mathrm{avg}}^{\psi}(\boldsymbol{\theta})\big] = C_{\mathsf{G}}^{\psi}(\boldsymbol{\theta}) = 2p^{\psi}. \tag{B86}$$

Now consider K different states $\{|\psi_1\rangle, \cdots |\psi_K\rangle\}$ are chosen uniformly at random from the distribution $\mathcal{S}$. The average value of the cost over K is

$$\widehat{C}_{\mathsf{G},\mathrm{avg}}(\boldsymbol{\theta}) = \frac{1}{K}\sum_{j=1}^{K}\widehat{C}_{\mathsf{G},\mathrm{avg}}^{\psi_j}(\boldsymbol{\theta}) = \frac{1}{LK}\sum_{j=1}^{K}\sum_{i=1}^{L}\widehat{C}_{\mathsf{G},i}^{\psi_j}(\boldsymbol{\theta}) \tag{B87}$$

with

$$\mathbb{E}[\widehat{C}_{\mathsf{G},\mathrm{avg}}(\boldsymbol{\theta})] = \frac{1}{K}\sum_{j=1}^{K}\frac{1}{\mathcal{N}}\int_{\mathcal{S}}\widehat{C}_{\mathsf{G},\mathrm{avg}}^{\psi}(\boldsymbol{\theta})\,d\psi = C_{\mathsf{G}}(\boldsymbol{\theta}). \tag{B88}$$

Using Höeffding's inequality [121], one can obtain a probabilistic bound on $|\widehat{C}_{\mathsf{G},\mathrm{avg}}(\boldsymbol{\theta}) - C_{\mathsf{G}}(\boldsymbol{\theta})|$,

$$\mathbb{P}[|\widehat{C}_{\mathsf{G},\mathrm{avg}}(\boldsymbol{\theta}) - C_{\mathsf{G}}(\boldsymbol{\theta})| \geqslant \epsilon'] \leqslant 2e^{-2KL\epsilon'^2}. \tag{B89}$$

Now, setting $2e^{-2KL\epsilon'^2} = \delta$ and solving for $L \times K$ gives

$$L \times K = O\left(\frac{1}{\epsilon'^2}\log\frac{2}{\delta}\right). \tag{B90}$$

■

Finally, we note that the SWAP test, in practice, is somewhat challenging to implement on NISQ devices, predominately due to the compilation overhead of compiling the 3-local controlled SWAP into the native gateset of a particular quantum hardware. Furthermore, in this case, we have a strict need for the copy of the input state $|\psi\rangle$ to be kept coherent while implementing the SWAP test, due to the equivalence between fidelity and overlap if one state is pure. This is due to the fact that for *mixed* quantum states, there is no known efficient method to compute the fidelity exactly [122] and one must resort to using bounds on it, perhaps also discovered variationally [107,123–125]. In light of this, one could use the shorter depth circuits to compute the overlap found using a variational approach similar to that implemented here [71].

APPENDIX C: QUANTUM KEY DISTRIBUTION AND CLONING ATTACKS

Here we provide further details of the cloning-based eavesdropping attacks on the BB84 protocol discussed in the main text.

First, let us clarify the types of attacks one may consider. The simplest attack by Eve is a so-called "incoherent" or individual attack, where Eve interacts with the quantum states only one at a time, and in the same fashion, and does so before the reconciliation phase of the protocol. In such attacks, the security condition states that a secret key can no longer be extracted if the fidelity of the states received by Bob and stored by Eve is the same compared to the original state sent by Alice. This criterion defines the critical value for the error rate of BB84 to be $D_{\mathrm{crit}}^{\mathrm{incoh}} = 1 - F_{\mathrm{L,opt}}^{\mathrm{PC,E}} = 14.6\%$, as mentioned in the main text.

Based on this simple fidelity calculation, one could claim that the VarQlone learned circuits for phase-covariant Fig. 2 achieve the same error rate since they saturate the same fidelity bound.

However, this criterion does not allow for comparison between a cloning machine using the ancilla, and one without. This is important since, as discussed above and in Ref. [16], a phase-covariant cloning machine with ancilla provides the optimal attack on both Alice and Bob. In contrast, one *without* an ancilla retains no information with which Eve can use to attack Bob's side of the protocol.

In order to have a better comparison we return to the general expression for the key rate in the main text:

$$R = I(A{:}B) - \min\{\chi(A:E_Q), \chi(B:E_Q)\}. \tag{C1}$$

To compute the critical error rate from this expression, D_{crit}, it is enough to calculate the Holevo quantity for Eve, set $R = 0$, and $I(A{:}B) = 1 - H(D_{\mathrm{crit}})$ and to solve the resulting equation for D_{crit}. We do this for the circuit in Fig. 2(c) only, since while the circuit in Fig. 2(d) achieves higher fidelities on the Aspen hardware, it does not actually make use of the ancillary qubit (the sequence of gates acting on it approximately resolve to the identity).

Now we compute the resulting mixed states outputted over all input states to the cloning machine, for each basis state: $\{|+\rangle, |-\rangle, |+i\rangle, |-i\rangle\}$ so ρ_E in Eq. (5) is given by

$$\rho_E := \tfrac{1}{4}\big(\rho_E^{+} + \rho_E^{-} + \rho_E^{+i} + \rho_E^{-i}\big). \tag{C2}$$

Similarly, ρ_E^0, ρ_E^1 in Eq. (5) are the mixed states encoding the symbol 0 (which have input $|+\rangle, |+i\rangle$) and the symbol 1 (which have input $|-\rangle, |-i\rangle$), so are given by

$$\rho_E^0 := \tfrac{1}{2}\big(\rho_E^{+} + \rho_E^{+i}\big) \quad \rho_E^1 := \tfrac{1}{2}\big(\rho_E^{-} + \rho_E^{-i}\big). \tag{C3}$$

Calculating the minimum Holevo quantity (denoted by $\chi_{\min}$) for the above density matrices outputted by the circuit in Fig. 2(c) numerically gives the following:

$$\begin{aligned} &1 - H(D_{\mathrm{crit}}) - \chi_{\min} = 0, \\ &\Rightarrow 1 - \chi_{\min} + [D_{\mathrm{crit}}\log_2 D_{\mathrm{crit}} \\ &\quad + (1 - D_{\mathrm{crit}})\log_2(1 - D_{\mathrm{crit}})] = 0, \\ &\Rightarrow D_{\mathrm{crit}} = 15.8\%, \end{aligned} \tag{C4}$$

which is very close to the optimal bound for the individual attack and as expected, is greater than the lower bound of 11% proved by Shor and Preskill [126]. Nevertheless as pointed out in Refs. [16,127], the same bound can be reached by a collective attack (where Eve defers any measurements until the end of the reconciliation phase and applies a general strategy to all collected states) where the individual quantum operations are still given by the optimal phase-covariant cloner. As such, the VarQlone learned attack can almost saturate the optimal collective bound as well.

APPENDIX D: QUANTUM COIN FLIPPING AND CLONING ATTACKS

With our primary objective being the improvement of practicality in attacking quantum secure communication-based protocols, this Appendix describes the explicit protocols whose security we analyze through the lens of VarQlone. In this Appendix, we focus on the primitive of *quantum coin flipping* [67,68] and the use of states which have a fixed overlap. Protocols of this nature are a nice case study for our purposes since they provide a testbed for cloning states with a fixed overlap, and in many cases explicit security analyses are missing. In particular in this work, to the best of our knowledge, we provide the first purely cloning-based attack on the protocols we analyze.

In the main text, we discussed one example of such a protocol (i.e., that of Mayers *et al.* [67]) and cloning-based attacks on it. In this section, we will introduce another such protocol (that of Aharonov *et al.* [68], alluded to in the main text), and several cloning-based attacks on it.

1. Quantum coin flipping

Let us first introduce (quantum) coin flipping in more detail. A "biased coin" in a coin-flipping protocol has one outcome more likely than the other, for example, with the following probabilities:

$$\begin{aligned} \Pr(y=0) &= 1/2+\epsilon, \\ \Pr(y=1) &= 1/2-\epsilon, \end{aligned} \tag{D1}$$

where y is a bit outputted by the coin. We can associate $y = 0$ to heads (H) and $y = 1$ to tails (T). The above coin is an ϵ-biased coin with a bias towards H. In contrast, a fair coin would correspond to $\epsilon = 0$.

It has been shown that it is impossible[2] in an information theoretic manner, to achieve a secure coin-flipping protocol with $\epsilon = 0$ in both the classical and quantum setting [67,128,129]. Furthermore, there are two notions of coin flipping studied in the literature: *weak* coin flipping (where it is *a priori* known that both parties prefer opposite outcomes) and *strong* coin flipping (where neither party knows the desired bias of the other party). In the quantum setting, the lowest possible bias achievable by any strong coin-flipping protocol is limited by $\sim$0.207 [130]. Although several protocols have been suggested for ϵ-biased strong coin flipping [2,67,68,131], the states used in them share a common structure. Here we introduce the more general form of these states which will be useful for us (a special case of which was introduced in the main text).

a. Quantum states for strong coin flipping

Multiple qubit coin flipping protocols utilize the following set of states (illustrated in Fig. 6):

$$|\phi_{x,a}\rangle = \begin{cases} |\phi_{x,0}\rangle = \cos\phi|0\rangle + (-1)^x \sin\phi|1\rangle \\ |\phi_{x,1}\rangle = \sin\phi|0\rangle + (-1)^{x\oplus 1}\cos\phi|1\rangle \end{cases}, \tag{D2}$$

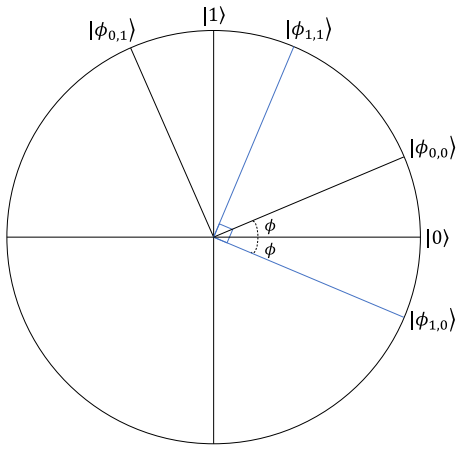

FIG. 6. States used for quantum coin flipping. The first bit represents the "basis," while the other represents one of the two orthogonal states.

where $x \in \{0, 1\}$.

Such coin-flipping protocols usually have a common structure. Alice will encode some random classical bits into some of the above states and Bob may do the same. They will then exchange classical or quantum information (or both) as part of the protocol. Attacks (attempts to bias the coin) by either party usually reduce to how much one party can learn about the classical bits of the other.

We explicitly treat two cases:

(1) The protocol of Mayers *et al.* [67] (protocol $\mathcal{P}_1$ in the main text) in which the states, $\{|\phi_{0,0}\rangle, |\phi_{1,0}\rangle\}$ are used [which have a fixed overlap $s = \cos(2\phi)$].

(2) The protocol of Aharonov *et al.* [68], which uses the full set, i.e., $\{|\phi_{x,a}\rangle\}$. We denote this protocol $\mathcal{P}_2$.

These set of states are all conveniently related through a reparameterization of the angle ϕ [98], which makes them easier to deal with mathematically.

In all strong coin-flipping protocols, the security or fairness of the final shared bit lies on the impossibility of perfect discrimination of the underlying nonorthogonal quantum states. In general, the protocol can be analyzed with either Alice or Bob being dishonest. Here we focus, for illustration, on a dishonest Bob who tries to bias the bit by cloning the nonorthogonal states sent by Alice.

For all of the below, the biases are computed assuming access to the *ideal* cloning machine (i.e., the one which clones the input states with the optimal, analytic fidelities). In Appendix E, we compare these ideal biases with those achievable using the quantum cloning machines learned by VarQlone.

2. Two-state coin-flipping protocol ($\mathcal{P}_1$)

In the main text, we gave a sketch the protocol of Mayers *et al.* [67] for a single round and a possible cloning attack on it. This was incidentally one of the first protocols proposed for strong quantum coin flipping. Here Alice utilizes the states[3]

[2]Meaning it is not possible to define a coin-flipping protocol such that *neither* party can enforce any bias.

[3]Since the value of the overlap is the only relevant quantity, the different parameterization of these states to those of Eq. (D2) does

Attack 1. Cloning Attack on $\mathcal{P}_1$ with $k = 1$.

Inputs. Random bit for Alice ($a \leftarrow_R \{0, 1\}$) and Bob ($b \leftarrow_R \{0, 1\}$). Bob receives a state $|\phi_c^i\rangle$.
Goal. A biased bit towards 0, i.e., $p(x = 0) > 1/2$.
The Attack:
1. for $i = 1, \ldots, n$:
   (a) **Step 1:** Alice announces $a \oplus c_i$. If $a \oplus c_i = 0$, Bob sends the second qubit of $|\phi_c^i\rangle$ to Alice, otherwise he sends the first qubit.
   (b) **Step 2:** Bob runs a $1 \to 2$ state-dependent cloner on the qubit he has to return to Alice, producing 2 approximate clones. He sends her one clone and keeps the other.
   (c) **Step 3:** Bob runs an optimal state discrimination on the remaining qubit and any other output of the cloner, and finds c_1 with a maximum success probability $P^{\text{opt}}_{\text{disc},\mathcal{P}_1}$. He then guesses a bit a' such that $P_{\text{succ},\mathcal{P}_1}(a' = a) := P^{\text{opt}}_{\text{disc},\mathcal{P}_1}$.
   (d) **Step 4:** If $a' \oplus b = 0$ he continues the protocol honestly and announces $b \oplus d_1$, otherwise he announces $a' \oplus d_1$. The remaining qubit on Alice's side is $|\phi_a^i\rangle$.

$|\phi_0\rangle := |\phi_{0,0}\rangle$ and $|\phi_1\rangle := |\phi_{1,0}\rangle$ such that the angle between them is $\phi := \frac{\pi}{18} \Rightarrow s := \cos(\frac{\pi}{9})$. In the following, we describe the general version of the protocol with k rounds. We also discuss the proposed attack in more detail, and prove the relevant theorems from the main text.

a. $\mathcal{P}_1$ with k rounds

With the k round version of the protocol, Alice and Bob now choose k random bits, $\{a_1, \ldots, a_k\}$ and $\{b_1, \ldots, b_k\}$, respectively. The final bit is now the XOR of input bits over all k rounds, i.e.,

$$x = \bigoplus_j a_j \oplus \bigoplus_j b_j. \tag{D3}$$

In each round $j = 1, \ldots, k$ of the protocol, and for every step $i = 1, \ldots, n$ within each round, Alice uniformly picks a random bit $c_{i,j}$ and sends the state $|\phi_c^{i,j}\rangle := |\phi_{c_{i,j}}\rangle \otimes |\phi_{\overline{c_{i,j}}}\rangle$ to Bob. Likewise, Bob uniformly picks a random bit $d_{i,j}$ and sends the state $|\phi_d^{i,j}\rangle := |\phi_{d_{i,j}}\rangle \otimes |\phi_{\overline{d_{i,j}}}\rangle$ to Alice. Hence, each party sends multiple copies of either $|\phi_0\rangle \otimes |\phi_1\rangle$ or $|\phi_1\rangle \otimes |\phi_0\rangle$.[4]

In the next step, for each j and i, Alice announces the value $a_j \oplus c_{i,j}$. If $a_j \oplus c_{i,j} = 0$, Bob returns the second state of the pair (i, j) back to Alice, and sends the first state otherwise. Similarly Bob announces $b_j \oplus d_{i,j}$, and Alice returns one of the states back to Bob accordingly. Now we come to why it is sufficient to consider only a single round in the protocol from the point of view of a cloning attack. This is because a dishonest Bob can bias the protocol if he learns about Alice's bit a_j, which he can do by guessing $c_{i,j}$ with a probability better than $1/2$. With this knowledge, Bob only needs to announce a single false $b_j \oplus d_{i,j}$ in order to cheat, and so this strategy can be deferred to the final round [67]. Hence a single round of the protocol is sufficient for analysis, and we herein drop the j index.

In the last phase of the protocol, after a and b are announced by both sides (so x can be computed by both sides), Alice measures the remaining states with the projectors, $(E_b, E_b^\perp)$ and the returned states by Bob with $(E_{\overline{a}}, E_{\overline{a}}^\perp)$ [Eq. (D4)]. She aborts the protocol if she gets the measurement result corresponding to $\perp$, and declares Bob as being dishonest. In this sense, the use of quantum states in this protocol is purely for the purpose of cheat detection.

$$E_l = |\phi_l\rangle\langle\phi_l|^{\otimes n}, \tag{D4}$$

$$E_l^\perp = \mathbb{1} - |\phi_l\rangle\langle\phi_l|^{\otimes n}, \quad l \in \{0, 1\}. \tag{D5}$$

b. A cloning attack on $\mathcal{P}_1$

Next, we present the explicit attack (illustrated in Attack 1) and calculation that can be implemented by Bob on $\mathcal{P}_1$. Without loss of generality, we assume that Bob wishes to bias the bit towards $x = 0$. For clarity, we give the attack for when Alice only sends one copy of the state ($n = 1$), but we discuss the general case in the next section.

Now we revisit the following theorem from the main text to get the success probability of the above attack:

Theorem 17. [Theorem 1 in main text.] Bob can achieve a bias of $\epsilon \approx 0.27$ using an ideal state-dependent cloning attack on the protocol, $\mathcal{P}_1$ with a single copy of Alice's state.

Proof. As mentioned in the previous section, the final measurements performed by Alice on her remaining n states, plus the n states returned to her by Bob allow her to detect his nefarious behavior. If he performed a cloning attack, the $\perp$ outcomes would be detected by Alice sometimes. We must compute both the probability that he is able to guess the value of Alice's bit a (by guessing the value of the bit c_1), and the probability that he is detected by Alice. This would provide us with Bob's final success probability in cheating, and hence the bias probability.

At the start of the attack, Bob has a product state of either $|\phi_0\rangle \otimes |\phi_1\rangle$ or $|\phi_1\rangle \otimes |\phi_0\rangle$ (but he does not know which). In step 2, depending on Alice's announced bit, Bob proceeds to clone one of the qubits, sends one copy to Alice, and keeps the other to himself. As mentioned in the main text we can assume, without loss of generality, that Alice's announced bit is 0. In this case, at this point in the attack, he has one of the following pairs: $|\phi_0\rangle\langle\phi_0| \otimes \rho_c^1$ or $|\phi_1\rangle\langle\phi_1| \otimes \rho_c^0$, where ρ_c^1 and ρ_c^0 are leftover clones for $|\phi_1\rangle$ and $|\phi_0\rangle$, respectively.

Bob must now discriminate between the following density matrices:

$$\rho_1 = |\phi_0\rangle\langle\phi_0| \otimes |\phi_1\rangle\langle\phi_1| \tag{D6}$$

$$\text{and} \quad \rho_2 = |\phi_1\rangle\langle\phi_1| \otimes \rho_c^0. \tag{D7}$$

not make a difference for our purposes. However, we note that explicit cloning unitary would be different in both cases.

[4]Note that if $c_{i,j}$ and $d_{i,j}$ are chosen independently of a_j and b_j, no information about the primary bits has been transferred.

Alternatively, if Alice announced $a \oplus c_i = 1$, he would have

$$\rho_1 = |\phi_1\rangle\langle\phi_1| \otimes |\phi_0\rangle\langle\phi_0| \tag{D8}$$

$$\text{and} \quad \rho_2 = |\phi_0\rangle\langle\phi_0| \otimes \rho_c^1. \tag{D9}$$

In either case, we have that the minimum discrimination error for two density matrices is given by the Holevo-Helstrom [79,80] bound as follows:[5]

$$P_{\text{disc}}^{\text{opt}} = \tfrac{1}{2} + \tfrac{1}{4}\|\rho_1 - \rho_2\|_{\text{Tr}} = \tfrac{1}{2} + \tfrac{1}{2}D_{\text{Tr}}(\rho_1, \rho_2). \tag{D10}$$

The ideal symmetric cloning machine for these states will have an output of the form

$$\rho_c = \alpha|\phi_0\rangle\langle\phi_0| + \beta|\phi_1\rangle\langle\phi_1| + \gamma(|\phi_0\rangle\langle\phi_1| + |\phi_1\rangle\langle\phi_0|), \tag{D11}$$

where α, β, and γ are functions of the overlap $s = \langle\phi_0|\phi_1\rangle = \cos\frac{\pi}{9}$. Now, using Eq. (D6), ρ_2 can be written as follows:

$$\begin{aligned}\rho_2 =& \alpha|\phi_1\rangle\langle\phi_1| \otimes |\phi_0\rangle\langle\phi_0| + \beta|\phi_1\rangle\langle\phi_1| \otimes |\phi_1\rangle\langle\phi_1| \\ &+ \gamma(|\phi_1\rangle\langle\phi_1| \otimes |\phi_0\rangle\langle\phi_1| + |\phi_1\rangle\langle\phi_1| \otimes |\phi_1\rangle\langle\phi_0|).\end{aligned} \tag{D12}$$

Finally by plugging in the values of the coefficients in Eq. (D11) for the optimal local cloning machine [78] and finding the eigenvalues of $\sigma := (\rho_1 - \rho_2)$, we can calculate the corresponding value for Eq. (D10), and recover the following minimum error probability:

$$P_{\text{fail},\mathcal{P}_1} = P_{\text{disc},\mathcal{P}_1}^{\text{er}} = 1 - P_{\text{disc},\mathcal{P}_1}^{\text{opt}} \approx 0.214. \tag{D13}$$

This means that Bob can successfully guess c_1 with $P_{\text{succ},\mathcal{P}_1}^1 = 78.5\%$ probability.

Now we look at the probability of a cheating Bob being detected by Alice. We note that whenever Bob guesses a successfully, the measurements $(E_b, E_b^\perp)$ will be passed with probability 1, hence we use $(E_{\bar{a}}, E_{\bar{a}}^\perp)$ where the states sent by Bob will be measured. Using Eq. (A4) with the value of overlap $s = \cos(\pi/9)$, the optimal fidelity is $F_{\text{L}} \approx 0.997$, and so the probability of Bob getting caught is at most 1%. Putting this together with Bob's guessing probability for a gives his overall success probability of 77.5%.

This implies that Bob is able to successfully create a bias of $\epsilon \approx 0.775 - 0.5 = 0.275$. ■

We also have the following corollary, for a general number of states, n exchanged, which shows the protocol can be completely broken and Bob can enforce an arbitrary bias:

Corollary 2. The probability of Bob successfully guessing a over all n copies has the property

$$\lim_{n\to\infty} P_{\text{succ},\mathcal{P}_1}^n = 1. \tag{D14}$$

Proof. If Bob repeats the above Attack 1 over all n copies, he will guess n different bits $\{a_i'\}_{i=1}^n$. He can then take a majority vote and announce b such that $a^* \oplus b = 0$, where we denote a^* as the bit he guesses in at least $\frac{n}{2} + 1$ of the rounds.

If n is even, he may have guessed a' to be 0 and 1 an equal number of times. In this case, the attack becomes indecisive and Bob is forced to guess at random. Hence we separate the success probability for even and odd n as follows:

$$P_{\text{succ},\mathcal{P}_1}^n = \begin{cases} \sum_{k=\frac{n+1}{2}}^n \binom{n}{k}(1-P_{\text{fail}})^k P_{\text{fail}}^{n-k} & n \text{ odd}, \\ \sum_{k=\frac{n}{2}+1}^n \binom{n}{k}(1-P_{\text{fail}})^k P_{\text{fail}}^{n-k} + \frac{1}{2}\binom{n}{n/2}(1-P_{\text{fail}})^{\frac{n}{2}} P_{\text{fail}}^{\frac{n}{2}} & n \text{ even}. \end{cases} \tag{D15}$$

By substituting the value of P_{fail} one can see that the function is uniformly increasing with n so $\lim_{n\to\infty} P_{\text{succ},\mathcal{P}_1}^n = 1$. ■

Although as Bob's success probability in guessing correctly increases with n, the probability of his cheating strategy getting detected by Alice will also increase. We also note that this strategy is independent of k, the number of different bits used during the protocol.

3. Four-state coin-flipping protocol ($\mathcal{P}_2$)

Another class of coin-flipping protocols are those which require all the four states in Eq. (D2). One such protocol was proposed by Aharonov *et al.* [68], where ϕ is set as $\frac{\pi}{8}$.

In protocols of this form, Alice encodes her bit in "basis information" of the family of states. More specifically, her random bit is encoded in the state $|\phi_{x,a}\rangle$. For instance, we can take $\{|\phi_{0,0}\rangle, |\phi_{1,0}\rangle\}$ to encode the bit $a = 0$ and $\{|\phi_{0,1}\rangle, |\phi_{1,1}\rangle\}$ to encode $a = 1$. The goal again is to produce a final "coin flip" $y = a \oplus b$, while ensuring that no party has biased the bit, y. A similar protocol has also been proposed using BB84 states [2] where $|\phi_{0,0}\rangle := |0\rangle$, $|\phi_{0,1}\rangle := |1\rangle$, $|\phi_{1,0}\rangle := |+\rangle$ and $|\phi_{1,1}\rangle := |-\rangle$. In this case, the states (also some protocol steps) are different but the angle between them is the same as with the states in $\mathcal{P}_2$. A fault-tolerant version of $\mathcal{P}_2$ has also been proposed in Ref. [131], which uses a generalized angle as in Eq. (D2).

The protocol proceeds as follows. First Alice sends one of the states, $|\phi_{x,a}\rangle$ to Bob. Later one of two things will happen. Either, Alice will send the bits x and a to Bob, who measures the qubit in the suitable basis to check if Alice was honest, *or* Bob is asked to return the qubit $|\phi_{x,a}\rangle$ to Alice, who measures it and verifies if it is correct. Now, example cheating strategies for Alice involve incorrect preparation of $|\phi_{x,a}\rangle$ and giving Bob the wrong information about (x, a), or for Bob in trying to determine the bits x, a from $|\phi_{x,a}\rangle$ before Alice has revealed them classically. We again focus only on Bob's strategies here to use cloning arguments. We note that the information theoretic achievable bias of $\epsilon = 0.42$ proven in Ref. [68] applies only to Alice's strategy since she has greater control of the protocol (she prepares the original state). In general, a cloning-based attack strategy by Bob will be able to achieve a lower bias, as we show. As above, Bob randomly

[5]This also is because the we assume a symmetric cloning machine for both $|\phi_0\rangle$ and $|\phi_1\rangle$. If this is not the case, the guessing probability is instead the average of the discrimination probabilities of both cases.

selects his own bit b and sends it to Alice. He then builds a QCM to clone all four states in Eq. (E21).

We next sketch the two cloning attacks on Bob's side of $\mathcal{P}_2$. Again, as with the protocol, $\mathcal{P}_1$, Bob can cheat using as much information as he gains about a and again, once Bob has performed the cloning, his strategy boils down to the problem of state discrimination. In both attacks, Bob will use a (variational) state-dependent cloning machine.

a. Cloning attacks on $\mathcal{P}_2$

In the first attack model [which we denote I; see Fig. 11(a) in Appendix E 3] Bob measures *all* the qubits outputted from the cloner to try and guess (x, a). As such, it is the *global* fidelity that will be the relevant quantity. This strategy would be useful in the first possible challenge in the protocol, where Bob is not required to send anything back to Alice. We discuss in Appendix D 3 b how the use of cloning in this type of attack can also reduce resources for Bob from a general POVM to projective measurements in the state discrimination, which may be of independent interest. The main attack here boils down to Bob measuring the global output state from his QCM using the projectors, $|v\rangle\langle v|$, $|v^\perp\rangle\langle v^\perp|$, and from this measurement, guessing a. These projectors are constructed explicitly relative to the input states using the Neumark theorem [132].

The second attack model [which we denote II; see Fig. 11(a) in Appendix E 3] is instead a *local* attack and as such will depend on the optimal local fidelity. It may also be more relevant in the scenario where Bob is required to return a quantum state to Alice. We note that Bob could also apply a global attack in this scenario but we do not consider this possibility here in order to give two contrasting examples. In the below, we compute a bias assuming he does not return a state for Alice for simplicity, and so the bias will be equivalent to his discrimination probability. The analysis could be tweaked to take a detection probability for Alice into account also. In this scenario, Bob again applies the QCM, but now he only uses one of the clones to perform state discrimination [given by the discriminator in Fig. 11(a)].

b. Attack I on $\mathcal{P}_2$

For attack I, which is a 4 state *global* attack on $\mathcal{P}_2$:

Theorem 18. [Ideal Cloning Attack (I) Bias on $\mathcal{P}_2$] Using a cloning attack on the protocol, $\mathcal{P}_2$, (in attack model I) Bob can achieve a bias:

$$\epsilon^{i}_{\mathcal{P}_2,\text{ideal}} \approx 0.35. \tag{D16}$$

We note first that this attack model (i.e., using cloning) can be considered a constructive way of implementing the optimal discrimination strategy of the states Alice is to send. In order to bias the bit, Bob needs to discriminate between the four pure states in Eq. (D2) or equivalently between the ensembles encoding $a = \{0, 1\}$, where the optimal discrimination is done via a set of POVM measurements.

However, by implementing a cloning-based attack, we can simplify the discrimination. This is because the symmetric state-dependent cloner (which is a unitary) has the interesting feature that for either case ($a = 0$ or $a = 1$), the cloner's output is a pure state in the two-qubit Hilbert space. As such, the states (after going through the QCM) can be optimally discriminated via a set of projective measurements $\{P_v, P_{v^\perp}\}$, rather than general POVMs (as would be the case if the QCM was not used). So using VarQlone to obtain optimal cloning strategies also is a means to potentially reduce resources for quantum state discrimination also. Now let us prove Theorem 18:

Proof. The attack involves the global output state of the cloning machine. For this attack we can use the fixed overlap $1 \to 2$ cloner with the global fidelity given by Eq. (A2):

$$F_{\mathsf{G}}^{\text{FO,opt}}(1,2) = \tfrac{1}{2}(1 + s^3 + \sqrt{1-s^2}\sqrt{1-s^4}) \approx 0.983, \tag{D17}$$

where $s = \sin(2\phi) = \cos(\frac{\pi}{4})$ for $\mathcal{P}_2$. Also alternatively we can use the four-state cloner which clones the two states with a fixed overlap plus their orthogonal set. For both of these cloners we are interested in the global state of the cloner which we denote as $|\psi^{1\to 2}_{x,a}\rangle$ for an input state $|\phi_{x,a}\rangle$.

In order for Bob to guess a he must discriminate between $|\phi_{0,0}\rangle$ (encoding $a = 0$) and $|\phi_{1,1}\rangle$ (encoding $a = 1$) or alternatively the pair $\{|\phi_{0,1}\rangle, |\phi_{1,0}\rangle\}$. This is since the pairs $\{|\phi_{0,0}\rangle, |\phi_{0,1}\rangle\}$ are orthogonal and $\{|\phi_{0,0}\rangle, |\phi_{1,0}\rangle\}$ both encode $a = 0$, so the only choice is to discriminate between $|\phi_{0,0}\rangle$ and $|\phi_{1,1}\rangle$. Due to the symmetry and without an ancilla, the cloner preserves the overlap between each pairs, i.e., $\langle\psi^{1\to 2}_{0,0}|\psi^{1\to 2}_{1,1}\rangle = \langle\phi_{0,0}|\phi_{1,1}\rangle = s$ (we also have $\langle\psi^{1\to 2}_{0,1}|\psi^{1\to 2}_{1,0}\rangle = s$).

Now we select the projective measurements $P_v = |v\rangle\langle v|$ and $P_{v^\perp} = |v^\perp\rangle\langle v^\perp|$ such that $\langle v|v^\perp\rangle = 0$. One can show that the discrimination probability is optimal when $|v\rangle$ and $|v^\perp\rangle$ are symmetric with respect to the target states [illustrated in Fig. 11(a)] according to the Neumark theorem. From the figure, we have that $\langle v|v^\perp\rangle = 0$ so $2\theta + 2\phi = \frac{\pi}{2} \Rightarrow \theta = \frac{\pi}{4} - \phi$. Finally, writing the cloner's states for $\{|\psi^{1\to 2}_{0,0}\rangle, |\psi^{1\to 2}_{1,1}\rangle\}$ in the basis $\{|v\rangle, |v^\perp\rangle\}$ gives

$$\begin{aligned}
|\psi^{1\to 2}_{0,0}\rangle &= \cos\left(\frac{\pi}{4} - \phi\right)|v\rangle + \sin\left(\frac{\pi}{4} - \phi\right)|v^\perp\rangle,\\
|\psi^{1\to 2}_{1,1}\rangle &= \cos\left(\frac{\pi}{4} - \phi\right)|v\rangle - \sin\left(\frac{\pi}{4} - \phi\right)|v^\perp\rangle,
\end{aligned} \tag{D18}$$

where it can be checked that $\langle\psi^{1\to 2}_{0,0}|\psi^{1\to 2}_{1,1}\rangle = \cos(\frac{\pi}{2} - 2\phi) = \sin(2\phi) = s$. Hence $|v\rangle$ and $|v^\perp\rangle$ can be explicitly derived. Note that these bases are also symmetric with respect to the other pair, i.e., $\{|\psi^{1\to 2}_{0,1}\rangle, |\psi^{1\to 2}_{1,0}\rangle\}$. Finally, the success probability of this measurement is then given by

$$P^{\text{opt},i}_{\text{disc},\mathcal{P}_2} = \tfrac{1}{2} + \tfrac{1}{2}\left|\langle\psi^{1\to 2}_{0,0}|\psi^{1\to 2}_{1,1}\rangle\right| = \tfrac{1}{2} + \tfrac{1}{2}\sin 2\phi = 0.853, \tag{D19}$$

which is the maximum cheating probability for Bob. From this, we derive the bias as

$$\epsilon^{i}_{\mathcal{P}_2,\text{ideal}} = P^{\text{opt},i}_{\text{disc},\mathcal{P}_2} - \tfrac{1}{2} = 0.353, \tag{D20}$$

which completes the proof. ■

c. Attack II on $\mathcal{P}_2$

Finally, we consider a second attack model (attack II) on the protocol, $\mathcal{P}_2$, which is in the form of a "local" attack. Here we further consider two scenarios:

(1) A cloning machine which is able to clone *all* four states $|\phi_{0,0}\rangle$, $|\phi_{1,1}\rangle$ *and* $|\phi_{0,1}\rangle$, $|\phi_{1,0}\rangle$,

(2) A cloning machine tailored only the two states, $|\phi_{0,0}\rangle$ and $|\phi_{1,1}\rangle$ (which Bob needs to discriminate between).

We focus on the former scenario, since it connects more cleanly with the VarQlone clone fidelities, but scenario 2 facilitates a more optimal attack (in the ideal scenario).

Scenario 1:

In this case, we can compute an exact discrimination probability, but it will result in a less optimal attack.

Theorem 19. [Ideal Cloning Attack (II) Bias on $\mathcal{P}_2$ in scenario 1.] Using a cloning attack on the protocol, $\mathcal{P}_2$, (in attack model II with four states) Bob can achieve a bias:

$$\epsilon^{\mathrm{II}}_{\mathcal{P}_2,\mathrm{ideal}} = 0.25. \tag{D21}$$

Proof. Considering the four states to be in the x-z plane of the Bloch sphere, the density matrices of each state can be represented as

$$\rho_{ij} = \tfrac{1}{2}\big(\mathbb{1} + m^x_{ij}\sigma_x + m^z_{ij}\sigma_z\big), \tag{D22}$$

where σ_x and σ_z are Pauli matrices and m_{ij} is a three-dimensional vector given by

$$\begin{aligned} m_{00} &:= [\sin(2\phi), 0, \cos(2\phi)],\\ m_{01} &:= [-\sin(2\phi), 0, -\cos(2\phi)],\\ m_{10} &:= [-\sin(2\phi), 0, \cos(2\phi)],\\ m_{11} &:= [\sin(2\phi), 0, -\cos(2\phi)]. \end{aligned} \tag{D23}$$

After the cloning (in the ideal case), the density matrix of each clone will become

$$\rho^c_{ij} = \tfrac{1}{2}\big(\mathbb{1} + \eta_x m^x_{ij}\sigma_x + \eta_z m^z_{ij}\sigma_z\big), \tag{D24}$$

where η_x and η_z are the shrinking factors in each direction given as follows:

$$\begin{aligned} \eta_x &= \sin^2(2\phi)\sqrt{\frac{1}{\sin^4(2\phi)+\cos^4(2\phi)}},\\ \eta_z &= \cos^2(2\phi)\sqrt{\frac{1}{\sin^4(2\phi)+\cos^4(2\phi)}}. \end{aligned} \tag{D25}$$

For the states used in $\mathcal{P}_2$, we have $\phi = \frac{\pi}{8}$ and hence $\eta_x = \eta_z := \eta = \frac{1}{\sqrt{2}}$. Again, we can return to the discrimination probability between the two ensembles encoding $a = 0$ and $a = 1$ in Eq. (D27). Here we have (let us define ρ^c to be the output clone that Bob chooses to use ($c \in \{1, 2\}$)

$$\begin{aligned} P^{\mathrm{opt,II}}_{\mathrm{disc},\mathcal{P}_2} &= \frac{1}{2} + \frac{1}{4}||\rho_{(a=0)} - \rho_{(a=1)}||_{\mathrm{Tr}}\\ &= \frac{1}{2} + \frac{1}{4}\left|\left|\frac{1}{2}\big[\big(\rho^c_{00} - \rho^c_{11}\big) + \big(\rho^c_{10} - \rho^c_{01}\big]\right|\right|_{\mathrm{Tr}}\\ &= \frac{1}{2} + \frac{1}{4}\left|\left|\frac{\eta}{4}\big(\big(m^x_{00} - m^x_{11} + m^x_{10} - m^x_{01}\big)\sigma_x\right.\right.\\ &\quad \left.\left.+\big(m^z_{00} - m^z_{11} + m^z_{10} - m^z_{01}\big)\sigma_z\right|\right|_{\mathrm{Tr}}\\ &= \frac{1}{2} + \frac{\eta\cos(2\phi)}{4}||\sigma_z||_{\mathrm{Tr}}\\ &= \frac{1}{2} + \frac{\eta\cos(2\phi)}{2} = \frac{3}{4}. \end{aligned}$$

Computing the bias in the same way as above completes the proof. ■

Scenario 2:

Here we give a bound on the success probabilities of Bob in terms of the local fidelities of the QCM where the cloning machine is only tailored to clone two fixed-overlap states. Here we rely on the fact that Bob can discriminate between the two ensembles of states (for $a = 0$, $a = 1$) with equal probabilities.

Theorem 20. The optimal discrimination probability for a cloning attack on the protocol, P_2, (in attack model II, with two states) is

$$0.619 \leqslant P^{\mathrm{opt,II}}_{\mathrm{disc},\mathcal{P}_2} \leqslant 0.823. \tag{D26}$$

Proof. For each of the input states, $|\phi_{i,j}\rangle$, in Eq. (E21), we denote ρ^c_{ij} to be a clone outputted from the QCM. Due to symmetry, we only need to consider one of the two output clones. We can now write the effective states for each encoding ($a = 0$, $a = 1$) as

$$\rho_{(a=0)} := \tfrac{1}{2}\big(\rho^c_{00} + \rho^c_{10}\big), \quad \rho_{(a=1)} := \tfrac{1}{2}\big(\rho^c_{01} + \rho^c_{11}\big). \tag{D27}$$

Dealing with these two states is sufficient since it can be shown that discriminating between these two density matrices, is equivalent to discriminating between the entire set of four states in Eq. (D2).

Again we use the discrimination probability from the Holevo-Helstrom bound:

$$P^{\mathrm{opt,II}}_{\mathrm{disc},\mathcal{P}_2} := P^{\mathrm{opt}}_{\mathrm{disc}}(\rho_{(a=0)}, \rho_{(a=1)}) := \tfrac{1}{2} + \tfrac{1}{2}D_{\mathrm{Tr}}(\rho_{(a=0)}, \rho_{(a=1)}). \tag{D28}$$

Now we have

$$\begin{aligned} D_{\mathrm{Tr}}(\rho_{(a=0)}, \rho_{(a=1)}) &= \tfrac{1}{2}||\rho_{(a=0)} - \rho_{(a=1)}||_{\mathrm{Tr}}\\ &= \tfrac{1}{2}\left|\left|\tfrac{1}{2}\big(\rho^c_{00} - \rho^c_{11}\big) + \tfrac{1}{2}\big(\rho^c_{10} - \rho^c_{01}\big)\right|\right|_{\mathrm{Tr}}\\ &\leqslant \tfrac{1}{4}\left|\left|\big(\rho^c_{00} - \rho^c_{11}\big)\right|\right|_{\mathrm{Tr}} + \left|\left|\big(\rho^c_{10} - \rho^c_{01}\big)\right|\right|_{\mathrm{Tr}}\\ &\leqslant \tfrac{1}{2}\big[D_{\mathrm{Tr}}\big(\rho^c_{00}, \rho^c_{11}\big) + D_{\mathrm{Tr}}\big(\rho^c_{01}, \rho^c_{10}\big)\big]\\ &\Rightarrow P^{\mathrm{opt}}_{\mathrm{disc}}(\rho_{(a=0)}, \rho_{(a=1)})\\ &\leqslant \tfrac{1}{2}\big(P^{\mathrm{opt}}_{\mathrm{disc}}\big(\rho^c_{00}, \rho^c_{11}\big) + P^{\mathrm{opt}}_{\mathrm{disc}}\big(\rho^c_{01}, \rho^c_{10}\big)\big)\\ &= P^{\mathrm{opt}}_{\mathrm{disc}}\big(\rho^c_{00}, \rho^c_{11}\big). \end{aligned} \tag{D29}$$

The last equality follows since for both ensembles, $\{|\phi_{0,0}\rangle, |\phi_{1,1}\rangle\}$ and $\{|\phi_{0,1}\rangle, |\phi_{1,0}\rangle\}$, we have that their output clones have equal discrimination probability:

$$P^{\mathrm{opt}}_{\mathrm{disc}}\big(\rho^c_{00}, \rho^c_{11}\big) = P^{\mathrm{opt}}_{\mathrm{disc}}\big(\rho^c_{01}, \rho^c_{10}\big). \tag{D30}$$

This is because the QCM is symmetric and depends only on the overlap of the states (we have in both cases $\langle\phi_{00}|\phi_{11}\rangle = \langle\phi_{01}|\phi_{10}\rangle = \sin(2\phi)$).

Furthermore, since the cloning machine can only lower the discrimination probability between two states, we have

$$P^{\mathrm{opt}}_{\mathrm{disc}}\big(\rho^c_{00}, \rho^c_{11}\big) \leqslant P^{\mathrm{opt}}_{\mathrm{disc}}\big(\rho^c_{00}, |\phi_{1,1}\rangle\langle\phi_{1,1}|\big) =: \overline{P^{\mathrm{opt}}_{\mathrm{disc}}}.$$

Now using the relationship between fidelity and the trace distance, we have the bounds

$$\begin{aligned} \tfrac{1}{2} + \tfrac{1}{2}\big(1 - \sqrt{\langle\phi_{1,1}|\rho^c_{00}|\phi_{1,1}\rangle}\big) &\leqslant \overline{P^{\mathrm{opt}}_{\mathrm{disc}}}\\ &\leqslant \tfrac{1}{2} + \tfrac{1}{2}\sqrt{1 - \langle\phi_{1,1}|\rho^c_{00}|\phi_{1,1}\rangle}. \end{aligned} \tag{D31}$$

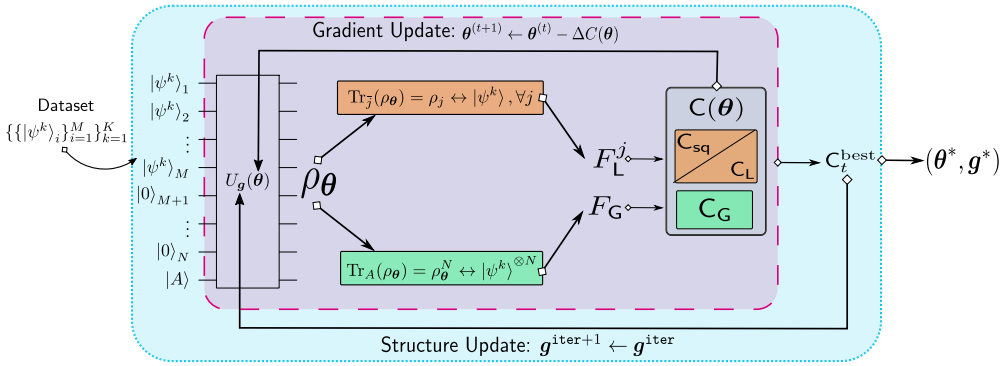

FIG. 7. Illustration of VarQlone for $M \to N$ cloning. A data set of K states is chosen from S, with M copies of each. These are fed with $N - M$ blank states, and possibly another ancilla, $|\phi\rangle_A$ into the variable structure Ansatz, $U_g(\boldsymbol{\theta})$. Depending on the problem, either the global, or local fidelities of the output state, $\rho_{\boldsymbol{\theta}}$, is compared to the input states, $|\psi^k\rangle$, and the corresponding local or global cost function, $C(\boldsymbol{\theta})$ is computed, along with its gradient. We have two optimization loops, one over the continuous parameters, $\boldsymbol{\theta}$, by gradient descent, and the second over the circuit structure, $\boldsymbol{g}$. Gradient descent over $\boldsymbol{\theta}$ in each structure update step outputs, upon convergence, the "minimum" cost function value, C_t^{best}, for the chosen cost function, $t \in \{\mathsf{L}, \mathsf{sq}, \mathsf{G}\}$.

By plugging in the observed density matrix for the output clone, we can find this discrimination probability. As in the previous section, the output density matrix from the QCM for an output clone can be written as Eq. (D11):

$$\rho_{00}^c = \alpha|\phi_{0,0}\rangle\langle\phi_{0,0}| + \beta|\phi_{1,1}\rangle\langle\phi_{1,1}| + \gamma(|\phi_{0,0}\rangle\langle\phi_{1,1}| + |\phi_{1,1}\rangle\langle\phi_{0,0}|), \tag{D32}$$

which has a local fidelity, $F_{\mathsf{L}} = \langle\phi_{0,0}|\rho_{00}^c|\phi_{0,0}\rangle = \alpha + s^2\beta + s\gamma$. On the other hand, we have $F(\rho_{00}^c, |\phi_{1,1}\rangle\langle\phi_{1,1}|) = \langle\phi_{1,1}|\rho_{00}^c|\phi_{1,1}\rangle = s^2\alpha + \beta + s\gamma$.

Combining these two, we then have

$$F\left(\rho_{00}^c, |\phi_{1,1}\rangle\langle\phi_{1,1}|\right) = F_{\mathsf{L}} + (s^2 - 1)(\alpha - \beta). \tag{D33}$$

Plugging in F_{L} from Eq. (A4), and $\alpha - \beta = \sqrt{\frac{1-s^2}{1-s^4}}$ (for an optimal state-dependent cloner), we get

$$\frac{1}{2} + \frac{1}{2}\left[1 - \sqrt{F_{\mathsf{L}} + (s^2 - 1)\sqrt{\frac{1-s^2}{1-s^4}}}\right] \leqslant P_{\mathrm{disc},\mathcal{P}_2}^{\mathrm{opt,II}} \leqslant \frac{1}{2} + \frac{1}{2}\sqrt{1 - F_{\mathsf{L}} - (s^2 - 1)\sqrt{\frac{1-s^2}{1-s^4}}}. \tag{D34}$$

To complete the proof, we use $F_{\mathsf{L}} \approx 0.989$ and $s = 1/\sqrt{2}$ which gives the numerical discrimination probabilities above. ■

APPENDIX E: ALGORITHM SPECIFICS AND SUPPLEMENTAL NUMERICAL RESULTS

In the main text and in the preceding sections, we discussed the VarQlone algorithm and a high-level overview of the numerical results. Here we revisit these numerics and dive into some specifics of the algorithm, in particular the Ansatz we choose. A cartoon illustration of the main ingredients can be seen in Fig. 7, which includes the cost functions discussed in Appendix B.

We will examine three different choices for the Ansätze in the VarQlone circuit. The first two options are *fixed structure* meaning the only trainable parameters are the continuous rotation angles in a fixed gate sequence. We then generalize to the primary Ansatz, which is that of a *variable structure* where both the continuous parameters *and* the gates in the Ansatz are optimized over.

1. Fixed structure Ansätze

To demonstrate the following two Ansätze, we use $1 \to 2$ phase-covariant cloning, which as a reminder requires cloning the following states:

$$|\psi_{xy}(\eta)\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle + e^{i\eta}|1\rangle\right). \tag{E1}$$

a. Phase-covariant cloning with a fixed ideal Ansatz

As discussed in the main text, the ideal circuit for performing phase-covariant cloning is given by Fig. 2(b). Here we learn the parameters of this fixed circuit. This gives us the opportunity to illustrate the effect of measurement noise in using the SWAP test to compute the fidelity. The results of this can be seen in Fig. 8. We compare the SWAP test in Fig. 8(a) to direct simulation of the qubit density matrices (using quantum state tomography [133] with the forest-benchmarking library [134]) to compute the fidelities. in Fig. 8(b) to compute the fidelity. The effect of measurement noise can be clearly seen in the latter case.

We note in the main text that we do not use the SWAP test when running the experiments on the Aspen QPU. This is because the test fails to output the fidelity since both states to compare will be mixed due to device noise. However, this essentially reproduces the findings of Ref. [81] in a slightly different scenario. Furthermore, this was only possible because we had prior knowledge of an optimal circuit to implement the cloning transformation from Ref. [17,75]. Of course, in generality this information is not available, and so we favor the variable structure Ansatz discussed above.

b. Phase-covariant cloning with a fixed hardware-efficient Ansatz

We also test a hardware-efficient fixed structure Ansatz for the sample problem as in the previous section. Here we introduce a number of layers in the Ansatz, K, in which each

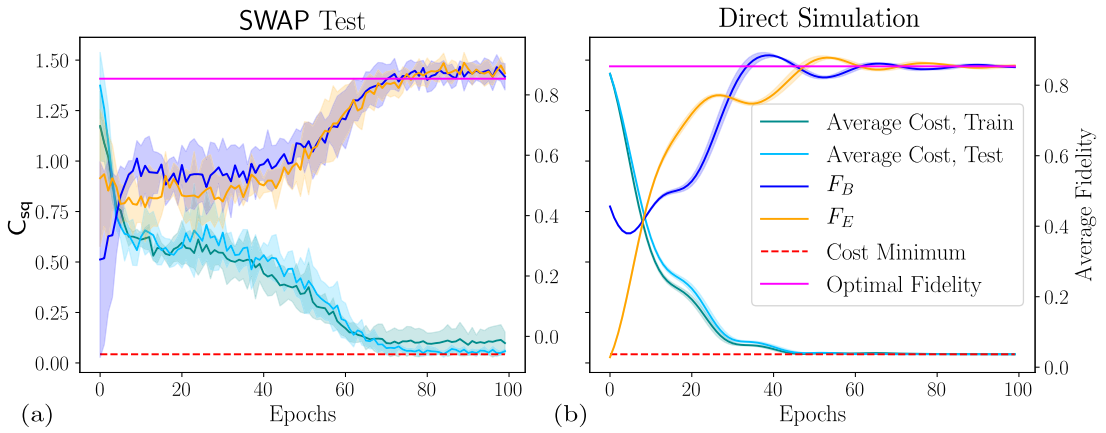

FIG. 8. Learning the parameters of the fixed circuit in Fig. 2(b). We use 30 random samples with an 80%–20% train-test split. To train, we use the analytic gradient, Eq. (B8), and the Adam optimizer with a batch size of 10 states and an initial learning rate of 0.05. In all cases, the error bars show mean and standard deviation over five independent training runs. Shown are the results when the fidelity is computed using (a) the SWAP test (with 50 measurement shots) and (b) using direct density matrix simulation. In both cases, we plot the average (squared) cost [Eq. (B1)] on the train and test set, and also the average fidelities of the output states of Bob, F_B, and Eve, F_E, corresponding to this cost function value. Also plotted are the theoretical optimal fidelities (magenta solid line) for this family of states, and the corresponding cost minimum (red dash line).

layer has a fixed structure. For simplicity, we choose each layer to have parameterized single-qubit rotations, $\mathrm{R}_y(\theta)$, and nearest-neighbor CZ gates. We deal again with $1 \to 2$ cloning, so we use three qubits and therefore we have two CZ gates per layer. We show the results for $K = 1$ layer to $K = 6$ layers in Fig. 9. Not surprisingly, we observe convergence to the minimum as the number of layers increases, saturating at $K = 3$.

Barren plateaus: Furthermore, we can examine VarQlone for the existence of barren plateaus in this scenario. We do this specifically for a local cost, given by

$$C_{\mathsf{L}} = \kappa \mathrm{Tr}[\mathsf{O}_{\mathsf{L}} U(\boldsymbol{\theta}) \rho U(\boldsymbol{\theta})^\dagger], \tag{E2}$$

$$\mathsf{O}_{\mathsf{L}} = c_0 \mathbb{1} + \sum_j c_j \mathsf{O}_j. \tag{E3}$$

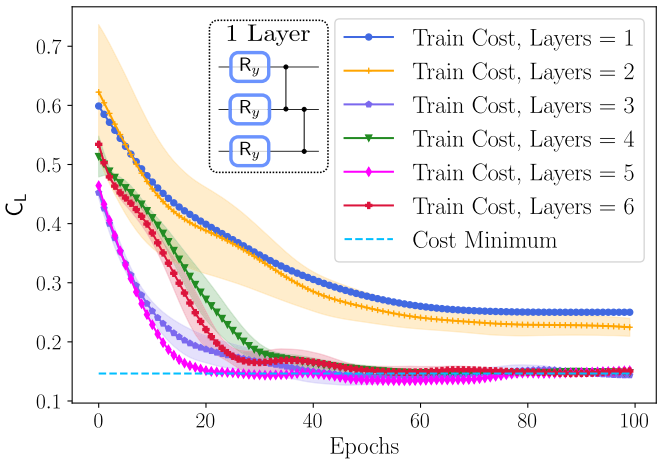

FIG. 9. Local cost, C_{L} minimized on a training set of 24 random phase-covariant states. We plot layers $L \in [1, \ldots, 6]$ of the hardware-efficient Ansatz shown in the inset. Fastest convergence is observed for $L = 5$ but $L = 3$ is sufficient to achieve a minimal cost value, which is the same number of entangling gates as in Fig. 2(b). Error bars shown mean and standard deviation over five independent training runs.

Note that taking $c_0 = 1, c_j = -1/N \, \forall j$ and $\kappa = 1$ recovers the specific form of our cost, Eq. (B2). We will prove that this cost does not exhibit barren plateaus for a sufficiently shallow alternating layered Ansatz, i.e., $U(\boldsymbol{\theta})$ contains blocks, W, acting on alternating pairs of qubits [70]. To do so, we first recall the following theorem from Ref. [70]:

Theorem 21 (Adapted from Theorem 2 in Ref. [70]). Consider a trainable parameter, θ^l in a block, W of an alternating layered Ansatz (denoted $U(\boldsymbol{\theta})$). Let $\mathrm{Var}[\partial_l C]$ be the variance of an m-local cost function, C with respect to θ^l. If each block in $U(\boldsymbol{\theta})$ forms a local 2-design, then $\mathrm{Var}[\partial_l C]$ is lower bounded by

$$G_N(K,k) \leqslant \mathrm{Var}[\partial_l C], \tag{E4}$$

$$\begin{aligned} G_N(K,k) = {} & \frac{2^{m(k+1)-1}}{(2^{2m}-1)^2(2^m+1)^{K+k}} \sum_{j \in j_{\mathcal{L}}} \sum_{\substack{(p,p') \in p_{\mathcal{L}_B} \\ p' \geqslant p}} \\ & \times c_j^2 D_{\mathrm{HS}}[\rho_{p,p'}, \mathrm{Tr}(\rho_{p,p'}) \mathbb{1}/d_{\rho_{(p,p')}}] D_{\mathrm{HS}} \\ & \times [\mathsf{O}_j, \mathrm{Tr}(\mathsf{O}_j)\mathbb{1}/d_{\mathsf{O}_j}]. \end{aligned} \tag{E5}$$

$j_{\mathcal{L}}$ are the set of j indices in the forward light cone $\mathcal{L}_B$ of the block W and $\rho_{p,p'}$ is the partial trace of the input state, ρ, down to the subsystems $S_p, S_{p+1}, \ldots, S_{p'}$. d_M denotes the dimension of a matrix M.

S_p in the above represents the qubit subsystem in which W acts. First, the operators O_j are all single-qubit projectors ($m = 1$ local), $|\psi\rangle\langle\psi|$, so we have

$$\begin{aligned} & D_{\mathrm{HS}}\left(\mathsf{O}_j, \mathrm{Tr}(\mathsf{O}_j)\frac{\mathbb{1}}{d}\right) \\ & = D_{\mathrm{HS}}\left(|\psi\rangle\langle\psi|, \mathrm{Tr}[|\psi\rangle\langle\psi|]\frac{\mathbb{1}}{2}\right) \\ & = \sqrt{\mathrm{Tr}\left[\left(|\psi\rangle\langle\psi| - \frac{\mathbb{1}}{2}\right)\left(|\psi\rangle\langle\psi| - \frac{\mathbb{1}}{2}\right)^\dagger\right]}, \end{aligned} \tag{E6}$$

$$= \sqrt{\mathrm{Tr}\left[|\psi\rangle\langle\psi| - \frac{|\psi\rangle\langle\psi|}{2} - \frac{|\psi\rangle\langle\psi|}{2} + \frac{\mathbb{1}}{4}\right]}$$

$$= \sqrt{\mathrm{Tr}\left(\frac{\mathbb{1}}{4}\right)} = \frac{1}{\sqrt{2}}. \tag{E7}$$

So $G(K,k)$ simplifies to

$$G_N(K,k) = \frac{2^k}{3^{K+k+2}\sqrt{2}N^2} \sum_{j\in j_{\mathcal{L}}} \sum_{\substack{(p,p')\in p\mathcal{L}_B \\ p'\geqslant p}} D_{\mathrm{HS}}(\rho_{p,p'}, \mathbb{1}/d_{\rho_{(p,p')}}). \tag{E8}$$

If we now define S_P to be the subsystems from p to p', the reduced state of ρ in S_P will be one of either $|\psi\rangle\langle\psi|^{\otimes|P|}$, $|0\rangle\langle 0|^{\otimes|P|}$ or $|\psi\rangle\langle\psi|^{\otimes q}|0\rangle\langle 0|^{\otimes|P|-q}$ for some $q<|P|$ where we denote $|P|$ to be the number of qubits in the reduced subsystem S_P. Since these are all pure states, we can compute $D_{\mathrm{HS}}(\rho_{p,p'}, \mathbb{1}/d_{\rho_{(p,p')}}) = \sqrt{1-1/d_{\rho_{(p,p')}}}$. Lower bounding the sum over j by 1 and $\sqrt{1-1/d_{\rho_{(p,p')}}}$ by $1/\sqrt{2}$ ($d_{\rho_{(p,p')}}$ is at least 2) gives

$$\frac{2^k}{3^{K+k+2}2N^2} \leqslant G_N(K,k). \tag{E9}$$

Finally, by choosing $K \in O[\log(N)]$, we have that $k, K+k \in O[\log(N)]$ and so $G_n(K,k) \in \Omega[1/\mathrm{poly}(N)]$. Since we have that if $G(K,k)$ vanishes no faster than $\Omega[1/\mathrm{poly}(N)]$, then so does the variance of the gradient and so will not require exponential resources to estimate. As a result, we can formalize the following corollary:

Corollary 3. [Absence of Barren Plateau in Local Cost] Given the local VarQlone cost function, C_{L} (Eq. (B2)) in $M \to N$ cloning, and a hardware-efficient fixed structure Ansatz, $U(\boldsymbol{\theta})$, made up of alternating blocks, W, with a depth $O[\log(N)]$, where each block forms a local 2-design. Then the variance of the gradient of C_{L} with respect to a parameter, θ_l can be lower bounded as

$$G_N := \min[G_N(K,k)] \leqslant \mathrm{Var}[\partial_l C],$$
$$G_N(K,k) \in \Omega[1/\mathrm{poly}(N)]. \tag{E10}$$

One final thing to note is that the Ansatz we choose in Fig. 9, does not form an exact local 2-design, but the same Ansatz is used in Ref. [70]) and is sufficient to exhibit a cost function-dependent barren plateau.

2. Variable structure Ansätze

Variations of the variable structure Ansatz approach have been proposed in Refs. [88,95] which could be easily incorporated, and we leave such investigation to future work. The approach of Ref. [71] (which we adopt) fixes the length, l, of the circuit sequence to be used, and as mentioned in the main text contains parameterized single-qubit gates, and unparameterized entangling gates, which we chose to be CZ for simplicity. For example, with a three-qubit chip, we have an example gatepool:

$$\mathcal{G} = \{\mathrm{R}_z^0(\theta), \mathrm{R}_z^1(\theta), \mathrm{R}_z^2(\theta), \mathrm{R}_x^0(\theta), \mathrm{R}_x^1(\theta), \mathrm{R}_x^2(\theta), \times \mathrm{R}_y^0(\theta), \mathrm{R}_y^1(\theta), \mathrm{R}_y^2(\theta), \mathrm{CZ}_{0,1}, \mathrm{CZ}_{1,2}, \mathrm{CZ}_{0,2}\}. \tag{E11}$$

We use the CZ gate as the entangler for two reasons. The first is that CZ is a native entangling gate on the Rigetti hardware. The second is that it simplifies our problem slightly, since it is symmetric on the control and target qubit, we do not need to worry about the ordering of the qubits: $\mathrm{CZ}_{i,j} = \mathrm{CZ}_{j,i}$. The fixed angle $\mathrm{R}_x(\pm\pi/2)$ and continuous angle $\mathrm{R}_z(\theta)$ gates are also native on the Rigetti hardware and we add the R_y gate for completeness, which can be compiled into the above as follows, $\mathrm{R}_y(\theta) = \mathrm{R}_x(\pi/2)\mathrm{R}_z(\theta)\mathrm{R}_x(-\pi/2)$. The unitary to be learned is given by

$$U_{\boldsymbol{g}}(\boldsymbol{\theta}) = U_{g_1}(\theta_1)U_{g_2}(\theta_2)\cdots U_{g_l}(\theta_l), \tag{E12}$$

where each gate is from the above set $\mathcal{G}$. The sequence, $\boldsymbol{g} := [g_1, \ldots, g_l]$, in Eq. (11) in the main text and Eq. (E12) above, corresponds to the indices of the gates in an ordered version of $\mathcal{G}$. So using $\mathcal{G}$ in Eq. (E11) as an example, $\boldsymbol{g} = [0, 6, 3, 2, 10]$ would give the unitary:

$$U_{\boldsymbol{g}}(\boldsymbol{\theta}) = \mathrm{R}_z^0(\theta_1)\mathrm{R}_y^1(\theta_2)\mathrm{R}_x^0(\theta_3)\mathrm{R}_z^2(\theta_4)\mathrm{CZ}_{0,1} \tag{E13}$$

and $\boldsymbol{\theta} := [\theta_1, \theta_2, \theta_3, \theta_4, 0]$ The procedure of Refs. [71,88,95,114] is intentionally flexible, and the gateset above Eq. (E15) can be swapped with any native gateset to fit on a particular quantum hardware.

At the beginning of the procedure, the gate sequence is chosen randomly (a random sequence, $\boldsymbol{g}$), and also the parameters ($\boldsymbol{\theta}$) therein.[6]

The optimization procedure proceeds over a number of epochs and iterations. In each iteration, $\boldsymbol{g}$ is perturbed by altering d gates, $\boldsymbol{g}^{\mathtt{iter}} \to \boldsymbol{g}^{\mathtt{iter}+1}$. The probability of changing d gates is given by $1/2^d$, and the probability of doing nothing (i.e., $\boldsymbol{g}^{\mathtt{iter}} = \boldsymbol{g}^{\mathtt{iter}+1}$) is

$$\Pr(d=0) = 1 - \sum_{d=1}^{l} \frac{1}{2^d} = 2 - \frac{1-\frac{1}{2^l}}{1-\frac{1}{2}} - \frac{1}{2^l}. \tag{E14}$$

The epochs correspond to optimization of the parameters $\boldsymbol{\theta}$ using gradient descent with the Adam optimizer, as throughout the main text. We typically set the maximum number of epochs to be 100 and iterations to be 50 in all this work. After each iteration, the best cost, C_t^{best} for a chosen cost: either the local, Eq. (B2) ($t = \mathsf{L}$), the global, Eq. (B3) ($t = \mathsf{G}$), squared, Eq. (B1) ($t = \mathsf{sq}$) or some other choice, is updated, if that iteration has found a circuit with a lower cost. As in Ref. [71], we repeatedly compress the sequence by removing redundant gates (e.g., combining $U_{g_i}(\theta_i)$ and $U_{g_{i+1}}(\theta_{i+1})$ if $g_i = g_i + 1$), and adding random gates to keep the sequence length fixed at g_l.

Figure 10 illustrates some results from this protocol. We find that with an increasing sequence length, the procedure is more likely to find circuits which achieve the minimum cost, and is able to first do so with a circuit with between 25 and 30 gates from the above gateset in Eq. (E15). We also plot the results achieved in a particular run of the protocol in Fig. 10(b). As the circuit learns, it is able to subsequently lower C_t^{best}, until it eventually finds a circuit capable of achieving the optimal cost for the problem.

[6] If some information is known about the problem beforehand, this could be used to initialize the sequence to improve performance.

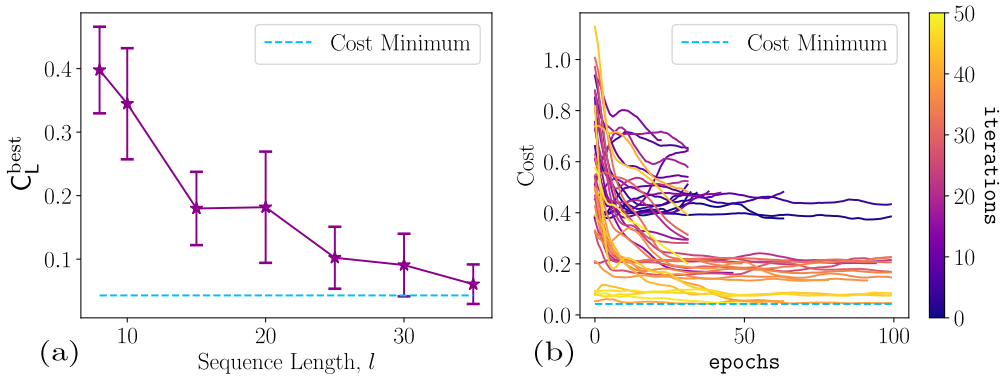

FIG. 10. (a) $C_{\mathsf{L}}^{\text{best}}$ as a function of sequence length, l in achieving the same task as Fig. 2(b), where the Bob's and Eve's clones appear in qubits 1 and 2. As l increases, the number of runs which successfully approach the theoretical minimum increases. Error bars shown mean and standard deviations for the minimum costs achieved over 20 independent runs with each sequence length. (b) Cost achieved for 50 iterations of the structure learning protocols, using a sequence length of $l = 35$. Each line corresponds to a slightly different circuit structure, $\boldsymbol{g}$. Early iterations (darker lines) are not able to find the minimum, but eventually, a circuit is found which has this capacity. For each $\boldsymbol{g}$, $\boldsymbol{\theta}$ is trained for 100 epochs of gradient descent, using the Adam optimizer. If an iteration has not converged close enough to $C_{\mathsf{L}}^{\text{best}}$ by 30 epochs, the iteration is ended.

a. Phase-covariant cloning

To begin, we will detail the procedure used to create Fig. 2. As a reminder, this is $1 \to 2$ cloning of phase-covariant states [Eq. (E1)] using the variable structure Ansätz described above. Here we allow three qubits (two output clones plus one ancilla) in the circuit. We also allow a fully connected (FC) gateset pool for this problem given by the following (indices represent qubits of the Aspen-8 sublattice):

$$\mathcal{G}_{\text{PC}} = \big\{\mathrm{R}_z^2(\theta), \mathrm{R}_z^3(\theta), \mathrm{R}_z^4(\theta), \mathrm{R}_x^2(\theta), \mathrm{R}_x^3(\theta), \mathrm{R}_x^4(\theta), \times \mathrm{R}_y^2(\theta), \mathrm{R}_y^3(\theta), \mathrm{R}_y^4(\theta), \mathrm{CZ}_{2,3}, \mathrm{CZ}_{3,4}, \mathrm{CZ}_{2,4}\big\}. \tag{E15}$$

Let us now discuss in greater detail the observations which can be drawn from Fig. 2. First, we notice that the ideal circuit in Fig. 2(b) suffers a degradation in performance when implemented on the QPU since it requires six entangling gates as it is attempting to transfer the information across the circuit. Furthermore, since the Aspen-8 chip does not have any three qubit loops in its topology, it is necessary for the compiler to insert SWAP gates.

Next we compare the ideal circuit to two examples learned by VarQlone. First, we force the qubit clones to appear in registers 2 and 3 [demonstrated in Fig. 2(c)] exactly as in Fig. 2(b). Second, we allow the clones to appear instead in registers 1 and 2 [demonstrated in Fig. 2(d), the circuit labeled "Rev." ("Reverse")]. The ability to make such a subtle change clearly demonstrates the advantage of our flexible approach. We notice that the restriction imposed in Fig. 2(c) results in only slightly improved performance over the ideal. However, by allowing the clones to appear in registers 1 and 2, VarQlone is able to find much more conservative circuits, having fewer entangling gates, and are directly implementable on a linear topology. This gives a significant improvement in the cloning fidelities, of about 15% when the circuit is run on the QPU, as observed in Fig. 2(a). For all results shown using a variable structure Ansatz, we use the forest-benchmarking library [134] to reconstruct the output density matrix in order to mitigate the effect of quantum noise.

Local vs global fidelities. As a final remark on this experiment, we can investigate the difference between the global and local fidelities achieved by the circuits VarQlone [i.e., in Fig. 2(c)] finds, versus the ideal one [shown in Fig. 2(b)]. Recall that in Appendix B 3 d, we showed that the "ideal" circuit achieves both the optimal local and global fidelities for this problem:

$$\text{Fig. 2(b)} \Rightarrow \begin{cases} F_{\mathsf{B}}^{(b)} = F_{\mathsf{E}}^{(b)} = F_{\mathsf{L}}^{\text{opt}} = \frac{1}{2}\big(1 + \frac{1}{\sqrt{2}}\big) \approx 0.853 \\ F_{\mathsf{G}}^{(b)} = F_{\mathsf{G}}^{\text{opt}} = \frac{1}{8}\big(1 + \sqrt{2}\big)^2 \approx 0.72 \end{cases}. \tag{E16}$$

In contrast, our learned circuit [Fig. 2(c)] maximizes the local fidelity, but in order to gain an advantage in circuit depth, compromises with respect to the global fidelity:

$$\text{Fig. 2(b)} \Rightarrow \begin{cases} F_{\mathsf{B}}^{(c)} \approx F_{\mathsf{E}}^{(c)} \approx F_{\mathsf{L}}^{\text{opt}} = 0.85 \\ F_{\mathsf{G}}^{(c)} \approx 0.638 < F_{\mathsf{G}}^{\text{opt}} \end{cases}. \tag{E17}$$

3. State-dependent cloning

Here we present the results of VarQlone when learning to clone the states used in the two coin-flipping protocols above. First, we focus on the states used in the original protocol, $\mathcal{P}_1$ for $1 \to 2$ cloning, and then move to the four-state protocol, $\mathcal{P}_2$. In the latter we also extend from $1 \to 2$ cloning to $1 \to 3$ and $2 \to 4$. These extensions will allow us to probe certain features of VarQlone, in particular explicit symmetry in the cost functions. In all cases, we use the variable structure Ansatz, and once a suitable candidate has been found, the solution is manually optimized further. The learned circuits used to produce the figures in this section are given in Appendix F.

a. Cloning $\mathcal{P}_1$ states

As a reminder, the two states used in this protocol are

$$|\phi_0\rangle := |\phi_{0,0}\rangle = \cos\left(\frac{\pi}{18}\right)|0\rangle + \sin\left(\frac{\pi}{18}\right)|1\rangle, \tag{E18}$$

$$|\phi_1\rangle := |\phi_{0,1}\rangle = \cos\left(\frac{\pi}{18}\right)|0\rangle - \sin\left(\frac{\pi}{18}\right)|1\rangle. \tag{E19}$$

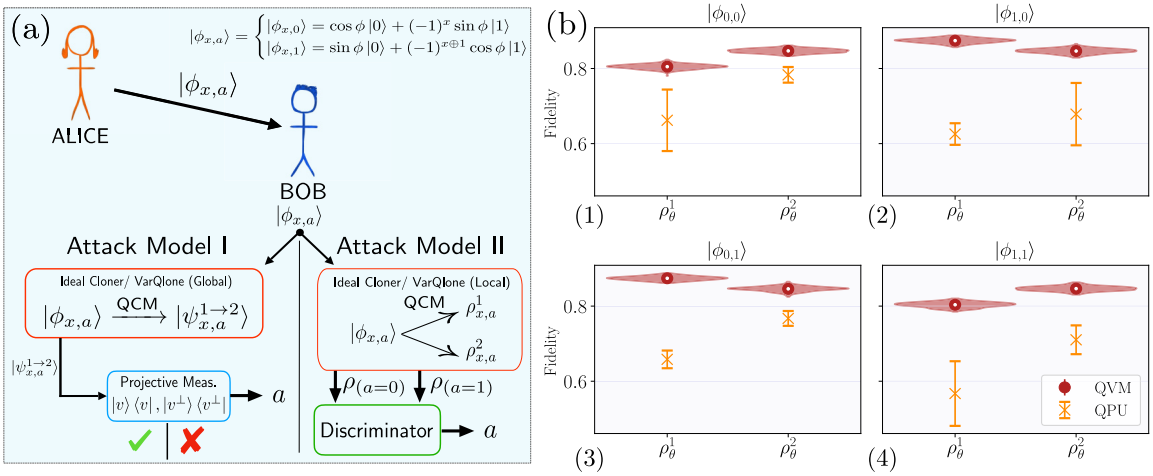

FIG. 11. Cloning attacks and numerical results for the protocol, $\mathcal{P}_2$.(a) The two cloning-based attacks we consider. In attack model I (left), Bob measures both output states with a set of fixed projective measurements, defined relative to the cloner output states, $|\psi^{1\to2}\rangle_{a,x}$ and guesses Alice's bit, a. In attack model II, Bob keeps one clone for either testing Alice later or to send back the deposit qubit requested by Alice. He uses then the other local clone to discriminate and guess a. (b) The fidelities achieved cloning the each state, $\{|\phi_{x,a}\rangle\}$ used in $\mathcal{P}_2$ with VarQlone. These numerics relate to scenario 1 from attack model II (see Appendix D 3 c). Each panel (1–4) shows both simulated (QVM, red circles) and on Rigetti hardware (QPU, orange crosses). We indicate the fidelities of the each clone received by Alice and Bob. For the QVM (QPU) results, 256 (3) samples of each state are used to generate statistics. Violin plots show complete distribution of outcomes and error bars show the means and standard deviations. Inset (i) shows the connectivity we allow in VarQlone for this example. The corresponding learned circuit is shown in Appendix F.

The fidelities achieved by the VarQlone learned circuit can be seen in Fig. 3 using the gate pool [Eq. (13)] which allows a linear entangling connectivity. A deviation from the optimal fidelity is observed in the simulated case, partly due to tomographic errors in reconstructing the cloned states. We note that the corresponding circuit for Fig. 3 actually used only two qubits (see Appendix F). This is because while VarQlone was *allowed* to use the ancilla, it *chose not* in this case by applying only identity gates to it. This mimics the behavior seen in the previous example of phase-covariant cloning. As such, we only use the two qubits shown in the inset (i) of the figure when running on the QPU to improve performance.

Now, returning to the attack on $\mathcal{P}_1$ above, we can compute the success probabilities using these fidelities. For illustration, let us return to the example in Eq. (D6), where instead the cloned state is now produced from our VarQlone circuit, $\rho_c^0 \to \rho^0_{\text{VarQlone}}$.

Theorem 22. [VarQlone Attack Bias on $\mathcal{P}_1$]

Bob can achieve a bias of $\epsilon \approx 0.29$ using a state-dependent VarQlone attack on the protocol, $\mathcal{P}_1$, with a single copy of Alice's state.

Theorem 22 can be proven by computing the success probability as in Appendix D 2:

$$P^{\text{VarQlone}}_{\text{succ},\mathcal{P}_1} = \tfrac{1}{2} + \tfrac{1}{4}\text{Tr}\left|\rho_1 - |\phi_1\rangle\langle\phi_1| \otimes \rho^0_{\text{VarQlone}}\right| \approx 0.804 \tag{E20}$$

The state $\rho_1 = |\phi_0\rangle\langle\phi_0| \otimes |\phi_1\rangle\langle\phi_1|$ as in Eq. (D6). Here we have a higher probability for Bob to correctly guess Alice's bit, a, but correspondingly the detection probability by Alice is higher than in the ideal case, due to a lower local fidelity of $F_{\text{L}}^{\text{VarQlone}} = 0.985$.

b. Cloning $\mathcal{P}_2$ states

Next, we turn to the family of states used in the four-state protocol, which are

$$|\phi_{x,a}\rangle = \begin{cases} \left|\frac{\pi}{8}_{x,0}\right\rangle = \cos\left(\frac{\pi}{8}\right)|0\rangle + (-1)^x \sin\left(\frac{\pi}{8}\right)|1\rangle \\ \left|\frac{\pi}{8}_{x,1}\right\rangle = \sin\left(\frac{\pi}{8}\right)|0\rangle + (-1)^{x\oplus1} \cos\left(\frac{\pi}{8}\right)|1\rangle \end{cases}. \tag{E21}$$

$1 \to 2$ *Cloning:* First, we repeat the exercise from above with the same scenario, using the same gateset and subset of the Aspen-8 lattice ($\mathcal{G}_{\mathcal{P}_2^{1\to2}} = \mathcal{G}_{\mathcal{P}_1^{1\to2}}$). We use the local cost, Eq. (B2), to train the model, with a sequence length of 35 gates. The results are seen in Fig. 11(b) both on the QVM and the QPU. We note that the solution exhibits some small degree of asymmetry in the output states, due to the form of the local cost function. This asymmetry is especially pronounced as we scale the problem size and try to produce N output clones, which we discuss in the next section.

Now, we can relate the performance of the VarQlone cloner to the attacks discussed in Appendix D 3. We do this by explicitly analyzing the output states produced in the circuits used to achieve fidelities shown in Fig. 11(b) and following the derivation in Appendix D for Theorem 23 and Theorem 24:

Theorem 23. [VarQlone Cloning Attack (I) Bias on $\mathcal{P}_2$]

Using a cloning attack on the protocol, $\mathcal{P}_2$, (in attack model I) Bob can achieve a bias:

$$\epsilon^{i}_{\mathcal{P}_2,\text{VarQlone}} \approx 0.345. \tag{E22}$$

Similarly, we have the bias which can be achieved with attack II:

Theorem 24. [VarQlone Cloning Attack (II) Bias on $\mathcal{P}_2$]

Using a cloning attack on the protocol, $\mathcal{P}_2$, (in attack model

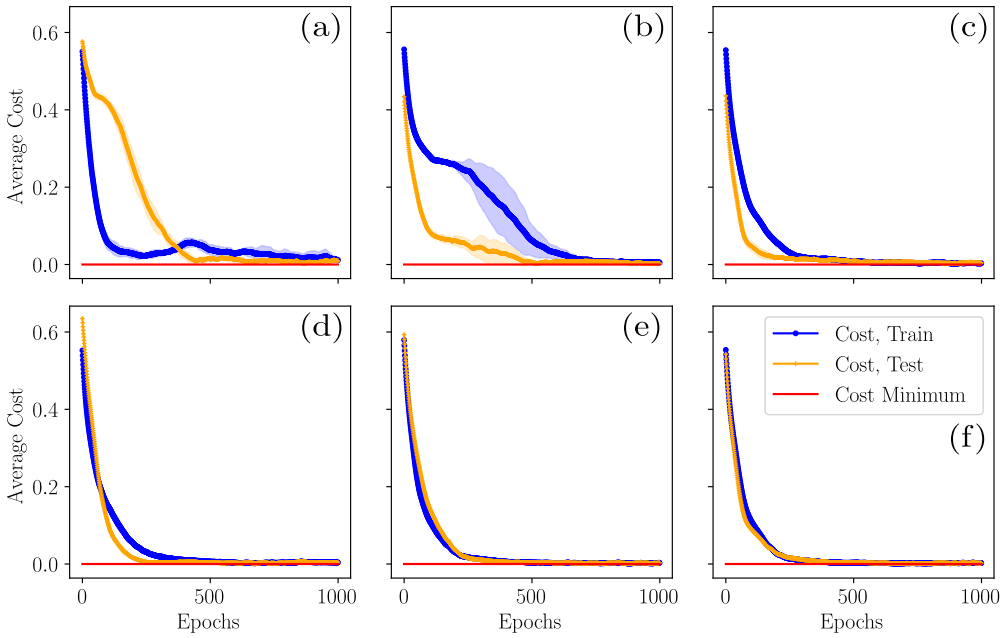

FIG. 12. Sample complexity of VarQlone using the squared cost. We begin with a random initialization of the structure learned circuit in Fig. 2(c) and optimize the parameters using different sizes in the training-text set, and different minibatch sizes. All of the following using a train-test split of 20%, and we denote the tuple (i, j, k) as i = number of training samples, j = number of test samples, k = batch size. (a) (1, 1, 1), (b) (4, 1, 2), (c) (8, 2, 5), (d) (16, 4, 8) (e) (40, 10, 15), (f) (80, 20, 20).

II) Bob can achieve a bias:

$$\epsilon^{\mathrm{II}}_{\mathcal{P}_2,\mathrm{VarQlone}} = 0.241. \tag{E23}$$

The discrepancy between these results and the ideal biases are primarily due to the small degree of asymmetry induced by the heuristics of VarQlone. However, we emphasize that these biases can now be achieved constructively.

$1 \to 3$ *and* $2 \to 4$ *Cloning:* Finally, we extend the above to the more general scenario of $M \to N$ cloning, taking $M = 1, 2$ and $N = 3, 4$. These examples are illustrative since they demonstrate strengths of the squared local cost function [Eq. (B1)] over the local cost function [Eq. (B2)]. In particular, we find the local cost function does not enforce symmetry strongly enough in the output clones, and using only the local cost function, suboptimal solutions are found. We particularly observed this in the example of $2 \to 4$ cloning, where VarQlone tended to take a shortcut by allowing one of the input states to fly through the circuit (resulting in nearly 100% fidelity for that clone), and then attempt to perform $1 \to 3$ cloning with the remaining input state. By strongly enforcing

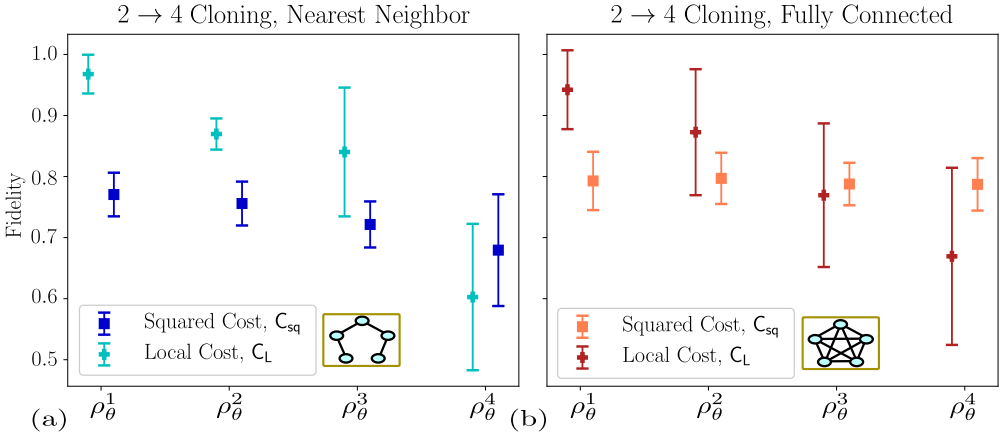

FIG. 13. Comparison between the local [Eq. (B2)] and squared [Eq. (B1)] cost functions for $2 \to 4$ cloning. (a) The nearest-neighbor (NN) and (b) a fully connected (FC) entanglement connectivity allowed in the variable structure Ansatz. Again, we use the family of states in the protocol $\mathcal{P}_2$. Plots show the mean and standard deviation of the optimal fidelities found by VarQlone over 10 independent runs (10 random initial circuit structures). A sequence length of 35 is used for $1 \to 3$ and 40 for $2 \to 4$, with 50 iterations of the variable structure Ansatz search in both cases. Here we use the same experiment hyperparameters as in Fig. 4.

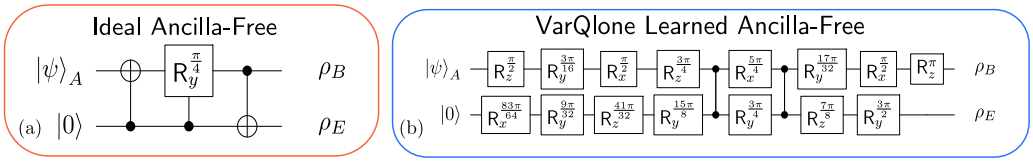

FIG. 14. Two qubit circuits to clone phase-covariant states, without ancilla. (a) Optimal circuit from Ref. [136], (b) circuit learned by VarQlone. In it can be checked that the ideal circuit in (a) can be compiled to *also* use two CZ plus single-qubit gates, so VarQlone has found something close to optimal. The average fidelities for B, E for the circuit in (b) is $F^{B,\mathrm{PC}}_{\mathsf{L},\mathrm{VarQlone}} \approx 0.854$ and $F^{E,\mathrm{PC}}_{\mathsf{L},\mathrm{VarQlone}} \approx 0.851$, respectively, over 256 input samples, $|\psi\rangle_A$ (comparing to the ideal fidelity of $F^{\mathrm{PC}}_{\mathsf{L},\mathrm{opt}} = 0.853$).

symmetry in the output clones using the squared cost, this can be avoided as we demonstrate explicitly in Appendix E 5.

We also test two connectivities in these examples, a fully connected (FC) and a nearest-neighbor (NN) architecture as allowed by the following gatesets:

$$\mathcal{G}^{\mathrm{NN}}_{\mathcal{P}_2^{1\to3}} = \left\{\mathrm{R}^i_z(\theta), \mathrm{R}^i_x(\theta), \mathrm{R}^i_y(\theta), \mathrm{CZ}_{2,3}, \mathrm{CZ}_{3,4}, \mathrm{CZ}_{4,5}\right\}, \quad \forall i \in \{2,3,4,5\}, \tag{E24}$$

$$\mathcal{G}^{\mathrm{FC}}_{\mathcal{P}_2^{1\to3}} = \left\{\mathrm{R}^i_z(\theta), \mathrm{R}^i_x(\theta), \mathrm{R}^i_y(\theta), \mathrm{CZ}_{2,3}, \mathrm{CZ}_{2,4}, \times \mathrm{CZ}_{2,5}, \mathrm{CZ}_{3,4}, \mathrm{CZ}_{3,5}, \mathrm{CZ}_{4,5}\right\}, \quad \forall i \in \{2,3,4,5\}. \tag{E25}$$

Note that for $1 \to 3$ ($2 \to 4$) cloning, we actually use four (five) qubits, with one being an ancilla. The results of these experiments are given in Fig. 4. We use the following hyperparameters for this experiment: (1) a sequence length of $l = 35$ for $1 \to 3$, and $l = 40$ for $1 \to 4$ with 50 iterations over $\boldsymbol{g}$ in both cases, and (2) the Adam optimizer with an initial learning rate of $\eta_{\mathrm{init}} = 0.05$, 3) 50 training samples. In all cases, we use the squared cost function, C_{sq}, to train and its gradients.

4. Training sample complexity

Here we study the sample complexity of the training procedure by retraining the continuous parameters of the learned circuit [Fig. 2(b)] starting from a random initialization of the parameters, $\boldsymbol{\theta}$ (illustrated in Fig. 12). As expected, as the number of training samples increases [i.e., the number of random choices of the phase parameter, η, in Eq. (3)], the generalization error (difference between training and test error) approaches zero. This is not surprising, since the training set will eventually cover all states on the equator of the Bloch sphere.

5. Local cost function comparison

In Fig. 13 we demonstrate the weakness of the local cost function, $\mathsf{C_L}$, in not enforcing symmetry strongly enough in the problem output, and how the squared cost function, $\mathsf{C_{sq}}$ can alleviate this, for $2 \to 4$ cloning specifically. Here we show the optimal fidelities found by VarQlone with a variable structure Ansatz, starting from a random structure. The local cost tends towards local minima, where one of the initial states (ρ^1_θ) ends up with high fidelity, while the last qubit (ρ^4_θ) has a low fidelity. This is alleviated with the squared cost function which is clearly more symmetric, on average, in the output fidelities. This is observed for both circuit connectivities we try (although a NN architecture is less able to transfer information across the circuit for a fixed depth).

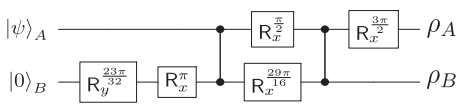

FIG. 15. Circuit learned by VarQlone in to clone states, $|\phi_0\rangle, |\phi_1\rangle$, with an overlap $s = \cos(\pi/9)$ in the protocol, $\mathcal{P}_1$. For example, ρ_A is the clone sent back to Alice, while ρ_B is kept by Bob.

APPENDIX F: VARQLONE LEARNED CIRCUITS

Here we give the explicit circuits learned by VarQlone and which give the results in the main text. We mention as above that these are only representative examples, and many alternatives were also found in each case.

1. Ancilla-free phase-covariant cloning

The circuits found in Fig. 2 to clone phase-covariant states are slightly more general than we may wish to use. In particular, the circuit Fig. 2(b) also has the ability to clone *universal* states, due to the addition of the ancilla, which can be used as a resource. However, it is known that phase-covariant cloning can be implemented economically, i.e., *without* the ancilla [16,135].[7] As such, we could compare against a shorter depth circuit which also does not use the ancilla. For example, the circuit from Ref. [136] shown in Fig. 14(a) is also able to achieve the optimal cloning fidelities (∼0.85). An example VarQlone learned circuit for this task can be seen in Fig. 14(b) which has two CZ gates. We note that this ideal circuit can be compiled to *also* use two CZ gates, so in this case VarQlone finds a circuit which is approximately comparable up to single-qubit rotations.

2. Mirror phase-covariant cloning

As mentioned in the main text, a variation of phase-covariant cloning exists called *mirror* phase-covariant cloning [77]. Here given states of the form

$$|\psi(\theta,\eta)\rangle = \cos(\theta)|0\rangle + e^{i\eta}\sin(\theta)|1\rangle, \tag{F1}$$

mirror cloning refers to the adversary having knowledge of $|\langle \mathsf{Z}\rangle|$, or equivalently known $\sin(\theta)$ (in contrast to known $\langle \mathsf{Z}\rangle$ (or θ) in the case of phase-covariant cloning.

[7] Although as discussed above in Appendix C, the economical version does not provide the optimal attack on related protocols.

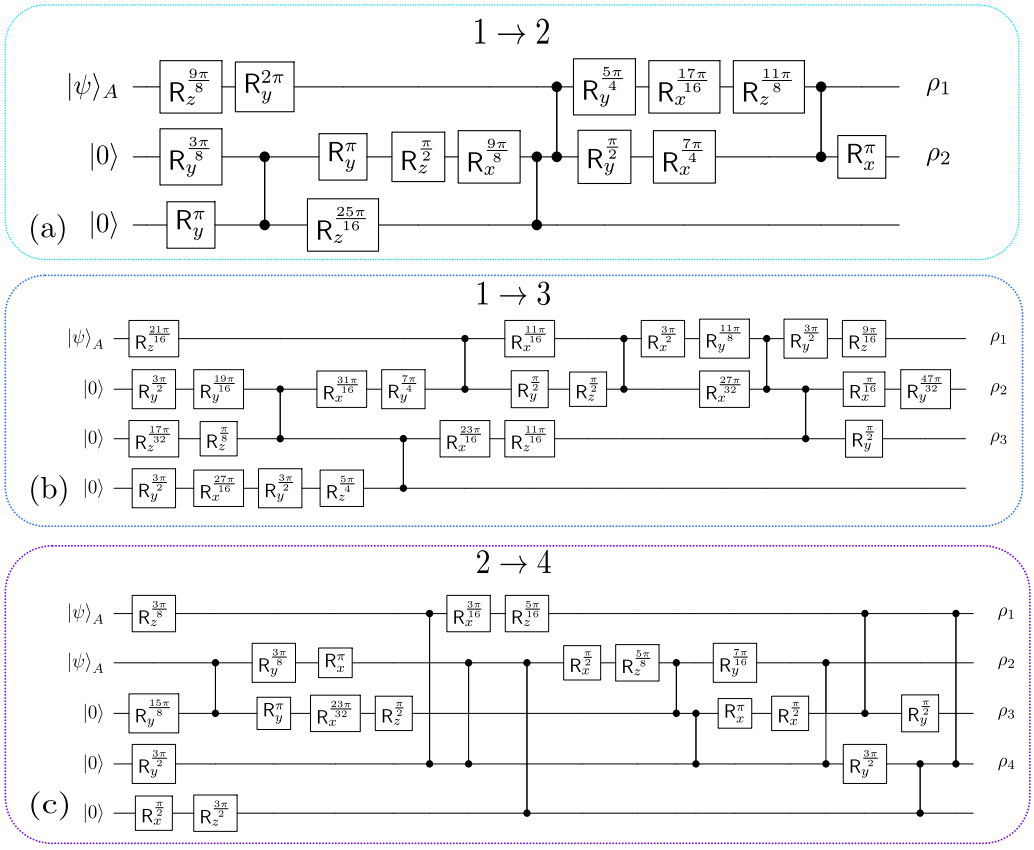

FIG. 16. Circuits learned by VarQlone to clone states from the protocol, $\mathcal{P}_2$ for (a) $1 \to 2$, (b) $1 \to 3$, and (c) $2 \to 4$ cloning. These specific circuits produce the fidelities in Fig. 11(b) for $1 \to 2$, (using the local cost function), and in Fig. 4 for $1 \to 3$ and $2 \to 4$ (using the squared cost function). We allow an ancilla for all circuits, and ρ_k indicates the qubit which will be the kth output clone.

Reference [77] provided optimal circuits for this family of states, which we show in Eq. (F2). However, while this circuit uses only four entangling gates, comparable to circuits (d) in Fig. 2 it still does not contain a fully nearest neighbor connectivity, and so hardware performance would suffer due to the need for SWAP operation:

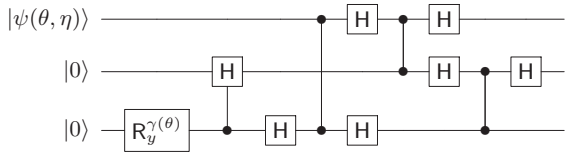

$$\gamma(\theta) := 2\cos^{-1}\left(\sqrt{\frac{1}{2} + \frac{\cos^2\theta}{2\sqrt{2 - 4\cos^2\theta + 3\cos^4\theta}}}\right) \tag{F2}$$

Interestingly, Ref. [77] also demonstrates a circuit to perform optimal mirror cloning, which incorporates a Hamiltonian evolution. This insight may be useful in future versions of VarQlone on quantum hardware platforms where the underlying Hamiltonians and their evolutions are accessible to the architecture search gate pool.

3. State-dependent cloning circuits

Figure 15 shows the circuit used to achieve the fidelities in the attack on $\mathcal{P}^1$ in the main text. In training, we still allowed an ancilla to aid the cloning, but the example in Fig. 15 did not make use of it (in other words, VarQlone only applied gates which resolved to the identity on the ancilla), so we remove it to improve hardware performance. This repeats the behavior seen for the circuits learned in phase-covariant cloning. We mention again, that some of the learned circuits did make use of the ancilla with similar performance.

Figure 16 shows the circuits learned by VarQlone and approximately clone all four states in Eq. (E21) in the protocol, $\mathcal{P}_2$, for $1 \to 2$, $1 \to 3$ and $2 \to 4$ cloning. These are the specific circuits used to produce the fidelities in Fig. 11(b) and Fig. 4.

[1] S. Wiesner, Conjugate coding, ACM SIGACT News **15**, 78 (1983).

[2] C. H. Bennett and G. Brassard, Quantum cryptography: Public key distribution and coin tossing, Theor. Comput. Sci. **560**, 7 (2014).

[3] A. Broadbent, J. Fitzsimons, and E. Kashefi, Universal blind quantum computation, in *Proceedings of the 2009 50th Annual IEEE Symposium on Foundations of Computer Science* (IEEE, New York, 2009), p. 517.

[4] D. Gottesman and I. Chuang, Quantum digital signatures, arXiv:quant-ph/0105032 (2001).

[5] S. Aaronson, Quantum copy-protection and quantum money, in *Proceedings of the 2009 24th Annual IEEE Conference on Computational Complexity* (IEEE, New York, 2009), p. 229.

[6] J. A. Vaccaro, J. Spring, and A. Chefles, Quantum protocols for anonymous voting and surveying, Phys. Rev. A **75**, 012333 (2007).

[7] J. S. Bell, On the Einstein Podolsky Rosen paradox, Phys. Phys. Fiz. **1**, 195 (1964).

[8] J. Yin, Y. Cao, Y.-H. Li, S.-K. Liao, L. Zhang, J.-G. Ren, W.-Q. Cai, W.-Y. Liu, B. Li, H. Dai *et al.*, Satellite-based entanglement distribution over 1200 kilometers, Science **356**, 1140 (2017).

[9] Q. Yao, M. Wang, Y. Chen, W. Dai, Y.-F. Li, W.-W. Tu, Q. Yang, and Y. Yu, Taking human out of learning applications: A survey on automated machine learning, arXiv:1810.13306 [cs, stat] (2019).

[10] J.-G. Ren, P. Xu, H.-L. Yong, L. Zhang, S.-K. Liao, J. Yin, W.-Y. Liu, W.-Q. Cai, M. Yang, L. Li *et al.*, Ground-to-satellite quantum teleportation, Nature (London) **549**, 70 (2017).

[11] S. Pirandola, U. L. Andersen, L. Banchi, M. Berta, D. Bunandar, R. Colbeck, D. Englund, T. Gehring, C. Lupo, C. Ottaviani, J. Pereira *et al.*, Advances in quantum cryptography, Adv. Opt. Photonics **12**, 1012 (2020).

[12] W. K. Wootters and W. H. Zurek, A single quantum cannot be cloned, Nature (London) **299**, 802 (1982).

[13] V. Bužek and M. Hillery, Quantum copying: Beyond the no-cloning theorem, Phys. Rev. A **54**, 1844 (1996).

[14] L.-M. Duan and G.-C. Guo, Two non-orthogonal states can be cloned by a unitary-reduction process, arXiv:quant-ph/9704020 (1997).

[15] L.-M. Duan and G.-C. Guo, Probabilistic Cloning and Identification of Linearly Independent Quantum States, Phys. Rev. Lett. **80**, 4999 (1998).

[16] V. Scarani, S. Iblisdir, N. Gisin, and A. Acín, Quantum cloning, Rev. Mod. Phys. **77**, 1225 (2005).

[17] H. Fan *et al.*, Quantum cloning machines and the applications, Phys. Rep. **544**, 241 (2014).

[18] Z.-X. Xiong *et al.*, General quantum key distribution in higher dimension, Phys. Rev. A **85**, 012334 (2012).

[19] J. Preskill, Quantum computing in the NISQ era and beyond, Quantum **2**, 79 (2018).

[20] T. A. Brun, Quantum error correction, arXiv:1910.03672 [quant-ph] (2019).

[21] S. J. Devitt, W. J. Munro, and K. Nemoto, Quantum error correction for beginners, Rep. Prog. Phys. **76**, 076001 (2013).

[22] A. Lamas-Linares, C. Simon, J. C. Howell, and D. Bouwmeester, Experimental quantum cloning of single photons, Science **296**, 712 (2002).

[23] J. Fiurášek, Optical implementations of the optimal phase-covariant quantum cloning machine, Phys. Rev. A **67**, 052314 (2003).

[24] H. Chen, X. Zhou, D. Suter, and J. Du, Experimental realization of $1 \rightarrow 2$ asymmetric phase-covariant quantum cloning, Phys. Rev. A **75**, 012317 (2007).

[25] K. Bartkiewicz, K. Lemr, A. Černoch, J. Soubusta, and A. Miranowicz, Experimental Eavesdropping Based on Optimal Quantum Cloning, Phys. Rev. Lett. **110**, 173601 (2013).

[26] K. Bartkiewicz *et al.*, Experimental quantum forgery of quantum optical money, npj Quantum Inf. **3**, 7 (2017).

[27] P. Wittek, *Quantum Machine Learning: What Quantum Computing Means to Data Mining* (Academic Press, 2014).

[28] J. Biamonte *et al.*, Quantum machine learning, Nature (London) **549**, 195 (2017).

[29] D. Kopczyk, Quantum machine learning for data scientists, arXiv:1804.10068 [quant-ph] (2018).

[30] M. Schuld and F. Petruccione, *Supervised Learning with Quantum Computers*, Quantum Science and Technology (Springer International Publishing, 2018).

[31] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, New J. Phys. **18**, 023023 (2016).

[32] J. Biamonte, Universal variational quantum computation, Phys. Rev. A **103**, L030401 (2021).

[33] S. Endo, Z. Cai, S. C. Benjamin, and X. Yuan, Hybrid quantum-classical algorithms and quantum error mitigation, J. Phys. Soc. Jpn. **90**, 032001 (2021).

[34] D. Wecker, M. B. Hastings, and M. Troyer, Progress towards practical quantum variational algorithms, Phys. Rev. A **92**, 042303 (2015).

[35] M. Cerezo *et al.*, Variational quantum algorithms, Nat. Rev. Phys. **3**, 625 (2021).

[36] P. W. Shor, Algorithms for quantum computation: Discrete logarithms and factoring, in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science* (IEEE, 1994), pp. 124–134.

[37] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, Nature (London) **574**, 505 (2019).

[38] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu *et al.*, Quantum computational advantage using photons, Science **370**, 1460 (2020).

[39] F. Centrone, N. Kumar, E. Diamanti, and I. Kerenidis, Experimental demonstration of quantum advantage for NP verification with limited information, Nat. Commun. **12**, 850 (2021).

[40] A. Peruzzo *et al.*, A variational eigenvalue solver on a photonic quantum processor, Nat. Commun. **5**, 4213 (2014).

[41] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, arXiv:1411.4028 [quant-ph] (2014).

[42] S. Sim, P. D. Johnson, and A. Aspuru-Guzik, Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms, Adv. Quantum Technol. **2**, 1900070 (2019).

[43] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, Parameterized quantum circuits as machine learning models, Quantum Sci. Technol. **4**, 043001 (2019).

[44] N. Killoran *et al.*, Continuous-variable quantum neural networks, Phys. Rev. Res. **1**, 033063 (2019).

[45] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, Quantum circuit learning, Phys. Rev. A **98**, 032309 (2018).

[46] E. Grant *et al.*, Hierarchical quantum classifiers, npj Quantum Inf. **4**, 65 (2018).

[47] L. G. Wright and P. L. McMahon, The capacity of quantum neural networks, in *Conference on Lasers and Electro-Optics*, OSA Technical Digest (Optica Publishing Group, 2020), paper JM4G.5.

[48] B. Coyle, D. Mills, V. Danos, and E. Kashefi, The Born supremacy: Quantum advantage and training of an Ising Born machine, npj Quantum Inf. **6**, 60 (2020).

[49] I. Cong, S. Choi, and M. D. Lukin, Quantum convolutional neural networks, Nat. Phys. **15**, 1273 (2019).

[50] M. E. S. Morales, T. Tlyachev, and J. Biamonte, Variational learning of Grover's quantum search algorithm, Phys. Rev. A **98**, 062333 (2018).

[51] S. Khatri *et al.*, Quantum-assisted quantum compiling, Quantum **3**, 140 (2019).

[52] T. Jones and S. C. Benjamin, Quantum compilation and circuit optimisation via energy dissipation, Quantum **6**, 628 (2022).

[53] K. Heya, Y. Suzuki, Y. Nakamura, and K. Fujii, Variational quantum gate optimization, arXiv:1810.12745 [quant-ph] (2018).

[54] C. Bravo-Prieto, Variational quantum linear solver: A hybrid algorithm for linear systems, arXiv:1909.05820 [quant-ph] (2019).

[55] X. Xu *et al.*, Variational algorithms for linear algebra, Sci. Bull. **66**, 2181 (2021).

[56] H.-Y. Huang, K. Bharti, and P. Rebentrost, Near-term quantum algorithms for linear systems of equations with regression loss functions, New J. Phys. **23**, 113021 (2021).

[57] M. Krenn, M. Malik, R. Fickler, R. Lapkiewicz, and A. Zeilinger, Automated Search for New Quantum Experiments, Phys. Rev. Lett. **116**, 090405 (2016).

[58] A. A. Melnikov *et al.*, Active learning machine learns to create new quantum experiments, Proc. Natl. Acad. Sci. USA **115**, 1221 (2018).

[59] L. O'Driscoll, R. Nichols, and P. A. Knott, A hybrid machine learning algorithm for designing quantum experiments, Quantum Mach. Intel. **1**, 5 (2019).

[60] R. Nichols, L. Mineh, J. Rubio, J. C. F. Matthews, and P. A. Knott, Designing quantum experiments with a genetic algorithm, Quantum Sci. Technol. **4**, 045012 (2019).

[61] J. Wallnöfer, A. A. Melnikov, W. Dür, and H. J. Briegel, Machine learning for long-distance quantum communication, PRX Quantum **1**, 010301 (2020).

[62] A. Arrasmith, L. Cincio, A. T. Sornborger, W. H. Zurek, and P. J. Coles, Variational consistent histories as a hybrid algorithm for quantum foundations, Nat. Commun. **10**, 1 (2019).

[63] G. Ateniese *et al.*, Hacking smart machines with smarter ones: How to extract meaningful data from machine learning classifiers, Intl. J. Security Netw. **10**, 137 (2015).

[64] H. Maghrebi, T. Portigliatti, and E. Prouff, Breaking Cryptographic implementations using deep learning techniques, in *Proc. of the Security, Privacy, and Applied Cryptography Engineering - 6th International Conference, SPACE 2016, Hyderabad, India*, edited by C. Carlet, M. A. Hasan, and V. Saraswat, Lecture Notes in Computer Science, Vol. 10076 (Springer, 2016), pp. 3–26.

[65] N. Papernot, P. McDaniel, A. Sinha, and M. Wellman, Towards the science of security and privacy in machine learning, arXiv:1611.03814 [cs] (2016).

[66] M. M. Alani, Applications of machine learning in cryptography: A survey, in *Proceedings of the 3rd International Conference on Cryptography, Security and Privacy, ICCSP '19* (Association for Computing Machinery, New York, 2019), pp. 23–27.

[67] D. Mayers, L. Salvail, and Y. Chiba-Kohno, Unconditionally secure quantum coin tossing, arXiv:quant-ph/9904078 (1999).

[68] D. Aharonov, A. Ta-Shma, U. V. Vazirani, and A. C. Yao, Quantum bit escrow, in *Proceedings of the Thirty-second Annual ACM Symposium on Theory of Computing, STOC '00* (Association for Computing Machinery, Portland, OR, 2000), pp. 705–714.

[69] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, Nat. Commun. **9**, 4812 (2018).

[70] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, Cost function dependent barren plateaus in shallow parametrized quantum circuits, Nat. Commun. **12**, 1791 (2021).

[71] L. Cincio, Y. Subaşı, A. T. Sornborger, and P. J. Coles, Learning the quantum algorithm for state overlap, New J. Phys. **20**, 113022 (2018).

[72] R. LaRose, Overview and comparison of gate level quantum software platforms, Quantum **3**, 130 (2019).

[73] R. Jozsa, Fidelity for mixed quantum states, J. Mod. Opt. **41**, 2315 (1994).

[74] D. Bruß, M. Cinchetti, G. Mauro D'Ariano, and C. Macchiavello, Phase-covariant quantum cloning, Phys. Rev. A **62**, 012302 (2000).

[75] V. Bužek, S. L. Braunstein, M. Hillery, and D. Bruß, Quantum copying: A network, Phys. Rev. A **56**, 3446 (1997).

[76] H. Fan, K. Matsumoto, X.-B. Wang, and M. Wadati, Quantum cloning machines for equatorial qubits, Phys. Rev. A **65**, 012304 (2001).

[77] K. Bartkiewicz, A. Miranowicz, and Ş K. Özdemir, Optimal mirror phase-covariant cloning, Phys. Rev. A **80**, 032306 (2009).

[78] D. Bruß *et al.*, Optimal universal and state-dependent quantum cloning, Phys. Rev. A **57**, 2368 (1998).

[79] A. S. Holevo, Statistical decision theory for quantum systems, J. Multivariate Anal. **3**, 337 (1973).

[80] C. W. Helstrom, Quantum detection and estimation theory, J. Stat. Phys. **1**, 231 (1969).

[81] J. Jašek *et al.*, Experimental hybrid quantum-classical reinforcement learning by boson sampling: How to train a quantum cloner, Opt. Express **27**, 32454 (2019).

[82] R. LaRose, A. Tikku, E. O'Neel-Judy, L. Cincio, and P. J. Coles, Variational quantum state diagonalization, npj Quantum Inf. **5**, 57 (2019).

[83] K. Sharma, S. Khatri, M. Cerezo, and P. J. Coles, Noise resilience of variational quantum compiling, New J. Phys. **22**, 043006 (2020).

[84] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, Evaluating analytic gradients on quantum hardware, Phys. Rev. A **99**, 032331 (2019).

[85] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, in *Proceedings of the 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015*, edited by Y. Bengio and Y. LeCun (2015).

[86] G. Fubini, Sulle metriche definite da una forma hermitiana: Nota, Atti Reale Istit. Veneto Sci. Lett. Arti **63**, 502 (1904).

[87] E. Study, Kürzeste Wege im komplexen Gebiet, Math. Ann. **60**, 321 (1905).

[88] L. Li, M. Fan, M. Coram, P. Riley, and S. Leichenauer, Quantum optimization with a novel Gibbs objective function and ansatz architecture search, Phys. Rev. Res. **2**, 023074 (2020).

[89] S.-X. Zhang, C.-Y. Hsieh, S. Zhang, and H. Yao, Differentiable quantum architecture search, arXiv:2010.08561 [quant-ph] (2020).

[90] H. Liu, K. Simonyan, and Y. Yang, DARTS: Differentiable architecture search, in *Proceedings of the 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6–9, 2019* (OpenReview.net, 2019), https://openreview.net/forum?id=S1eYHoC5FX.

[91] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, An adaptive variational algorithm for exact molecular simulations on a quantum computer, Nat. Commun. **10**, 3007 (2019).

[92] M. Ostaszewski, E. Grant, and M. Benedetti, Structure optimization for parameterized quantum circuits, Quantum **5**, 391 (2021).

[93] D. Chivilikhin *et al.*, MoG-VQE: Multiobjective genetic variational quantum eigensolver, arXiv:2007.04424 [cond-mat, physics:quant-ph] (2020).

[94] M. Pirhooshyaran and T. Terlaky, Quantum circuit design search, arXiv:2012.04046 [quant-ph] (2021).

[95] Y. Du, T. Huang, S. You, M.-H. Hsieh, and D. Tao, Quantum circuit architecture search: Error mitigation and trainability enhancement for variational quantum solvers, arXiv:2010.10217 [quant-ph] (2020).

[96] N. Gisin and S. Massar, Optimal Quantum Cloning Machines, Phys. Rev. Lett. **79**, 2153 (1997).

[97] D. Bruss, A. Ekert, and C. Macchiavello, Optimal Universal Quantum Cloning and State Estimation, Phys. Rev. Lett. **81**, 2598 (1998).

[98] D. Bruß and C. Macchiavello, Approximate quantum cloning, in *Lectures on Quantum Information* (John Wiley & Sons, New York, 2006), pp. 53–71.

[99] V. Bergholm *et al.*, PennyLane: Automatic differentiation of hybrid quantum-classical computations, arXiv:1811.04968 [physics, physics:quant-ph] (2020).

[100] M. Broughton *et al.*, TensorFlow Quantum: A software framework for quantum machine learning, arXiv:2003.02989 [cond-mat, physics:quant-ph] (2020).

[101] E. R. Anschuetz, J. P. Olson, A. Aspuru-Guzik, and Y. Cao, Variational quantum factoring, arXiv:1808.08927 [quant-ph] (2018).

[102] J. Carolan *et al.*, Variational quantum unsampling on a quantum photonic processor, Nat. Phys. **16**, 322 (2020).

[103] C. Bravo-Prieto, D. García-Martín, and J. I. Latorre, Quantum singular value decomposer, Phys. Rev. A **101**, 062310 (2020).

[104] G. Verdon, J. Marks, S. Nanda, S. Leichenauer, and J. Hidary, Quantum Hamiltonian-based models and the variational quantum thermalizer algorithm, arXiv:1910.02071 [quant-ph] (2019).

[105] E. Grant, L. Wossnig, M. Ostaszewski, and M. Benedetti, An initialization strategy for addressing barren plateaus in parametrized quantum circuits, Quantum **3**, 214 (2019).

[106] A. Arrasmith, M. Cerezo, P. Czarnik, L. Cincio, and P. J. Coles, Effect of barren plateaus on gradient-free optimization, Quantum **5**, 558 (2021).

[107] M. Cerezo, A. Poremba, L. Cincio, and P. J. Coles, Variational quantum fidelity estimation, Quantum **4**, 248 (2020).

[108] W. Vinci and A. Shabani, Optimally stopped variational quantum algorithms, Phys. Rev. A **97**, 042346 (2018).

[109] J. Stokes, J. Izaac, N. Killoran, and G. Carleo, Quantum natural gradient, Quantum **4**, 269 (2020).

[110] R. LaRose and B. Coyle, Robust data encodings for quantum classifiers, Phys. Rev. A **102**, 032420 (2020).

[111] J. I. Colless *et al.*, Computation of Molecular Spectra on a Quantum Processor with an Error-Resilient Algorithm, Phys. Rev. X **8**, 011021 (2018).

[112] J. Schmidhuber, Deep learning in neural networks: An overview, Neural Netw. **61**, 85 (2015).

[113] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016).

[114] L. Cincio, K. Rudinger, M. Sarovar, and P. J. Coles, Machine learning of noise-resilient quantum circuits, PRX Quantum **2**, 010324 (2021).

[115] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th ed. (Cambridge University Press, Cambridge, 2010).

[116] V. Bužek and M. Hillery, Universal Optimal Cloning of Arbitrary Quantum States: From Qubits to Quantum Registers, Phys. Rev. Lett. **81**, 5003 (1998).

[117] R. F. Werner, Optimal cloning of pure states, Phys. Rev. A **58**, 1827 (1998).

[118] M. Keyl and R. F. Werner, Optimal cloning of pure states, testing single clones, J. Math. Phys. **40**, 3283 (1999).

[119] N. Cerf, S. Iblisdir, and G. Van Assche, Cloning and cryptography with quantum continuous variables, Eur. Phys. J. D **18**, 211 (2002).

[120] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf, Quantum Fingerprinting, Phys. Rev. Lett. **87**, 167902 (2001).

[121] W. Hoeffding, Probability inequalities for sums of bounded random variables, J. Am. Stat. Assoc. **58**, 13 (1963).

[122] J. Watrous, Limits on the power of quantum statistical zero-knowledge, in *Proccedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science* (IEEE Computer Society, Los Alamitos, CA, 2002), pp. 459–468.

[123] J.-L. Chen, L. Fu, A. A. Ungar, and X.-G. Zhao, Alternative fidelity measure between two states of an N-state quantum system, Phys. Rev. A **65**, 054304 (2002).

[124] P. E. M. F. Mendonça, R. d. J. Napolitano, M. A. Marchiolli, C. J. Foster, and Y.-C. Liang, Alternative fidelity measure between quantum states, Phys. Rev. A **78**, 052330 (2008).

[125] Z. Puchała and J. A. Miszczak, Bound on trace distance based on superfidelity, Phys. Rev. A **79**, 024302 (2009).

[126] P. W. Shor and J. Preskill, Simple Proof of Security of the BB84 Quantum Key Distribution Protocol, Phys. Rev. Lett. **85**, 441 (2000).

[127] A. Ferenczi and N. Lütkenhaus, Symmetries in quantum key distribution and the connection between optimal attacks and optimal cloning, Phys. Rev. A **85**, 052310 (2012).

[128] M. Blum, Coin flipping by telephone a protocol for solving impossible problems, ACM SIGCT News **15**, 23 (1983).

[129] H.-K. Lo and H. F. Chau, Why quantum bit commitment and ideal quantum coin tossing are impossible, Physica D **120**, 177 (1998).

[130] A. Kitaev. Quantum coin-flipping, *Quantum Information Processing* (MSRI, Simons Auditorium, 2002).

[131] G. Berlín, G. Brassard, F. Bussières, and N. Godbout, Fair loss-tolerant quantum coin flipping, Phys. Rev. A **80**, 062321 (2009).

[132] J. Bae and L.-C. Kwek, Quantum state discrimination and its applications, J. Phys. A: Math. Theor. **48**, 083001 (2015).

[133] M. G. D'Ariano, M. G. A. Paris, and M. F. Sacchi, Quantum tomography, in *Advances in Imaging and Electron Physics*, Advances in Imaging and Electron Physics, Vol. 128, edited by P. W. Hawkes (Elsevier, 2003), p. 205.

[134] K. Gulshen *et al.*, Forest benchmarking: QCVV using PyQuil, https://doi.org/10.5281/zenodo.3455847 (2019).

[135] C.-S. Niu and R. B. Griffiths, Two-qubit copying machine for economical quantum eavesdropping, Phys. Rev. A **60**, 2764 (1999).

[136] J. Du *et al.*, Experimental Quantum Cloning with Prior Partial Information, Phys. Rev. Lett. **94**, 040505 (2005).