# The Efficient Preparation of Normal Distributions in Quantum Registers 

Arthur G. Rattew, Yue Sun, Pierre Minssen, and Marco Pistoia<br>Future Lab for Applied Research and Engineering, JPMorgan Chase \& Co.

The efficient preparation of input distributions is an important problem in obtaining quantum advantage in a wide range of domains. We propose a novel quantum algorithm for the efficient preparation of arbitrary normal distributions in quantum registers. To the best of our knowledge, our work is the first to leverage the power of Mid-Circuit Measurement and Reuse (MCMR), in a way that is broadly applicable to a range of state-preparation problems. Specifically, our algorithm employs a repeat-until-success scheme, and only requires a constant-bounded number of repetitions in expectation. In the experiments presented, the use of MCMR enables up to a $862.6 \times$ reduction in required qubits. Furthermore, the algorithm is provably resistant to both phase-flip and bit-flip errors, leading to a first-of-its-kind empirical demonstration on real quantum hardware, the MCMR-enabled Honeywell System Models H0 and H1-2.

## 1 Introduction

The efficient preparation of input distributions is particularly important for a wide range of quantum algorithms, such as those for amplitude estimation [1], option pricing [2], principal-component analysis [3], matrix inversion [4], and machine learning [5-7], which all offer the potential for quantum advantage, notably in financial applications [8], so long as their initial distributions may be generated without introducing computational bottlenecks.

Constructing an arbitrary quantum state necessitates exponential-depth circuits [9]. As a result, any efficient state preparation technique must either be approximate in nature, or exploit information specific to the distribution being generated. Moreover, the quantum computers of the foreseeable future belong to the class of Noisy Intermediate-Scale Quantum (NISQ) hardware, characterized by small qubit counts, limited two-qubit gate fidelity, and short coherence times [10]. As a result, for an algorithm to offer quantum advantage in the near future, it is even more important that state-generation procedures use as shallow circuits with as few ancillary qubits as possible, and produce high-fidelity states even in the presence of low gate-execution fidelity.

Initially proposed by Lloyd and Weedbrook in 2018 for the purpose of generating quantum states, quantum generative adversarial networks (QGANs) employ two agents: a generator and a discriminator [3]. The generator is tasked with producing the desired distribution, which is then evaluated by the discriminator [11]. Numerous papers have since built upon this work, demonstrating QGANs on real quantum processors, and extending QGANs for both classical sampling as well as for loading coherent quantum states [12, 13].

However, these techniques are approximate, and are not necessary in cases where efficient circuits may be theoretically derived.

In the creation of exact distributions, or those motivated analytically, a wide range of approaches have been explored. Grover and Rudolph published a procedure for generating efficiently integrable (e.g., log-concave) probability-density functions [14]. While there are some conflicting opinions regarding the theoretical asymptotic efficiency of the described procedure $[1,15]$ ), all agree that this work does not offer an approach that is efficiently realizable in practice on extant NISQ hardware. Kitaev and Webb built upon the aforementioned work by Grover and Rudolph, describing a method of generating multivariate normal distributions through resampling [16], similar to the qubit-scaling procedure we present in this paper, but fundamentally different in a way that causes their approach to rapidly accumulate error. Moreover, their resampling technique only allows for the efficient creation of Gaussian distributions with small variances, and requires the use of multiple ancillary qubits, limiting its practicality in the NISQ era. Our approach has no such limitations. Häner et al. presented a number of quantum circuits with polynomial depth, implementing commonly encountered functions through piece-wise polynomial approximations [17]. Nevertheless, this approach comes with a significant overhead in both required ancillary qubits, and resulting circuit depths [1]. Moreover, we do not expect this approach to efficiently implement a range of functions, such as the exponential function (and thus normal distributions), as the piece-wise polynomial approach would require an infeasible number of pieces to create accurate models on near-term hardware.

Non-unitary transformations may enable distributions to be obtained more efficiently [18, 19]. Two common ways to implement non-unitary transformations are (1) scaling to higher dimensional spaces through the use of ancillary qubits [17]), and (2) introducing non-linearities by performing partial measurements throughout the coherent execution of a circuit. The use of an asymptotic number of ancillary qubits is not desirable for the purpose of introducing non-unitary transformations, as qubit counts are limited in the NISQ era. As such, we believe it is preferable to use Mid-Circuit Measurement and Reuse (MCMR) [20]. The use MCMR has been demonstrated to generate more NISQ-friendly circuits with lower resource footprints. Specifically, Yalovetzky et al. leveraged MCMR and quantum conditional logic for Quantum Phase Estimation (QPE), to reduce the number of ancilla qubits and two-qubit gates, and lower the requirement on qubit connectivity [21]. On the hardware, this QPE variant produced results with higher fidelity than the orignal QPE. In the context of this work, the MCMR-free approach requires a number of ancillas that scales asymptotically in the variance of the target distribution being generated. As a result, in the experiments performed, the MCMR-based solution we propose uses up to $862.6 \times$ fewer total qubits and $8617 \times$ fewer ancillary qubits than its MCMR-free counterpart, thereby making it a more usable state-preparation procedure in the NISQ era.

Repeat-until-success (RUS) paradigms are often used by approaches that introduce non-linearities through the use of mid-circuit partial measurements [22]. In these methodologies, ancillary flag qubits are entangled with primary data registers, and the flags are then measured to determine if a particular operation was applied successfully. If the operation was not applied successfully, it may be possible to apply an operation conditioned on the classical measurement to then obtain the desired output. Alternatively, the circuit may simply be re-executed until the desired transformation is obtained. Guerreschi observes that RUS approaches have been used in applications such as implementing quantum neurons with non-linear activation functions [23, 24], and in the synthesis of arbitrary single-qubit rotations [25]. Moreover, RUS has been applied to a number of
quantum arithmetic problems [26]. Our work appears to be the first demonstration of RUS in a state-preparation algorithm. Moreover, we also demonstrate that our usage of RUS gives our algorithm an intrinsic robustness against hardware noise, further enhancing the procedure's viability in the NISQ era.

## 2 The Algorithm

Our algorithm prepares a quantized normal distribution in a quantum register, and receives as its input five parameters: $\hat{\mu}, \hat{\sigma}^{2}, x_{0}, l$ and $n$. Here, $\hat{\mu}$ and $\hat{\sigma}^{2}$ specify the normal distribution $\mathcal{N}\left(\hat{\mu}, \hat{\sigma}^{2}\right), x_{0}$ and $l$ specify the interval $\left[x_{0}, x_{0}+l\right]$ upon which the distribution is produced, and $2^{n}$ is the desired resolution. We assume that $\left[x_{0}, x_{0}+l\right]$ contains the mean of the normal distribution, as well as effectively all of the probability density. First, we introduce notation to make explicit the distinction between continuous and discrete state spaces. We allow $\{|x\rangle\}_{x}$ to index the set of continuous states, with $x \in\left[x_{0}, x_{0}+l\right]$, and we allow $\left\{\left|x_{j}\right\rangle\right\}_{j}$ to index the set of discrete states with $j \in \mathbb{Z}_{N}$. We assume we have an $n$ qubit system, with $N=2^{n}$ states. Then, we allow $\Delta x=\frac{l}{N}$ to be the interval between our discrete states. The continuous and discrete states are related with,

$$
\begin{equation*}
\left|x_{j}\right\rangle=\left|x_{0}+j \Delta x\right\rangle . \tag{1}
\end{equation*}
$$

We define the set of integer-valued computational (standard) basis states $\{|j\rangle\}_{j}$ with $j \in \mathbb{Z}_{N}$, which are isomorphic to the discretized set of states $\left\{\left|x_{j}\right\rangle\right\}_{j}$ through the mapping $|j\rangle=\left|\left(x_{j}-x_{0}\right) \Delta x^{-1}\right\rangle$. The algorithm aims to produce a quantum state $|\psi\rangle$ such that the amplitude of each standard basis vector $|j\rangle$ (up to a constant normalization factor) is given by,

$$
\langle j \mid \psi\rangle=\left\langle x_{j} \mid \psi\right\rangle=\frac{1}{\sqrt{2 \pi \hat{\sigma}^{2}}} \int_{x_{j}}^{x_{j+1}} \exp \left[-\frac{(x-\hat{\mu})^{2}}{2 \hat{\sigma}^{2}}\right] d x .
$$

Moreover, we will assume that $\mu$ and $\sigma^{2}$ are in units corresponding to the integer-valued basis states, while $\hat{\mu}$ and $\hat{\sigma}^{2}$ are in units corresponding to the continuous input domain. As such, we have that $\hat{\mu}=\frac{l}{N} \mu$ and $\hat{\sigma}^{2}=\frac{l^{2}}{N^{2}} \sigma^{2}$.

### 2.1 Generating Normal Distributions through Discrete Random Walks

Consider a random walk on the set of integers $\{j\}_{j}$ with $j \in \mathbb{Z}_{N}$ with $t$ steps, where in each step there is an equal probability of transitioning from $j$ to $j$ and $j+1^{1}$. The probability of such a walk terminating on any given integer clearly follows a binomial distribution, which tends to the normal distribution as $N \rightarrow \infty$ by the central limit theorem. As such, an operator implementing the quantum dynamics,

$$
\begin{equation*}
\mathcal{H}|j\rangle=|j\rangle+|j+1\rangle, \tag{2}
\end{equation*}
$$

coherently performs such a random walk, and thus produces a quantum state with amplitudes following a binomial distribution. Note that the final procedure presented doesn't produce a binomial distribution, but rather directly produces a discrete normal distribution. Therefore, we may simulate such a discrete random walk with $t$ steps by applying

[^0]$\mathcal{H} t$ times to $\left|\psi_{i}\right\rangle$, hence by constructing $\mathcal{H}^{t}\left|\psi_{i}\right\rangle$ for some starting basis state $\left|\psi_{i}\right\rangle$. For those familiar with Galton machines, it may be more clear to consider this random walk in the context of simulating a Galton machine, and we indeed consider this perspective in Supplementary Information (SI) A, where we also provide a more comprehensive overview of the intuition behind the algorithm.

### 2.1.1 Deriving $t-\hat{\sigma}^{2}$ and $t-\hat{\mu}$ Relationships

Applying $\mathcal{H}^{t}$ to some initial state produces a binomial distribution (on the amplitudes) equivalent to that produced by the aforementioned discrete random walk with $t$ steps. As such, the mean of the produced distribution is given by $\mu=\frac{t}{2}$, while the standard deviation is given by $\sigma=\frac{1}{2} \sqrt{t}$. Note that if we wish to obtain a given variance in the probability distribution rather than in the amplitudes, by property of a normal distribution, we may simply produce $2 \times$ the desired probability variance as the amplitude variance. Throughout the remainder of this document, unless we explicitly state otherwise (e.g. by writing a probability mass function explicitly), we assume that our distributions are in terms of the amplitudes of quantum states rather than their probabilities. Then, we may compute the number of applications of $\mathcal{H}$ required to obtain a normal distribution with variance corresponding to $\hat{\sigma}^{2}$ as follows,

$$
\begin{gather*}
\hat{\sigma}^{2}=\frac{l^{2}}{N^{2}} \sigma^{2}=\frac{l^{2}}{4 N^{2}} t \\
\Longrightarrow t=\frac{4 N^{2}}{l^{2}} \hat{\sigma}^{2} . \tag{3}
\end{gather*}
$$

### 2.1.2 Qubit Scaling Procedure

Equation 3 may appear to suggest that to obtain a distribution on $n$ qubits corresponding to a variance $\hat{\sigma}^{2}$, a number of applications of $\mathcal{H}$ scaling with $O\left(2^{2 n}\right)$ would be required. This would be problematic, as even if each application of $\mathcal{H}$ could be implemented in polynomial depth (as we later show, it can) the overall circuit would necessitate exponential depth to produce a normal distribution with arbitrary variance.

Fortunately, as we show in SI D, it is possible to avoid this exponential blowup. In essence, we do so by first loading a nearly exact distribution on some small constant number of qubits, $n_{1}$, by applying $\mathcal{H} t_{1}$ times to the $n_{1}$ qubit $|0\rangle$ state.

We then iteratively add a new qubit in the $|+\rangle$ state (a uniform superposition on one qubit) as the least significant qubit ${ }^{2}$, obtaining a distribution on $n_{i+1}$ qubits where adjacent states share the amplitude of the corresponding state in the $n_{i}$ qubit register. This effectively doubles the "resolution" of the $n_{i}$-qubit distribution without adding any additional information. We then obtain the correct $n_{i+1}$ qubit distribution by applying $\mathcal{H}$ a total of $t_{i+1}$ times $^{3}$, we call these iterations "correction" iterations. We repeat this procedure until the desired distribution on the final number of qubits, $n_{m}$, is obtained. Thus, the procedure requires the calculation of a list $\left[t_{1}, t_{2}, \ldots, t_{m}\right]$ specifying the number of iterations of $\mathcal{H}$ to apply at each qubit count. As shown in SI D , the final $n_{m}$ qubit

[^1]variance obtained with any such list is given by,
\[

$$
\begin{equation*}
\sigma^{2}=\frac{4^{n_{m}-1}-1}{12}+\frac{1}{4} \sum_{k=1}^{m} 4^{m-k} t_{k} . \tag{4}
\end{equation*}
$$

\]

Moreover, given a desired input variance on the real input grid $\hat{\sigma}^{2}$, it is straightforward to workout a schedule $\left[t_{1}, t_{2}, \ldots, t_{m}\right]$ producing the corresponding variance $\sigma^{2}$ in the computational basis. In general, an effective initial setting for $t_{i}$ is $t_{2}=t_{3}=\ldots=t_{m}=c$ for some small constant integer $c$ (for example, we often take $c=2$ ), and then setting $t_{1}$ to whatever value produces the final desired variance. Such an assignment produces the shortest circuits possible, as it maximizes the number of iterations performed on the smallest qubit count, where the fewest number of iterations of $\mathcal{H}$ has the greatest impact on the variance obtained. As such, we may compute $t_{1}$ with,

$$
\begin{equation*}
t_{1} \leftarrow \frac{\sigma^{2}-\frac{4^{n_{m}-1}-1}{12}-\frac{c}{4} \sum_{k=2}^{m} 4^{m-k}}{4^{n_{m}-1}} . \tag{5}
\end{equation*}
$$

However, when $t_{1}$ is not an integer some error is incurred as only an integer number of iterations can be applied, and so some minor tuning of $t_{i}$ for greater values of $i$ may be required to get the exact desired variance (noting that it is always possible and straightforward to do so). Finally, we shift the distribution (by adding a constant amount to all states) to obtain the correct mean. The shift may be implemented using a simple adder circuit, and the shift amount is calculated as the difference of the obtained mean,

$$
\begin{equation*}
\mu_{\text {obtained }}=\frac{1}{2}\left(2^{n_{m}-1}-1\right)+\frac{1}{2} \sum_{k=1}^{m} 2^{m-k} t_{k} \tag{6}
\end{equation*}
$$

(also derived in SI D) and the desired input mean (in integer units).
The correctness of this procedure, as well as the analysis of the variance as a function of the number of correction iterations, is provided in SI D. Furthermore, the error in this qubit-scaling approximation is entirely corrected, meaning that no additional error is introduced by this optimization.

### 2.1.3 Implementing $\mathcal{H}$

To ensure the efficiency of this procedure, it is essential that $\mathcal{H}$ be implementable with a polynomial-depth circuit. We have explored a number of approaches, which may be broken into two high-level categories: MCMR-based and MCMR-free. The MCMR-based approach requires only a single ancillary qubit to produce the desired distribution at the cost of requiring a constant-bounded number of expected circuit evaluations, as proven in SI F. The circuit for the MCMR-based approach is derived and presented in SI B. The MCMR-free approach requires a number of ancilla qubits scaling linearly in the number of applications of $\mathcal{H}$ required. This approach is also briefly discussed in SI B, and is discussed in more detail in SI C. The complexity of $\mathcal{H}$ is simply the complexity of the adder gate implementation used. A number of possible implementations for the adder gate are discussed in SI C.1, but in summary, this work uses a QFT adder which requires $\mathcal{O}(n)$ depth (if already in Fourier space), and no additional ancilla qubits. Alternatively, to avoid the use of the QFT the adder may be implemented with $\mathcal{O}(n \log n)$ depth if $\mathcal{O}(n)$ ancilla qubits are available.

### 2.1.4 Noise Resistance

Finally, the primary variant of the algorithm, the MCMR-based approach, has provable resistance to hardware errors by virtue of the formulation of $\mathcal{H}$. This property is discussed in detail in SI E, and follows from the probability of the ancilla qubit collapsing to the $|1\rangle$ state in each application of $\mathcal{H}$. In summary, the $|0\rangle$ and $|1\rangle$ ancilla states are used to produce a non-unitary transformation in the main register. When the $|0\rangle$ ancilla state is measured, it is known that the operation was applied successfully. Of essence to the efficiency of this procedure is that the probability of measuring the $|1\rangle$ state decays exponentially in the number of iterations applied. The amplitudes produced on the $|1\rangle$ state cancel near-optimally, and so if any hardware errors occur, with limited exceptions, they reduce the amount of destructive interference, thus increasing the probability of measuring the ancilla in the $|1\rangle$ state, and thereby increasing the probability of discarding the error-affected execution. Essentially, this follows from the fact that the cancellation property relies on the continuity of the distribution, and the occurrence of errors (in any noise model) almost always results in the introduction of discontinuities that cause the destructive interference on the $|1\rangle$ state not to occur. This discussion is treated more rigorously in the aforementioned entry in SI E.

### 2.2 The Procedure

We provide a pseudocode sketch of the procedure in Algorithm 1; here we assume the use of mid-circuit partial measurements. Here, $n_{1}$ is the initial constant number of qubits on which the scaling procedure begins, while $n_{m}$ is the final target number of qubits upon which the distribution is to be prepared. The iteration counts $t_{i}$ can be defined much more flexibility than prescribed by this algorithm, but for simplicity we provide this recommended setting as it usually produces quite shallow circuits. Additionally, $c$ is some small constant number; $c=2$ often works well. The subscript on a ket denotes the number of qubits in that quantum state. Moreover, $A^{\alpha}$ is an adder gate that adds $\alpha$ to all states (as described in SI C.3), while $A_{n+1}^{+1}$ is a conditonal adder gate conditioned on qubit $n+1$ (as described in SI C.2). As usual, $H$ represents a single-qubit Hadamard gate. Finally, as proven in SI F the algorithm's probability of success is constant bounded.

All together, we have shown how to produce arbitrary normal distributions in quantum registers, with asymptotic complexity logarithmic in the desired resolution of the distribution. Our procedure requires a polynomial circuit depth and a constant number of expected circuit executions, and can efficiently load a normal distribution with arbitrary variance onto a quantum register - all with provable resistance to hardware noise.

## 3 Experimental Demonstrations

In this section, we first present and analyze the results from running the MCMR version of the algorithm on the Honeywell System Model H0 [20] and Honeywell System Model H1-2, then we present the results comparing various variants of the algorithm in numerical simulations, and finally we present a discussion of the fidelity of the algorithm in theoretically ideal conditions.

### 3.1 Quantum Hardware Experiments

The results from the hardware experiment are presented in Figure 1, with the configuration described in the corresponding caption. We use an optimized version of the MCMR-based

```
Algorithm 1 Generating Normal Distributions with Mid-Circuit Partial Measurement.
- We use the notation \(|\psi\rangle_{k}\) to make explicit that \(|\psi\rangle_{k}\) is a \(k\)-qubit state.
- Here \(|+\rangle\) refers to a single qubit uniform superposition.
- Arrows are used to update or define quantum states.
    procedure GenerateNormalDistribution \(\left(\hat{\mu}, \hat{\sigma}^{2}, x_{0}, l, n_{1}, n_{m}\right)\)
    \(N=2^{n_{m}}\)
    \(\sigma^{2}=\frac{N^{2}}{l^{2}} \hat{\sigma}^{2}\)
    \(\mu=\frac{N}{l} \hat{\mu}\)
    \(|\psi\rangle_{n_{1}} \leftarrow|0\rangle_{n_{1}}\)
    \(t_{2}=t_{3}=\ldots=t_{m}=c\)
    \(t_{1}=\left(\sigma^{2}-\frac{4^{n_{m}-1}-1}{12}-\frac{c}{4} \sum_{k=2}^{m} 4^{m-k}\right) / 4^{n_{m}-1} \quad \triangleright\) I.e. Equation 5
    for \(i=1\) to \(m\) do
        for \(j=1\) to \(t_{i}\) do
            Create the new state \(|\phi\rangle_{n_{i+1}}\) by applying \(\left(I^{\otimes n_{i}} \otimes H\right) A_{n_{i+1}}^{+1}\) to \(|\psi\rangle_{n_{i}} \otimes|+\rangle\).
                Measure qubit \(n_{i+1}\) of state \(|\phi\rangle_{n_{i+1}}\) in the standard basis
                Let \(x \in\{0,1\}\) be the result of the measurement
                Let \(|\phi\rangle_{n_{i}}\) be the resulting partial state
                if \(x=1\) then
                    Procedure failed, discard result, repeat procedure until success.
                end if
                \(|\psi\rangle_{n_{i}} \leftarrow|\phi\rangle_{n_{i}}\)
            end for
            if \(i<m\) then
                \(|\psi\rangle_{n_{i+1}} \leftarrow|\phi\rangle_{n_{i}}|+\rangle\)
            end if
    end for
    \(\alpha \leftarrow \mu-\frac{1}{2}\left(2^{n_{m}-1}-1+\sum_{k=1}^{n_{m}} 2^{n_{m}-k} t_{k}\right)\)
    \(|\psi\rangle_{n_{m}} \leftarrow A^{\alpha}|\psi\rangle_{n_{m}} \quad \triangleright\) Shifts the produced mean to the target mean
    return \(|\psi\rangle_{n_{m}} \quad \triangleright\) The procedure has succeeded
    end procedure
```

qubit-scaling algorithm, as presented in SI B.3. A total of 2500 samples were taken in this experiment, among which 532 occurred with all of the ancillary measurements yielding the $|0\rangle$ state. This corresponds to a post-selection rate of $21.28 \%$. In contrast, theoretical calculations indicate that performing this experiment in noiseless simulation would result in a post-selection rate of $\approx 31.93 \%$. Therefore, there is approximately a $10 \%$ deviation from the theoretical prediction and the experimental observations. In the right panel of Figure 1, the y-axis shows the selection rate at the measurement of the $t^{\text {th }}$ application of $\mathcal{H}$. The selection rate at $t$ is defined as the probability of measuring a $|0\rangle$ at that specific iteration given that all prior measurements in the same experiment yielded $|0\rangle$. In general, the selection rates observed on the ancillary qubit match the theoretical predictions, especially for the first three applications of $\mathcal{H}$, with the relatively small deviations well explained by the intrinsic stochastic variation given the limited quantity of samples collected. After the third application of $\mathcal{H}$, the theoretical results and the experimental results start to differ more significantly, insofar as the experimental results reject more samples than error-free analysis predicts. With each application of $\mathcal{H}$, additional CX gates are executed, increasing the probability of an error occurring. As predicted in SI E as the probability of obtaining an error increases, the probability of discarding a circuit execution


Figure 1: Samples Drawn from a 5-Qubit Implementation of the Algorithm on the Honeywell System Model H0. In this experiment, the algorithm is configured with $n_{1}=2, n_{2}=3, n_{3}=4$, and $t_{1}=t_{2}=t_{3}=2$. The distribution is then loaded into a quantum register using the proposed algorithm's MCMR variant, and then samples are drawn from the output. Thus, a total of 5 qubits are used in the experiment. The resulting circuit has 75 U 1 gates, 60 CX gates, 22 U 2 gates, and 10 measurements. A total of 2500 samples are taken, of which only 532 are kept in the algorithm's post-selection process. The left panel presents the post-selected results from the hardware-execution (bar-plot), the ideal statevector distribution (blue line), and the corresponding Gaussian distribution (orange line). Both plots are only assigned values directly on the integer grid points. The panel on the right compares the experimental and theoretical probabilities of obtaining the $|0\rangle$ state on the ancilla at the $t^{t h}$ iteration, given that all preceding iterations yielded $|0\rangle$ ancillary measurements. The colors of the circles indicate different qubit scaling stages in the algorithm ( $n_{1}$ - blue, $n_{2}$ - orange, $n_{3}$ - green).
also increases, thus explaining the deviations observed. A similar argument also explains the $10 \%$ experimental deviation from the theoretical prediction for the total number of shots kept. Indeed, this experimental deviation from the theoretical prediction supports the claim of the algorithm being noise resistant.

The left panel of Figure 2 presents the results of the same experiment as Figure 1, with the only difference being that the first two iterations of the procedure on 2 qubits were jump-started (defined precisely shortly), and the circuit was executed on Honeywell System Model H1-2 instead of Model H0. We performed jump-starting to demonstrate the significant boost it can have on the percentage of circuits kept after post-selection, and we used Honeywell System Model H1-2 instead of Honeywell System Model H0 as it was the hardware available when running the experiment. The right panel of Figure 2 presents the results of the same experiment as the left panel, only with an additional quantum Fourier transform applied at the end (which is actually implemented by removing the final inverse quantum Fourier transform, and appropriately reordering the bits). The purpose of the experiment shown in the right panel is to highlight the fact that the algorithm's resistance to hardware errors results in it not only drawing samples according to the correct probability distribution, but also actually produces the correct quantum state (with the correct relative phases, etc).

We first discuss the left panel of Figure 2. A $k$-jump start removes the first $k$ applications of $\mathcal{H}$, and instead some other circuit (e.g. a universal circuit when $n_{i}$ is small) is used to exactly produce the state that would have otherwise been created. As the algorithm's probability of failure exponentially decreases after each application of $\mathcal{H}$, setting $k$ to a small value such as 2 can significantly increase the probability of success, as it removes the iterations of the procedure where most failures would otherwise have occurred. Moreover, as the algorithm assumes that $n_{1}$ is some small constant number of qubits, such a


Figure 2: Samples Drawn from a 5-Qubit Implementation of the Algorithm on the Honeywell System Model H1-2 Measured in Two Bases, With Jump-Starting. In both experiments, the algorithm is configured with $n_{1}=2, n_{2}=3, n_{3}=4$, and $t_{1}=t_{2}=t_{3}=2$, and uses the MCMR variant (and thus requires one additional ancillary qubit). Additionally, in both experiments the two iterations at $n_{1}$ are performed with a deterministic procedure producing the exact two-qubit distribution, a process we call jump-starting. In the left panel, samples are drawn in the computational basis, while in the right panel samples are drawn in the Fourier basis. In the left panel, a total of 10000 samples are taken, of which 7343 are kept in post-selection. In the right panel, a total of 10000 samples are taken, of which 6631 are kept in post-selection. The blue-plot represents the ideal algorithm output in the absence of hardware noise, while the orange plot represent the exact corresponding normal distribution.
jump-starting circuit has a small constant depth, and thus doesn't significantly increase the overall depth of the circuit. To ensure that the jump-starting circuit need only act on a small number of qubits, it is not recommended to use a value of $k$ greater than $t_{1}$. The theory explaining the benefits of jump starting is clear, and in practice we obtain the post-selection rate of $73.42 \%$ in the left panel of Figure 2 (as opposed to the post-selection rate of $21.28 \%$ in Figure 1), neatly demonstrating the significant benefits the technique incurs. As an additional note, theoretical calculations indicate that in the absence of noise, a 2 -jump start in this experiment would yield a post-selection rate of $\approx 85.16 \%$, approximately an $11.74 \%$ deviation from the experimental results.

We now discuss the right panel of Figure 2. The purpose of this experiment is to demonstrate that our procedure not only produces samples from the correct probability distribution when executed on real quantum hardware, but also produces the correct amplitudes on the quantum states. In particular, this experiment exploits the fact that the Fourier transform of a normal distribution is another normal distribution, with a new mean and a new variance. If the normal distribution produced in the quantum register by the algorithm were correct and free of any phase errors (which would actually be bit-flip errors in Fourier space), then the resulting distribution after a final Fourier transform would be as shown by the blue and orange curves in the right panel of Figure 2. Indeed, sampling after this final Fourier transform yields a probability distribution matching that predicted by the theory, confirming that our algorithm is producing a state free of both bitflip and phase-flip errors (and thus that we are not only producing the correct probability distribution, but also the correct amplitudes). Note that this is consistent with our proof that the algorithm is resistant to both phase-flip and bit-flip errors.


Figure 3: Simulated Experiment Comparing the MCMR-based and MCMR-free Approaches for the Exact and Approximate Methods. The left panel of this figure displays the results obtained in theoretical simulation for the approximate method with $n_{1}=5$ scaling up to $n_{5}=9$, and with $t_{1}=32$ and $t_{2}=t_{3}=t_{4}=t_{5}=4$. The exact approach (i.e. no qubit scaling) is executed using $n=9$ qubits, and with an equivalent effective iteration count of $t=8617$. The normal distribution shown is produced by computing the variance and mean corresponding to $t=8617$, assuming that $l=N$. The panel on the right shows the total number of qubits and the number of ancillary qubits required to execute each configuration of the algorithm. The bars labelled (a) represent the total qubit counts for the approximate algorithm, the bars labelled (b) represent the ancillary qubit counts required for the approximate algorithm, the bars labelled (c) represent the total qubit counts required for the exact algorithm, and the bars labelled ( $d$ ) represent the ancillary qubit counts required for the exact algorithm.

### 3.2 Simulated Experiment

In Figure 3, we compare the exact and approximate techniques, along with their MCMRbased and MCMR-free variants. Here, we call the approach without qubit scaling the exact approach as it produces an exact Binomial distribution, and we call the qubit scaling approach approximate as it does not exactly produce a Binomial distribution (rather it just directly approximates a normal distribution through the central limit theorem). We first configure the approximate algorithm with $n_{1}=5$ scaling up to $n_{5}=9$, and with $t_{1}=32$ and $t_{2}=t_{3}=t_{4}=t_{5}=4$. We run the exact algorithm on $n=9$ qubits, and with an iteration count of $t=8617$. Both approaches are expected to generate a distribution with $\sigma^{2}=2154.25$ and $\mu=256$. Results from the simulations are plotted along with a normal distribution with the same mean and variance. All of the produced curves are shown on top of each other, demonstrating the effective equivalence of all three distributions (as expected, from a central limit theorem approximation when $N=2^{9}$ ). The panel on the right of Figure 3 shows the number of ancillary and total qubits required for the exact and approximate simulators with and without MCMR. The exact simulator requires $8617 \times$ fewer ancillary qubits and $862.6 \times$ fewer total qubits when using MCMR. The approximate simulator requires $35 \times$ fewer ancillary qubits, and $4.4 \times$ fewer total qubits. Clearly, these reductions in required qubit resources will only increase as the variance of the desired distribution increases.

### 3.3 Fidelity Analysis

In Figure 4 we conduct an analysis of the fidelity of the states produced by Algorithm 1 for a constant $c$ (in the absence of hardware error). Error is measured as the KL divergence between the distribution produced on the amplitudes and the corresponding exact normal


Figure 4: Theoretical Fidelity Analysis for Algorithm. This figure shows the error between the distribution produced by the algorithm and the corresponding exact analytical normal distribution. Error is measured as the KL-Divergence between the produced distribution (treating the amplitudes as probabilities) and the exact distribution. The algorithm is configured with $n_{1}=6, n_{2}=7, n_{3}=8, n_{4}=$ $9, n_{5}=10$, and with $t_{1}=t_{2}=t_{3}=t_{4}=t_{5}=c$ for $c$ an integer in the domain [2, 180].
distribution. Note that a similar plot is produced when error is measured as $1-\langle\psi \mid \phi\rangle$ (where $|\psi\rangle$ is the exact distribution and $|\phi\rangle$ is the approximate distribution) only with all errors decreased by about one order of magnitude. We treat the amplitudes as probabilities for the purpose of this calculation, a valid assumption since the algorithm is restricted to the real plane, and each normal distribution in amplitudes is isomorphic to a normal distribution with a different variance in the corresponding probability space. As stated in the caption, we configure the algorithm with $n_{1}=6, n_{2}=7, n_{3}=8, n_{4}=9, n_{5}=10$, and with $t_{1}=t_{2}=t_{3}=t_{4}=t_{5}=c$ for $c$ an integer in the domain [2,180]. The plot shows a linear relationship between the log-error and log-iteration-count. Running a linear regression on the log-log data, we find a line of best fit with slope -1.990334 , an $r$-value of -0.999996 , and a p-value of 0.0 . This supports the conclusion that in the model where a constant $c$ iterations are performed at each qubit count, the error scales as $\Theta\left(\frac{1}{t^{2}}\right)$ (where $t=\sum_{i=1}^{m} t_{i}=c m$ ). It is possible that this error dependence could be improved (potentially asymptotically) by allocating iterations at various qubit counts non-uniformly (i.e. not using a constant $c$ ). However, it is not obvious how this would work, as the error from the central limit theorem approximation to the normal distribution depends on the number of random variables being summed (i.e. the number of applications of $\mathcal{H}$ ) and does not necessarily appear to depend on which qubit count an iteration of the procedure is applied. That is to say, the error of the procedure may simply depend on the total number of iterations performed, irrespective of the schedule for each $t_{i}$ that is selected. As such, to obtain a given variance, it may make sense to perform as many iterations as possible at the final qubit count $n_{m}$ so as to produce the desired fidelity (as iterations at the final qubit count have the least impact on the produced variance). We leave such explorations as topics for future investigations.

It is worth briefly explaining that the only source of error in the qubit-scaling procedure comes from the central limit theorem approximation to the normal distribution. To begin, as proven in the SI, we are clearly producing some distribution with the correct mean and variance (or as close as the selected discretization allows). Thus, the remaining source of error comes from how "close" the produced distribution is to being normal. Clearly, if it
were exactly a normal distribution, there would be no error in the produced distribution (as it would be a normal distribution with the correct mean and variance). However, our distribution is only approximately normal (with the approximation coming from the central limit theorem as we are essentially summing a number of random variables) and thus the error is only that incurred by the central limit theorem approximation.

### 3.4 Probability of Success

When the algorithm is configured such that all of its $t$ iterations are performed at a single qubit count (i.e. without the qubit scaling procedure), the probability of success scales with $\Omega\left(t^{-1 / 2}\right)$ as shown in SI F.1. In the absence of the qubit scaling procedure, Equation 3 clearly states the relationship between $t$ and the user-specified input parameters, allowing us to write this bound on the probability of success as $\Omega\left(\frac{l}{2^{n} \hat{\sigma}}\right)$.

The analysis for the probability of success becomes more complicated when the qubit scaling procedure is utilized. The main insight required to understand the superior scaling of this approach comes from the observation that the probability of success of a single application of $\mathcal{H}$ doubles when $\mathcal{H}$ is applied to the state $|\psi\rangle|+\rangle$ as opposed to just the state $|\psi\rangle$ (where the amplitudes of $|\psi\rangle$ are some discretization of a continuous distribution). As a result, each time the qubit count is increased by adding a qubit in the $|+\rangle$ state as the least-significant qubit, the probability of success of the next application of $\mathcal{H}$ doubles (meaning the probability of failure of any given iteration decays exponentially in the number of qubits added). Moreover, by assumption, $n_{1}$ is selected to be some small constant number of qubits, and thus any value of $t_{1}$ which we select will also be a small constant. Observing that we need only perform a constant-bounded number of iterations at each qubit count in the scaling procedure to obtain an arbitrary variance (since $t_{1}$ at $n_{1}$ gives the bulk of the "shape" of the distribution, and all $t_{j}$ with $j>1$ simply correct the error incurred from adding qubits in the $|+\rangle$ state) we have $t_{2}, \ldots, t_{m}<c$ for some constant $c$ (and again $t_{1}$ is another constant). As a result, Equation 25 from SI F. 2 says that the expected number of trials before the procedure succeeds once, $E[T]$, is constant bounded as $E[T]<2^{t_{\max }+t_{1}}$ (where $t_{\max }=\max _{2 \leq k \leq m} t_{k}$ ) as $t_{\text {max }}$ and $t_{1}$ are both constant bounded. Therefore, the cumulative probability of success is bounded from below as, $\Omega\left(\frac{1}{2^{c+t_{1}}}\right)$. Since this is a constant, it is not necessary to provide this bound in terms of the user-specified input parameters.

To obtain an arbitrary normal distribution, we have shown that our approach requires a constant-bounded number of circuit executions (i.e. has a constant lower-bound on the success rate). However, we note that there is an apparent trade-off between the successrate of the algorithm and the fidelity of the distributions obtained. In particular, as previously discussed, the fidelity of the algorithm roughly depends on the total number of times that $\mathcal{H}$ has been applied. However, our argument for the probability of success relies on the fact that we perform a constant-bounded number of iterations at each qubit count, limiting the ability of the user to fine-tune the fidelity of the obtained normal distribution. We simply observe that each additional iteration at the final qubit count $n_{m}$ has an exponentially small impact on the probability failure, and so the fidelity can be fine tuned by setting $t_{m}$ to a sufficiently large value (which we would not expect to substantially impact the overall probability of success) and then appropriately reducing $t_{i}$ for $i<m$.

Finally, we note one significant observation which can substantially improve the algorithm's overall probability of success. The probability of failure of an application of $\mathcal{H}$ decays exponentially in the number qubits added in the scaling procedure. As a result, the majority of the probability of failure comes from the first few iterations of the algorithm
(namely from $t_{1}$ ). By assumption, $t_{1}$ is applied on a small constant number of qubits $n_{1}$, and as a result we can simply use a deterministic universal quantum state preparation procedure to prepare the exact $n_{1}$ qubit state that would have been produced by $t_{1}$ applications of $\mathcal{H}[9]$. That is to say, we can use a universal quantum circuit to perform a $k$-jump start. Asymptotically, the cost of this procedure would roughly be $O\left(2^{n_{1}}\right)$, which again is a constant.

## 4 Conclusion

This work presents a novel quantum algorithm for the generation of normal distributions in quantum registers, the latter of which runs with $\mathcal{O}\left(n^{2}\right)$ circuit depth (with $n$ being the number of qubits in the output quantum register) and a single ancilla, or with $\mathcal{O}(n \log n)$ circuit depth and $\mathcal{O}(n)$ ancilla qubits. The algorithm uses a repeat-until-success scheme combined with MCMR technology, with a constant-bounded rate of success. Thus, to obtain a normal distribution with this approach, one need only run the procedure a constant number of times in expectation, independent of any of the user input parameters. In addition, this work also demonstrates the potential of MCMR technology in the NISQ era by highlighting how it can enable existing qubit resources to be used asymptotically more efficiently in the introduction of non-unitary transformations, and by demonstrating how it allows qubit-efficient error detection and error mitigation techniques.

## Acknowledgments

We would like to thank Saori Pastore (Washington University in St. Louis) for the insightful discussions on the material presented in this paper. Special thanks to Tony Uttley, Brian Neyenhuis and the rest of the Honeywell Quantum Solutions team for assisting us on the execution of the experiments on the Honeywell System Model H0.

## Disclaimer

This paper was prepared for information purposes by the Future Lab for Applied Research and Engineering (FLARE) group of JPMorgan Chase Bank, N.A.. This paper is not a product of the Research Department of JPMorgan Chase \& Co. or its affiliates. Neither JPMorgan Chase \& Co. nor any of its affiliates make any explicit or implied representation or warranty and none of them accept any liability in connection with this paper, including, but limited to, the completeness, accuracy, reliability of information contained herein and the potential legal, compliance, tax or accounting effects thereof. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction.

## Author Contributions

A.G.R. conceived of the project and designed the proposed algorithms. Y.S., A.G.R., P.M., and M.P. contributed further optimizations to the algorithms, and developed the theory. A.G.R., Y.S., and M.P. implemented the algorithm. All authors contributed to the manuscript.

## References

[1] Almudena Carrera Vazquez and Stefan Woerner. "Efficient State Preparation for Quantum Amplitude Estimation". In: Physical Review Applied 15.3 (Mar. 2021). ISSN: 2331-7019. DOI: 10.1103/physrevapplied.15.034027.
[2] Nikitas Stamatopoulos et al. "Option pricing using quantum computers". In: Quantum 4 (2020), p. 291. DOI: 10.22331/q-2020-07-06-291.
[3] Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. "Quantum principal component analysis". In: Nature Physics 10.9 (2014), pp. 631-633. Doi: 10.1038/nphys3029.
[4] Aram W Harrow, Avinatan Hassidim, and Seth Lloyd. "Quantum algorithm for linear systems of equations". In: Physical review letters 103.15 (2009), p. 150502. DOI: $10.1103 /$ physrevlett.103.150502.
[5] Carlo Ciliberto et al. "Quantum machine learning: a classical perspective". In: Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences 474.2209 (2018), p. 20170551. DOI: 10.1098/rspa.2017.0551.
[6] Kosuke Mitarai et al. "Quantum circuit learning". In: Physical Review A 98.3 (2018), p. 032309. DOI: 10.1103/physreva.98.032309.
[7] Patrick Rebentrost, Masoud Mohseni, and Seth Lloyd. "Quantum support vector machine for big data classification". In: Physical review letters 113.13 (2014), p. 130503. DOI: 10.1103/physrevlett.113.130503.
[8] Marco Pistoia et al. "Quantum Machine Learning for Finance". In: Proceedings of the 40th IEEE/ACM International Conference on Computer Aided Design (ICCAD) 2021. Invited Special Session Paper. Munich, Germany. URL: https://arxiv. org/abs/2109. 04298.
[9] Martin Plesch and Časlav Brukner. "Quantum-state preparation with universal gate decompositions". In: Physical Review A 83.3 (2011), p. 032302. Doi: 10.1103/physreva.83.032302.
[10] John Preskill. "Quantum Computing in the NISQ era and beyond". In: Quantum 2 (2018), p. 79. DOI: $10.22331 / \mathrm{q}-2018-08-06-79$.
[11] Zhaoqing Pan et al. "Recent progress on generative adversarial networks (GANs): A survey". In: IEEE Access 7 (2019), pp. 36322-36333. DOI: 10.1109/ACCESS.2019.2905015.
[12] Christa Zoufal, Aurélien Lucchi, and Stefan Woerner. "Quantum generative adversarial networks for learning and loading random distributions". In: npj Quantum Information 5.1 (2019), pp. 1-9. DOI: 10.1038/s41534-019-0223-2.
[13] Ling Hu et al. "Quantum generative adversarial learning in a superconducting quantum circuit". In: Science advances 5.1 (2019), eaav2761. DoI: 10.1126/sciadv.aav2761.
[14] Lov Grover and Terry Rudolph. "Creating superpositions that correspond to efficiently integrable probability distributions". In: arXiv preprint (2002). URL: https: //arxiv.org/abs/quant-ph/0208112.
[15] Adam Holmes and AY Matsuura. "Efficient Quantum Circuits for Accurate State Preparation of Smooth, Differentiable Functions". In: arXiv preprint (2020). URL: https://arxiv.org/abs/2005.04351.
[16] Alexei Kitaev and William A Webb. "Wavefunction preparation and resampling using a quantum computer". In: arXiv preprint arXiv:0801.0342 (2008). URL: http s://arxiv.org/abs/0801.0342.
[17] Thomas Häner, Martin Roetteler, and Krysta M Svore. "Optimizing quantum circuits for arithmetic". In: arXiv preprint (2018). URL: https://arXiv.org/abs/ 1805.12445.
[18] Bahram Ahansaz and Abbas Ektesabi. "Quantum speedup, non-Markovianity and formation of bound state". In: Scientific reports 9.1 (2019), pp. 1-12. DOI: 10.1038/s41598-019-51290-x.
[19] Viv Kendon and Olivier Maloyer. "Optimal computation with non-unitary quantum walks". In: Theoretical computer science 394.3 (2008), pp. 187-196. DOI: 10.1016/j.tcs.2007.12.011.
[20] J. M. Pino et al. "Demonstration of the trapped-ion quantum CCD computer architecture". In: Nature 592.7853 (Apr. 2021), pp. 209-213. ISSN: 1476-4687. DOI: 10.1038/s41586-021-03318-4.
[21] Romina Yalovetzky et al. NISQ-HHL: Portfolio Optimization for Near-Term Quantum Hardware. 2021. URL: https://arxiv.org/abs/2110.15958.
[22] Adam Paetznick and Krysta M Svore. "Repeat-Until-Success: Non-deterministic decomposition of single-qubit unitaries". In: arXiv preprint (2013). URL: https : // arXiv.org/abs/1311.1074.
[23] Yudong Cao, Gian Giacomo Guerreschi, and Alán Aspuru-Guzik. "Quantum neuron: an elementary building block for machine learning on quantum computers". In: arXiv preprint (2017). URL: https://arXiv.org/abs/1711.11240.
[24] Wei Hu. "Towards a real quantum neuron". In: Natural Science 10.3 (2018), pp. 99109. DOI: 10.4236/ns.2018.103011.
[25] Alex Bocharov, Martin Roetteler, and Krysta M Svore. "Efficient synthesis of universal repeat-until-success quantum circuits". In: Physical review letters 114.8 (2015), p. 080502. DOI: 10.1103/physrevlett.114.080502.
[26] Nathan Wiebe and Martin Roetteler. "Quantum Arithmetic and Numerical Analysis Using Repeat-until-Success Circuits". In: Quantum Info. Comput. 16.1-2 (Jan. 2016), pp. 134-178. ISSN: 1533-7146. DOI: 10.26421/QIC16.1-2-9.
[27] Sam McArdle et al. "Variational ansatz-based quantum simulation of imaginary time evolution". In: npj Quantum Information 5.1 (2019), pp. 1-6. DOI: 10.1038/s41534-019-0187-2.
[28] Sheldon Natenberg. Option volatility and pricing: Advanced trading strategies and techniques. McGraw-Hill Education, 2014. ISBN: 9780071818780.
[29] Michael A Nielsen and Isaac Chuang. Quantum computation and quantum information. 2002. DOI: 10.1017/CBO9780511976667.
[30] Thomas G Draper. "Addition on a quantum computer". In: arXiv preprint (2000). URL: https://arxiv.org/abs/quant-ph/0008033.
[31] Thomas G Draper et al. "A logarithmic-depth quantum carry-lookahead adder". In: arXiv preprint (2004). URL: https://arxiv.org/abs/quant-ph/0406142.
[32] Yunseong Nam, Yuan Su, and Dmitri Maslov. "Approximate quantum Fourier transform with O (n log (n)) T gates". In: NPJ Quantum Information 6.1 (2020), pp. 1-6. DOI: 10.1038/s41534-020-0257-5.
[33] Hayato Goto. "Resource requirements for a fault-tolerant quantum Fourier transform". In: Physical Review A 90.5 (2014), p. 052318. Doi: 10.1103/PhysRevA.90.052318.
[34] Stephane Beauregard. Circuit for Shor's algorithm using 2n+3 qubits. 2002. urL: https://arxiv.org/abs/quant-ph/0205095.
[35] Adriano Barenco et al. "Elementary gates for quantum computation". In: Physical review $A 52.5$ (1995), p. 3457. DOI: 10.1103/physreva.52.3457.
[36] Hale F Trotter. "On the product of semi-groups of operators". In: Proceedings of the American Mathematical Society 10.4 (1959), pp. 545-551. Doi: 10.1090/S0002-9939-1959-0108732-6.
[37] Ying Li and Simon C Benjamin. "Efficient variational quantum simulator incorporating active error minimization". In: Physical Review X 7.2 (2017), p. 021050. doi: 10.1103/physrevx.7.021050.

## Supplementary Information: <br> Approximate and Exact Quantum Simulation of Galton Machines

## A Motivating the Algorithm

Continuing with the notation introduced in Section 2, we now provide the intuition motivating the state generating algorithm. According to the de Moivre-Laplace theorem, a binomial distribution $B(t, p)$ converges to a normal distribution with mean $t p$ and standard deviation $\sqrt{t p(1-p)}$ as $t \rightarrow \infty$. This property is exploited by Galton machines, such as the one shown in Figure 5, to generate normal distributions by dropping balls through a sequence of rows, where in each row the falling ball may have its position shifted by either one bin to the left, or one bin to the right. As a result, the transition dynamics of a single row in a Galton machine may be described by the transition matrix $\mathcal{H}$ defined as (up to a normalization factor)

$$
\mathcal{H}|j\rangle=|j-1\rangle+|j+1\rangle .
$$

As previously mentioned, the quality of the approximation of the normal distribution produced by a Galton machine improves as more balls are dropped through the system. Moreover, as our quantum Galton approach simulates an infinite number of balls falling through the system (by virtue of operating on superpositions of "bins"), we avoid this type of error entirely. As a result, combined with the exponential growth in the number of bins in terms of the number of qubits, the distributions we produce are indistinguishable from the target normal distributions (i.e. our only source of error is that incurred by the central limit theorem approximation, and the central limit theorem error vanishes exponentially in the number of qubits, and so our approximation rapidly becomes exact).

It is worth briefly mentioning that we explored another variant of the algorithm simulating $e^{\mathcal{H} t}$ (instead of $\mathcal{H}^{t}$ ) to obtain the desired normal distribution. This approach utilizes Imaginary Time Evolution (ITE) to simulate $e^{\mathcal{H} t}$, as proposed by McArdle et al. in 2019 [27], and is discussed in greater depth in SI H.

## B Implementing $\mathcal{H}$ as a Quantum Circuit with Mid-Circuit Measurement and Reuse (MCMR)

In order to make the implementation of $\mathcal{H}$ as a quantum circuit more straight forward, we will redefine the transition matrix. Instead of performing the transformation $\mathcal{H}|j\rangle=$


Figure 5: Galton Machine Visualization. Visual demonstration of a Galton machine generating a normal distribution. Figure adapted from Option Pricing and Volatility, by Sheldon Natenberg [28].
$|j-1\rangle+|j+1\rangle$, we will now perform the mapping $\mathcal{H}|j\rangle=|j\rangle+|j+1\rangle$. In so doing, we must pay additional attention to ensure that the mean of the generated distribution is configured correctly, and in exchange, we reduce the complexity of certain sections of the analysis. First, observe that this transition matrix still generates a normal distribution in the same way as the first matrix, only that instead of alternating between zero amplitude even and odd states with each application of $\mathcal{H}$, all states are utilized without the need for the input state to contain a superposition of adjacent states (or without the need of additional ancillary qubits). We will now define an addition operator, $A_{n+1}^{+1}$, where the subscript indicates the control qubit, and the superscript indicates the quantity added to each state in the first $n$-qubit register. For example, $A_{n+1}^{+1}|j\rangle|1\rangle=|j+1\rangle|1\rangle$, and $A_{n+1}^{+1}|j\rangle|0\rangle=|j\rangle|0\rangle$. We then derive the quantum circuit as follows. First we apply a Hadamard gate, $H$, on the ancilla qubit,

$$
I^{\otimes n} \otimes H|j\rangle|0\rangle=|j\rangle(|0\rangle+|1\rangle)
$$

Then we apply the +1 gate on the main register, controlled on the ancilla,

$$
A_{n+1}^{+1}|j\rangle(|0\rangle+|1\rangle)=|j, 0\rangle+|j+1,1\rangle
$$

Finally, we apply another Hadamard gate on the ancilla,

$$
\begin{aligned}
& I^{\otimes n} \otimes H(|j, 0\rangle+|j+1,1\rangle) \\
& =|j, 0\rangle+|j, 1\rangle+|j+1,0\rangle-|j+1,1\rangle \\
& =(|j\rangle+|j+1\rangle)|0\rangle+(|j\rangle-|j+1\rangle)|1\rangle
\end{aligned}
$$

We see that upon measuring the ancilla qubit in the $|0\rangle$ state, the main register will contain the state $|j\rangle+|j+1\rangle$ as desired. However, if we measure the ancilla in the $|1\rangle$ state, the main register will be found in the incorrect state $|j\rangle-|j+1\rangle$. Initially, methods for correcting this error were explored, however, we realized that $|j\rangle-|j+1\rangle$ actually represents a pattern of destructive interference that exponentially approaches net zero amplitude as a function of the number of iterations applied. As a result, a single iteration of the circuit implementation of $\mathcal{H}$ may be produced by following the procedure just described, and by measuring the ancilla qubit, continuing if the desirable ancilla


Figure 6: Amplitude Cancellation Demonstration. This figure demonstrates the interference pattern produced when an initial distribution $f(x)$ is mapped to $f(x)-f(x+1)$ at various qubit counts.
state is measured, and terminating execution if the undesirable state is measured. As will be proved shortly, this procedure requires an increase in the expected number of circuit evaluations growing sub-linearly in the number of iterations performed, and therefore results in a polynomially bounded total number of circuit executions for the algorithm as a whole. These claims are motivated in SI B.1, and proven in SI F. In contrast, in the MCMR-free approach, as the single ancilla is not reused subsequently to being measured, a new ancilla qubit must be added for each application of $\mathcal{H}$. Each ancilla may be measured at the end of the application of $\mathcal{H}$, or by the principle of deferred measurement [29], all added ancillas may be simultaneously measured at the end of the circuit's execution. A more detailed explanation of the MCMR-free approach is presented in SI C.

## B. 1 Ancilla $|1\rangle$ State Cancellation Analysis

The implementation of $\mathcal{H}$ just described relies upon performing a partial measurement on the state, $(|j\rangle+|j+1\rangle)|0\rangle+(|j\rangle-|j+1\rangle)|1\rangle$. If the ancilla is measured in the $|0\rangle$ state, the desired transformation yielding $|j\rangle+|j+1\rangle$ has been obtained in the main register. If the ancilla is measured in the $|1\rangle$ state, the undesired state of $|j\rangle-|j+1\rangle$ is obtained. We claim, and prove in SI F, that the probability of measuring the $|1\rangle$ ancillary state rapidly vanishes as a function of the number of applications of $\mathcal{H}$ performed, and that as a result, we only need to repeat the algorithm's execution a number of times sub-linear in the number of applications of $\mathcal{H}$ (and in the case of the qubit scaling approach, a constantbounded number of times). To understand this, it may be more clear to understand an equivalent statement. Given a distribution on $n$ qubits obtained by computing $\mathcal{H}^{t}$, with $t$ computed taking into account $n$ and the input variance, the behavior of the mapping $|j\rangle \rightarrow|j\rangle-|j+1\rangle$ may be similarly expressed in terms of functions of the amplitudes as $f\left(x_{j}\right) \rightarrow f\left(x_{j}\right)-f\left(x_{j}+\Delta x\right)$, where $\Delta x=\frac{l}{2^{n}}$. As such, given a fixed variance $\hat{\sigma}^{2}$, as we increase the number of qubits, we are exponentially decreasing the relative shift between distributions $f\left(x_{j}\right)$ and $-f\left(x_{j}+\Delta x\right)$, and so the two distributions rapidly approach $f\left(x_{j}\right)$ and $-f\left(x_{j}\right)$, meaning that the interference pattern on the $|1\rangle$ ancilla state exponentially approaches complete destructive interference as the number of qubits increase. Indeed, in SI G we have shown that this cancellation property holds for any continuous function $f$. This principle is clearly illustrated in Figure 6, where the interference pattern produced rapidly approaches complete destructive interference as the number of qubits increase up to 8 .


Figure 7: Quantum circuit implementation of $\mathcal{H}$, with the definition $\mathcal{H}|j\rangle=|j\rangle+|j+1\rangle$ (up to a normalization factor). Note that circuit execution only proceeds if the ancilla is measured in the $|1\rangle$ state.


Figure 8: Quantum circuit implementing the quantum Galton machine with qubit scaling.

## B. 2 Quantum Circuit Implementation of $\mathcal{H}$

Again, the mapping we wish to perform is given by $\mathcal{H}|j\rangle=|j\rangle+|j+1\rangle$. To perform such a mapping, we use four components: the QFT, the (Fourier space) adder gate, the Hadamard gate, and mid circuit measurement and reuse. Assuming we are in the Fourier space, (i.e. a QFT has already been performed previously in the circuit), $\mathcal{H}$ is implemented as shown in Figure 7. As the adder gate may be implemented with $O(n)$ two-qubit gates and depth, one iteration according to $\mathcal{H}$ is also implemented in linear depth and with linear two-qubit gates. The adder gate and its quantum circuit, along with the corresponding controlled adder gate, is presented in SI C. 2 and SI C.3.

## B. 3 Algorithm Quantum Circuit Implementation

We first present a basic circuit used to implement $\mathcal{H}^{t}$, proceed to give the implementation of the full quantum circuit, including the qubit scaling procedure. The naive circuit (which is used for the exact simulation of Galton machines) uses two quantum registers: one with $n=n_{m}$ qubits for the output and the other 1-qubit register as the ancilla. The two registers are initialized in joint state $|0\rangle_{n+1}$. We then apply QFT on the $n$-qubit register to bring the quantum register into Fourier space. Subsequently, we apply $\mathcal{H}$ which consists of a Hadamard gate on the ancilla qubit, followed by a controlled +1 gate $\left(A_{j}^{+1}\right)$ controlled on the ancilla, a second Hadamard gate on the ancilla, and a measurement on the ancilla. The application of $\mathcal{H}$ is repeated $t$ times, with the value of $t$ given by Equation 3. In the end, we apply an inverse-QFT on the output register to go back to the real space. The success of the circuit depends on the measurement results on the ancilla qubit. We repeat the circuit until all measurement results come out as 0 .

As discussed at the beginning of Section 2, the naive implementation without qubit scaling described above would require exponential number of $\mathcal{H}$ applications, which renders the algorithm inefficient. To address this issue, we use a qubit scaling scheme which significantly reduces the circuit depth with only nominal sacrifices in accuracy. The circuit is illustrated in Figure 8. We start with a small, constant, number of qubits $n_{1}$, and apply
$\mathcal{H} t_{1}$ times to get close to the desired variance of the distribution. Then we add one qubit to the least significant end in the sate $|+\rangle$, and apply $\mathcal{H} t_{2}$ times. We repeat the process of adding qubits and applying $\mathcal{H}$ untill we reach the desired number of qubits $n_{m}$. The values of $t_{1}, t_{2}, \ldots, t_{m}$ are determined by the scaled variance $\hat{\sigma}^{2}$ of the distribution, the formulas of which are given in SI D. Note that the total number of applications of $\mathcal{H}$, i.e. $\sum_{r=1}^{m} t_{r}$ scales as $O\left(n_{m}\right)$, as discussed at the beginning of Section 2, and in SI D.

To avoid going back and forth between the real space and the Fourier space when increasing the number of qubits, we perform QFT on the entire $n_{m}$ qubit register and inverse-QFT at the very end. Accordingly, the controlled +1 gates in the $r$-th stage with $n_{r}$ qubits would be changed to a controlled $+2^{m-r}$ gate on the $n_{m}$ qubit register. Note that this still amounts to $n_{r}$ controlled- $U 1$ operations in the controlled adder gate, as $U 1$ gates with angles that are multiples of $2 \pi$ amount to the identity gate.

In the case where the mean of the outcome distribution needs to be adjusted, an additional adder gate $+x_{0}$ may be added on the $n_{m}$-qubit register in the Fourier space, as shown in Figure 8.

## C Implementing $\mathcal{H}$ as a Quantum Circuit without Mid-Circuit Measurement and Reuse (MCMR-free)

We now describe the implementation of the transition matrix $\mathcal{H}$ defined as $\mathcal{H}|j\rangle=$ $|j\rangle+|j+1\rangle$ without the use of Mid-Circuit Measurement and Reuse technology. We first perform a derivation similar to that performed in SI B. Examining one application of $\mathcal{H}$,

$$
|j\rangle|+\rangle \xrightarrow{A_{n+1}^{+1}}|j\rangle|0\rangle+|j+1\rangle|1\rangle \xrightarrow{I^{\otimes n} \otimes H}(|j\rangle+|j+1\rangle)|0\rangle+(|j\rangle-|j+1\rangle)|1\rangle
$$

When mid-circuit measurement is used, the preceding final quantum state may simply be measured, as described in a preceding SI entry. However, when MCMR technology is not available, the ancilla cannot be measured until the end of the circuit execution, and so to perform another iteration by $\mathcal{H}$ we must add another ancillary qubit so as to preserve the state of the first ancilla. In order to reveal some of the properties of this MCMR-free approach, we will now examine the state of the system after performing a second application of $\mathcal{H}$. First, we add another ancillary qubit initialized in the $|+\rangle$ state to obtain,

$$
(|j\rangle+|j+1\rangle)|0\rangle|+\rangle+(|j\rangle-|j+1\rangle)|1\rangle|+\rangle
$$

We subsequently apply the adder +1 gate conditioned on the second ancilla, $A_{n+2}^{+1}$, to obtain the state,
$(|j\rangle+|j+1\rangle)|00\rangle+(|j+1\rangle+|j+2\rangle)|01\rangle+(|j\rangle-|j+1\rangle)|10\rangle+(|j+1\rangle-|j+2\rangle)|11\rangle$.
Simplifying and applying $I^{\otimes n+1} \otimes H$ then yields,

$$
\begin{aligned}
& (|j\rangle+2|j+1\rangle+|j+2\rangle)|00\rangle+(|j\rangle-|j+2\rangle)|01\rangle \\
& +(|j\rangle-|j+2\rangle)|10\rangle-(|j\rangle-2|j+1\rangle+|j+2\rangle)|11\rangle
\end{aligned}
$$

It now becomes clear that upon adding a new ancillary qubit, and performing the standard operations required to implement an additional iteration according to $\mathcal{H}$, that the probability of measuring the $|1\rangle$ state of any preceding ancilla qubits remains unchanged. As such, the analysis regarding the probability of measuring all ancilla qubits in the $|0\rangle$ state, and thus of obtaining the correct state in the primary register, follows the same analysis as the MCMR-based version of the algorithm.

## C. 1 Adder Gate Complexity

All approaches require the implementation of an adder-plus-one gate $\left(A^{+1}\right)$, which can be implemented either with additional ancillary qubits, or with no additional ancillary qubits. In the NISQ context, Table 1 and Table 2 will provide estimate on the constant factor of the asymptotic circuit depth for different approaches described for the QFT and adder gate respectively. Note as we already know one of the two number in the additions, further optimization could be found in the described methods. When not using ancillary qubits, the adder-plus-one gate may be implemented in Fourier space (requiring a Quantum Fourier Transform (QFT) at the start of the overall circuit, and an inverseQFT at the end) resulting in $\mathcal{O}(n)$ CX gates per application of $\mathcal{H}$ [30]. Alternatively, without additional ancillas, the adder-plus-one gate may be implemented with $\mathcal{O}\left(n^{2}\right) \mathrm{CX}$ gates per application of $\mathcal{H}$. As a result, we use the QFT adder approach throughout this document, resulting in a circuit depth of $\mathcal{O}\left(n^{2}+n t\right)$, which due to the qubit-scaling procedure already discussed is bounded by $\mathcal{O}\left(n^{2}\right)$. Alternatively, Draper et al. present an adder requiring $\mathcal{O}(n)$ ancilla qubits, and scaling with $\mathcal{O}(\log n)$ circuit depth, without the need to enter the Fourier space, enabling our algorithm to scale with $\mathcal{O}(t \log n)$ circuit depth, if $\mathcal{O}(n)$ ancilla qubits are available [31]. One could adapt this space-depth trade-off to the specific hardware used. This reduction in depth would not only be seen on total complexity but also asymptotic complexity if we were to consider the approximate QFT [32]. In a fault-tolerance context and with an error $\epsilon$, we would require for the adder and QFT, $\mathcal{O}\left(\left(n t+n^{2}\right) \log (1 / \epsilon)\right)$ gates or just $\mathcal{O}((n t+n) \log (1 / \epsilon))$ gates using the approximate QFT [33].

| Component | QFT | Approximate QFT |
| :---: | :---: | :---: |
| Depth | $\mathcal{O}\left(n^{2}\right)$ | $\mathcal{O}(n \log n)$ |
| Constant estimate | 2 | 30 |
| Number of qubits | $n$ | $n+3\lceil\log n\rceil-4$ |

Table 1: QFT in NISQ context

| Component | Phase Adder | Draper |
| :---: | :---: | :---: |
| Depth | $\mathcal{O}(t n)$ | $\mathcal{O}(t \log n)$ |
| Constant estimate | 2 | 4 |
| Number of qubits | $2 n$ | $3 n$ |

Table 2: $t$ adders in NISQ context

## C. 2 The Adder Gate

The circuit implementation of the $+d$ gate is shown in Figure 9. The general form of the $+d$ gate, with $d$ being an integer, is $\bigotimes_{k=1}^{n} U 1\left(2 \pi d / 2^{k}\right)$, where $n$ is the number of qubits in the main register [34]. Moreover, as our adder gates are conditioned on the state of the ancillary qubit, the general form of the controlled adder gate $A_{j}^{+d}$ is given by,

$$
A_{j}^{+d}=\prod_{k=1}^{n} U 1_{j}^{k}\left(2 \pi \frac{d}{2^{k}}\right)
$$

where a $U 1_{a}^{b}$ gate is a $U 1$ gate with a control on qubit $a$, a target on qubit $b$, and identity on all other qubits. Moreover, a U1 gate is defined as follows,

$$
U 1(\lambda)=\left(\begin{array}{cc}
1 & 0 \\
0 & e^{i \lambda}
\end{array}\right)
$$

and is equivalent to an $R_{z}(\lambda)$ gate up to a global phase of $e^{i \frac{\lambda}{2}}$, as defined by Barenco et al. in 1995 [35]. The circuit for $A_{j}^{+d}$ is shown in Figure 10. For a more comprehensive discussion of the adder gate, see SI C.3.

## C. 3 The General Adder Gate

Let $x \in\{0,1\}^{n}$. When applying the QFT to such a binary-labelled standard basis vector, we obtain:

$$
\frac{1}{2^{n / 2}} \bigotimes_{l=0}^{n-1}\left(|0\rangle+e^{2 \pi i x \frac{2^{l}}{2^{n}}}|1\rangle\right)
$$

Thus, the action of the +1 gate is given by:

$$
\begin{aligned}
& \frac{1}{2^{n / 2}} \bigotimes_{l=1}^{n} U 1\left(2 \pi i \frac{2^{l}}{2^{n}}\right)\left(|0\rangle+e^{2 \pi i x \frac{2^{l}}{2^{n}}}\right)|1\rangle \\
& =\frac{1}{2^{n / 2}} \bigotimes_{l=1}^{n}\left[|0\rangle+\exp \left(2 \pi i(x+1) \frac{2^{l}}{2^{n}}\right)|1\rangle\right]
\end{aligned}
$$

A subsequent application of the inverse QFT would yield the desired state of $|x+1\rangle$. Note that multiple adder gates may be chained together while in the amplitude space, meaning that only one QFT would be required at the start of the circuit, and only one inverse QFT is requried at the end of the circuit. It is also worth noting that the periodicity of the imaginary exponential implies that this addition is actually addition modulo $2^{n}$, and so our $\mathcal{H}$ operation encounters a boundary condition when there are non-zero amplitude states $|x=0\rangle$ or $\left|x=2^{n}\right\rangle$ (in practice, the boundary condition only becomes significant when the amplitudes of these extreme states are much greater than zero). However, it would not be expected for any algorithm to be able to circumvent this limitation, as it is intrinsic to the finite quantum mechanical system. Finally we can model a $A^{+k}$ gate (adder $+k$ ) as follows,

$$
A^{+k}=\bigotimes_{l=1}^{n} U 1\left(2 \pi i \frac{2^{l}}{2^{n}} k\right)
$$



Figure 9: Quantum circuit implementing the Fourier space adder gate $+d$.


Figure 10: Quantum circuit implementing the Fourier space controlled adder gate $A_{\text {control }}^{+d}$.

## D Efficient Scaling Through Qubit Addition

As discussed, if we wish to execute $\mathcal{H}^{t}$ on the quantum processor, naively we would expect that to obtain a given variance we require a number of applications of $\mathcal{H}$ that grow exponentially in the number of qubits in the final quantum register, $n_{m}$, in accordance with Equation 3. However, instead of performing all $t$ iterations in the space of $n_{m}$ qubits, we can produce a distribution with the same variance $\hat{\sigma}^{2}$ on a small, constant number of qubits. In doing so, we essentially produce a "low-resolution" version of the distribution that has the overall correct shape of the final distribution. We subsequently add a qubit in the $|+\rangle$ state to the least significant qubit, apply a number of correction iterations implemented directly as the $\mathcal{H}$ operator, and repeat this addition and iteration procedure $n_{m}-n_{1}$ times. To understand how applying $\mathcal{H}$ some number of times increases the resolution of the distribution after adding the $|+\rangle$ qubit, it is important to better understand the distribution obtained after the addition. Figure 11 demonstrates both the distribution obtained immediately after adding a qubit in the $|+\rangle$ state (in the panel on the left), and the distribution obtained after adding a qubit in the $|+\rangle$ state followed by performing two correction iterations (directly implemented as applications of $\mathcal{H}$ ). As expected, by virtue of mapping the amplitude of state $|x\rangle$ to states $|2 x\rangle$ and $|2 x+1\rangle$, each pair of two states shown in the figure have the same amplitude, resulting in the step-pattern shown. The panel on the right illustrates that it only takes two correction iterations to almost perfectly match the exact distribution. This is consistent with the theory discussed shortly. Moreover, while not shown in the figure, applying a couple more correction iterations would make the scaling approximation and exact distributions completely indistinguishable.

First, we define the $i$-qubit grid as, $\left\{X_{i}\right\}$. Then, the grid on $i+1$-qubits (obtained simply by adding a qubit in the $|+\rangle$ state) is given by, $\left\{2 X_{i}, 2 X_{i}+1\right\}$. Allow the $i$-qubit mean to given by $\mu_{i}$. Upon doubling the number of qubits (with adjacent even and odd states sharing their amplitudes), the mean becomes $\mu_{i+1}=2 \mu_{i}+0.5$. From the definition of variance, we have

$$
\sigma_{i}^{2}=\sum_{x \in X_{i}}\left(x-\mu_{i}\right)^{2} P(x) .
$$



Figure 11: Demonstration of Correction Iterations in the Efficient Scaling Procedure. Two Gaussians generated by the algorithm on 6 qubits are shown. The scaling approximation curve in the left panel is obtained with $t_{1}=70$ and $t_{2}=0$, meaning that no correction iterations are performed after the scaling qubit is added to the least significant bit. The scaling approximation curve in the right panel is obtained with $t_{1}=70$ and $t_{2}=2$, meaning that two correction iterations are applied after the scaling qubit is added. In both panels, the exact curve is generated by applying $t_{\text {eff }}$ iterations on only the final qubit count (meaning that the scaling procedure is not used), where $t_{\text {eff }}=281$ in the panel on the left, and $t_{\text {eff }}=283$ in the panel on the right.

Thus we can compute the variance of the resulting distribution with,

$$
\begin{aligned}
\sigma_{i+1}^{2} & =\sum_{x \in X_{i+1}}\left(x-\mu_{i+1}\right)^{2} P(x) \\
& =\frac{1}{2}\left(\sum_{x \in X_{i}}\left(2 x-\mu_{i+1}\right)^{2} P(x)+\sum_{x \in X_{i}}\left(2 x+1-\mu_{i+1}\right)^{2} P(x)\right) \\
& =\frac{1}{2} \sum_{x \in X_{i}}\left[\left(2 x-2 \mu_{i}-0.5\right)^{2}+\left(2 x-2 \mu_{i}+0.5\right)^{2}\right] P(x) \\
& =\sum_{x \in X_{i}}\left[4\left(x-\mu_{i}\right)^{2}+0.5^{2}\right] P(x) \\
& =4\left[\sum_{x \in X_{i}}\left(x-\mu_{i}\right)^{2} P(x)\right]+0.5^{2},
\end{aligned}
$$

Therefore,

$$
\begin{equation*}
\sigma_{i+1}^{2}=4 \sigma_{i}^{2}+0.25 \tag{7}
\end{equation*}
$$

Moreover, define $t_{i}$ to be the number of iterations performed on $n_{i}$ qubits, with $n_{1} \leq n_{i} \leq$ $n_{m}$, and let $\left|\psi_{i}\right\rangle$ be the state obtained on $n_{i}$ qubits after $t_{i}$ iterations have been performed. Thus, assuming that $\mathcal{H}$ is acting on the appropriately sized quantum system $\left|\psi_{i+1}\right\rangle$ and $\left|\psi_{i}\right\rangle$ are related as follows,

$$
\left|\psi_{i+1}\right\rangle=\mathcal{H}^{t_{i+1}}\left|\psi_{i}\right\rangle|+\rangle
$$

Moreover, the variance of the distribution represented by $\left|\psi_{i}\right\rangle|+\rangle$ is stated in Equation 7. First, the grid on $n_{i}$ qubits is constant, given by $\left\{X_{i}\right\}$. Before applying $\mathcal{H}$, our distribution
is simply over $\left\{X_{i}\right\}$, with a mean and variance given by

$$
\mu_{i}=\sum_{x \in X_{i}} x P(x), \sigma_{i}^{2}=\sum_{x \in X_{i}}\left(x-\mu_{i}\right)^{2} P(x)
$$

Upon applying $\mathcal{H}$, we obtain two distributions, one equal to the original distribution, and one equal to the original distribution shifted by one state to the right. The probability mass function of the new distribution would therefore be given by,

$$
F(x)=p P(x-1)+(1-p) P(x)
$$

Therefore, the mean of the new distribution would be given by,

$$
\mu_{i}^{\prime}=\sum_{x \in X_{i}} x p P(x-1)+x(1-p) P(x)=p\left(\mu_{i}+1\right)+(1-p) \mu_{i}=p+\mu_{i}
$$

and the variance would be given by,

$$
\begin{aligned}
\sigma_{i}^{2 \prime}= & \sum_{x \in X_{i}}\left(x-\mu_{i}^{\prime}\right)^{2} F(x) \\
= & \sum_{x \in X_{i}}\left(x-\mu_{i}-p\right)^{2}[p P(x-1)+(1-p) P(x)] \\
= & p \sum_{x \in X_{i}}\left(x-\mu_{i}\right)^{2} P(x-1)+(1-p) \sum_{x \in X_{i}}\left(x-\mu_{i}\right)^{2} P(x) \\
& +\sum_{x \in X_{i}}\left[-2 x p+2 \mu_{i} p+p^{2}\right][p P(x-1)+(1-p) P(x)] \\
= & \sigma_{i}^{2}+p \sum_{x \in X_{i}} P(x)-2 p \sum_{x \in X_{i}} p P(x)+\sum_{x \in X_{i}} p^{2} P(x) \\
= & \sigma_{i}^{2}+p(1-p) .
\end{aligned}
$$

Thus, the mean and variance on the $n_{i}$ qubits corresponding to the distribution obtained after $t_{i}$ iterations of $\mathcal{H}$ are given by

$$
\begin{align*}
\mu_{i}^{\prime}=\mu_{i}+t_{i} p & =2 \mu_{i-1}+0.5+t_{i} p  \tag{8}\\
\sigma_{i}^{2 \prime}=\sigma_{i}^{2}+t_{i} p(1-p) & =4 \sigma_{i-1}^{2}+0.25+t_{i} p(1-p) \tag{9}
\end{align*}
$$

Or equivalently,

$$
\begin{gathered}
\mu_{i}^{\prime}=0.5 \times\left(2^{i-1}-1\right)+\sum_{k=1}^{i} 2^{i-k} p t_{k} \\
\sigma_{i}^{2 \prime}=0.25 \times \frac{4^{i-1}-1}{3}+\sum_{k=1}^{i} 4^{i-k} p(1-p) t_{k}
\end{gathered}
$$

It is straight forward to use this equation to both compute the effective variance obtained after performing a number of iterations at various qubit counts, and to compute the iterations required at various qubit counts to obtain a given variance. It is important to emphasize that Equation 9 implies that the qubit scaling procedure incurs no additional error beyond that already incurred by the central limit theorem approximation to the normal distribution, as it allows for an arbitrary exact variance to be obtained. Intuitively, if you wish to obtain a distribution corresponding to $t$ iterations performed on $n_{1}$ qubits, if you first performed $t$ iterations on the $n_{1}$ qubits, and then performed some number of iterations on subsequent qubit counts, the total variance produced would exceed the
desired variance. This expression describes how the number of iterations performed on the first qubit count may be reduced, and the number of removed iterations may then be translated to a number of correction iterations at higher qubit counts, allowing for the exact desired distribution to be obtained whilst still allowing for the correction iterations in the qubit scaling procedure to be made. Finally, Equation 9 shows that an arbitrary variance on $n$-qubits may be obtained while setting each $t_{i}$ to essentially an arbitrary constant value, so long as $t_{1}$ is set to give the desired variance on $n_{1}$ qubits. As such, it is clear that this approach implies a total number of iterations scaling polynomially in the number of qubits to obtain an arbitrary variance.

## E Resistance to Hardware Noise

The variant of the algorithm utilizing MCMR is also expected to be resistant to both bitflip and, to a lesser extent, phase-flip errors by automatically discarding results in which such errors occur with high probability. The subsequent analysis assumes a simple noise model where errors in the execution of the algorithm are dominated by the fidelity of twoqubit gate executions. Allow $\epsilon$ to represent the two-qubit gate infidelity such that when executing any given two-qubit gate either a $X$ or a $Z$ gate (implementing a bit-flip and phase-flip error, respectively) gets subsequently randomly executed on either the control or target qubit. The probability of executing $d$ such gates without obtaining any errors is therefore given by $(1-\epsilon)^{d}$.

For expository clarity, the following argument will concern itself with logical operations, such as mapping the state $|j\rangle$ to $|j+1\rangle$, and we will not consider the Fourier basis entered by the QFT. Precisely, we will assume that the QFT and inverse QFT are applied in a noiseless channel, and that in-between these two operations the channel is subject to both phase-flip and bit-flip errors. As a result, a bit-flip error occurring in the hardware is experienced as a phase-flip error in the conceptual analytical space, and a phase-flip error on the hardware is similarly experienced as a bit-flip error in the conceptual space. This is similar to how a bit-flip error correcting code may be applied to correct a phase-flip error by applying a set of Hadamard transforms before and after the noisy channel.

Finally, we model the execution of a single iteration of $\mathcal{H}$ as a single quantum gate, and so an error occurring in the execution of $\mathcal{H}$ is modelled as one of $Z_{j}$ or $X_{j}$ being executed after $\mathcal{H}$ with $j$ selected uniformly at random. Note that here the notation $A_{j}$ means gate $A$ is executed on qubit $j$, and an identity gate is executed on all other qubits. Moreover, we assume that such errors occur with probability equal to the probability of any of the $2 n$ CX gates implementing $\mathcal{H}$ experiencing an error. Thus, the infideltiy of the $\mathcal{H}$ operation is given by $1-(1-\epsilon)^{2 n}$. In the following analysis, we will use this model to examine the impact of the first occurrence of an error after $t$ successful applications of $\mathcal{H}$. We will not examine the case where multiple errors occur in detail, but a similar argument holds in such cases nevertheless. We will now proceed by considering the two types of conceptual errors separately.

Analysis of Phase-Flip Errors Throughout the remainder of this document, unless we explicitly state otherwise, we allow $\{|x\rangle\}_{x}$ to represent the set of integer-valued standard basis vectors (i.e. $x \in\{0,1\}^{n}$ ), identically to how we previously used $\{|j\rangle\}_{j}$. We now assume that we are given an initial state $\left|\psi_{i}\right\rangle$, apply $t$ iterations of $\mathcal{H}$ to obtain some state $\mathcal{H}^{t}\left|\psi_{i}\right\rangle$, apply an error $Z_{j}$ on some qubit with $0 \leq j<n$, and finally apply one last iteration of $\mathcal{H}$, giving the final state $\mathcal{H} Z_{j} \mathcal{H}^{t}\left|\psi_{i}\right\rangle$. Figure 12 shows the result of this experiment, showing the probability of measuring the $|1\rangle$ ancillary state plotted for various


Figure 12: 8 Qubit Demonstration of Rejection Probability as Function of Phase-Flip Errors. In this simulated experiment, the operator $\mathcal{H} Z_{j} \mathcal{H}^{t}$ is applied to an initial state $|0\rangle$. The probability of measuring the $|0\rangle$ ancillary state in the final application of $\mathcal{H}$, and thus the probability of rejecting a circuit execution in which an error occurred, is shown as $p(|1\rangle)$.
$j$ and $t$. As can be seen in this figure, the occurrence of any error strictly increases the probability of measuring the $|1\rangle$ ancillary state, and thus of discarding the obtained erroraffected result in post selection. We will now provide a general analysis of this type of error, and use the analysis to explain the plotted figure in greater detail. First, we define two new operators, $A_{m}=Z \otimes I \otimes \ldots \otimes I$, where there are $m-1$ identity gates and one $Z$ gate, and $B_{k}=I \otimes \ldots \otimes I$ where there are $k$ identity gates, and such that $k+m=n$. We may then write,

$$
Z_{j}=B_{j} \otimes A_{n-j}
$$

where we implicitly assume that $B_{0} \otimes A_{n}=A_{n}$, and that $m \geq 1$, as there must be at least one $Z$ gate in the $Z_{j}$ operator. First, we observe that $A_{m}=\operatorname{diag}(11 \ldots-1-1)$, where there are $2^{m-1}$ consecutive +1 elements from the start, followed by $2^{m-1}$ elements with value -1 . Furthermore, $B_{j}=\operatorname{diag}\left(\begin{array}{lll}1 & \ldots\end{array}\right)$ where there are $2^{j}$ elements with value +1 . By definition of the tensor product, $B_{j} \otimes A_{n-1}$ thus represents a diagonal sign pattern consisting of $2^{j}$ repetitions of the pattern given by $A_{m}$ - meaning that the period of the pattern is $2^{m}$. For example, in the extreme cases, $Z_{0}=A_{n}$ represents a $2^{n} \times 2^{n}$ diagonal matrix where the first half of the diagonal terms are all 1 and the second half of the diagonal terms are all -1 (thus having a period of $2^{n}$ ), and $Z_{n-1}$ represents a $2^{n} \times 2^{n}$ diagonal matrix with the pattern $\operatorname{diag}(1-1)$ repeated $2^{n-1}$ times. It is clear that the application of $Z_{j}$ to a state $\mathcal{H}^{t}\left|\psi_{i}\right\rangle$ will introduce a discontinuity to the wave function (beyond the negligible discontinuity already imposed by the discretization) at each pair of states $x, x+1$ for which $\operatorname{sign}\left(\langle x| Z_{j}|x\rangle\right) \neq \operatorname{sign}\left(\langle x+1| Z_{j}|x+1\rangle\right)$. In particular, there is one such sign change in each pattern given by the diagonal of $A_{n-j}$, and so there are a maximum of $2^{j}$ possible discontinuities introduced by $Z_{j}$. Moreover, after $t$ iterations according to $\mathcal{H}$, only states $|0\rangle$ through $|t-1\rangle$ will have non-zero amplitudes, and so the number of discontinuities introduced by $Z_{j}$ is given by $\left\lfloor\frac{t}{2^{n-j}}\right\rfloor$. This may be made more clear by considering an example. After the application of $Z_{j}$ a possible state associated with the $|0\rangle$ ancillary state may look like $|\psi\rangle=|x\rangle-2|x+1\rangle+|x+2\rangle$, and so $|\psi\rangle-|\psi+1\rangle=|x\rangle-3|x+1\rangle+3|x+2\rangle-|x+3\rangle$, as opposed to the error free state of $|x\rangle+|x+1\rangle-|x+2\rangle-|x-3\rangle$, clearly illustrating how the sign-flip error leaves greater amplitude on the $|1\rangle$ state because of the shift discontinuous negated distribution. Indeed,


Figure 13: 8 Qubit Demonstration of Rejection Probability as Function of Bit-Flip Errors. In this simulated experiment, the operator $\mathcal{H} X_{j} \mathcal{H}^{t}$ is applied to an initial state $|0\rangle$. The probability of measuring the $|0\rangle$ ancillary state in the final application of $\mathcal{H}$, and thus the probability of rejecting a circuit execution in which an error occurred, is shown as $p(|1\rangle)$.

SI G, shows that the cancellation property of the $|1\rangle$ ancilla state only holds if the wave function is continuous. This analysis well explains the data shown in Figure 12. Of note, as $j$ increases, as predicted by the increase in the number of discontinuities introduced, the probability of measuring the $|1\rangle$ ancillary state increases, up to a near- 1 probability of discarding the execution when a $Z_{7}$ error occurs (in an 8 -qubit system, numbered $[0,7]$ ). Furthermore, when the period of the error operator's sign pattern is greater than $t$ (in particular, when it is greater than the number of states with amplitude $\gg 0$ ) the action of $Z_{j}$ is that of identity and so no error is incurred, hence why each of the $Z_{j}$ curves follows the error-free curve for some number of $t$. For example, this explains why the $Z_{0}$ operator follows the error-free curve until $t \approx 220$ (observing that $2^{8}=256$ ).

Moreover, by applying the normal approximation for the distribution produced by $t$ iterations of $\mathcal{H}$, it would be straight-forward to continue this analysis and obtain a closedform expression for the probability of measuring the $|1\rangle$ state after experiencing a $Z_{j}$ error.

Analysis of Bit-Flip Errors The analysis for conceptual bit-flip errors follows a similar argument as that presented for the conceptual phase-flip errors, only with a couple small differences. Primarily, the discontinuities introduced by inserting an $X_{j}$ operator come from swapping the amplitudes of two states $\left|x_{0} \ldots x_{j} \ldots x_{n-1}\right\rangle$ and $\left|x_{0} \ldots \bar{x}_{j} \ldots x_{n-1}\right\rangle$. As such, as $j$ increases $X$ is applied to less-and-less significant qubits, meaning that the distance between swapped states decreases exponentially. As a result, we would expect the error introduced by the $X_{j}$ operator to increase as $j$ decreases, and thus for the probability of measuring the $|1\rangle$ ancillary state to increase as $j$ decreases. In the extreme, for large $n$, the distribution given by $\mathcal{H}^{t}\left|\psi_{i}\right\rangle$ is essentially continuous, and so swapping the amplitude on states $|x\rangle$ and $|x+\delta\rangle$ is essentially equivalent to an identity transformation, hence why the $X_{7}$ operator is identical to the error-free line in Figure 13. The fact that $X_{j}$ introduces more subtle discontinuities than $Z_{j}$ is also reflected in Figure 13, noting that the probability of measuring the $|1\rangle$ ancillary state (and thus of discarding a result in which an error occurred) is, in general, less than the corresponding probability for the $Z_{j}$ operator. We expect that an expression for the net cancellation incurred for a given $X_{j}$ and $t$ may be obtained by following the preceding analysis.

## F Asymptotic analysis on the expected number of trials

In this section we analytically derive the expected number of trials to successfully load the desired distribution into a quantum register.

## F. 1 Without qubit scaling

First we look at the case where we do not perform qubit scaling in the state generation process. Suppose that we have a $n$-qubit quantum register originally in the computational basis state $\left|x_{0}\right\rangle$, where $x_{0}$ is an integer in the range of $\left[0,2^{n}-1\right]$. Assume that $x_{0}+t<2^{n}$, after $t-1$ iterations of $\mathcal{H}$ on the quantum register, the state becomes

$$
\mathcal{H}^{t-1}\left|x_{0}\right\rangle=\frac{1}{\sqrt{c_{t-1}}} \sum_{k=0}^{t-1}\binom{t-1}{k}\left|x_{0}+k\right\rangle,
$$

where

$$
c_{t-1}=\sum_{k=0}^{t-1}\binom{t-1}{k}^{2}
$$

is the normalization factor. On the next iteration, the state we have before measuring the ancilla qubit is

$$
\begin{array}{r}
\frac{1}{2 c_{t-1}} \sum_{k=0}^{t-1}\left[\binom{t-1}{k}\left|x_{0}+k\right\rangle+\binom{t}{k}\left|x_{0}+k+1\right\rangle\right]|0\rangle \\
+\frac{1}{2 c_{t-1}} \sum_{k=0}^{t-1}\left[\binom{t-1}{k}\left|x_{0}+k\right\rangle-\binom{t}{k}\left|x_{0}+k+1\right\rangle\right]|1\rangle \\
=\frac{1}{2 c_{t-1}} \sum_{k=0}^{t}\binom{t}{k}\left|x_{0}+k\right\rangle|0\rangle+\frac{1}{2 c_{t-1}} \sum_{k=0}^{t-1}\left[\binom{t-1}{k}\left|x_{0}+k\right\rangle-\binom{t}{k}\left|x_{0}+k+1\right\rangle\right]|1\rangle
\end{array}
$$

The probability of measuring $|0\rangle$ in the ancilla qubit is

$$
P_{t}(|0\rangle)=\frac{1}{4 c_{t-1}^{2}} \sum_{k=0}^{t}\binom{t}{k}^{2}=\frac{c_{t}^{2}}{4 c_{t-1}^{2}}
$$

Using the Chu-Vandermonde identity, we have

$$
c_{t}^{2}=\sum_{k=0}^{t}\binom{t}{k}\binom{t}{k}=\binom{2 t}{t} .
$$

Therefore,

$$
\begin{equation*}
P_{t}(|0\rangle)=\frac{\binom{2 t}{t}}{4\binom{2 t-2}{t-1}}=1-\frac{1}{2 t} . \tag{10}
\end{equation*}
$$

Consequently, the probability of getting a successful state with $t$ applications of $\mathcal{H}$ is

$$
\begin{equation*}
p_{s}(t)=\prod_{k=1}^{t}\left(1-\frac{1}{2 k}\right) \tag{11}
\end{equation*}
$$

and the expected number of trials to get one successful run with $t$ iterations is given by the inverse of $p_{s}(t)$, i.e.

$$
\mathrm{E}_{t}[T]=\frac{1}{p_{s}(t)}=\frac{1}{\prod_{k=1}^{t}\left(1-\frac{1}{2 k}\right)} .
$$

We take the logarithm of $\mathrm{E}_{t}[T]$ and expand the additive terms into Taylor series at $\infty$,

$$
\begin{aligned}
\log \left(\mathrm{E}_{t}[T]\right) & =-\sum_{k=1}^{t} \log \left(1-\frac{1}{2 k}\right) \\
& =\sum_{k=1}^{t}\left(\frac{1}{2 k}+O\left(\frac{1}{k^{2}}, k \rightarrow \infty\right)\right) \\
& =\frac{1}{2} \sum_{k=1}^{t} \frac{1}{k}+O(1, t \rightarrow \infty)
\end{aligned}
$$

The last equality holds because the series of the inverse squares and higher order powers converge absolutely. Then with the asymptotic development of harmonic series, we have

$$
\log \left(\mathrm{E}_{t}[T]\right)=\frac{1}{2} \log (t)+O(1, t \rightarrow \infty)
$$

Therefore, we have

$$
\begin{equation*}
\mathrm{E}_{t}[T]=C \sqrt{t}(1+O(1, t \rightarrow \infty)) \tag{12}
\end{equation*}
$$

We will now consider an upper bound, as for any $x>-1, \log (1+x) \geq \frac{x}{1+x}$. In our case, $\log \left(1-\frac{1}{2 k}\right) \geq \frac{-1}{2 k-1}$ Thus,

$$
\begin{aligned}
\log \left(\mathrm{E}_{t}[T]\right) & =-\sum_{k=1}^{t} \log \left(1-\frac{1}{2 k}\right) \\
& \leq \sum_{k=1}^{t} \frac{1}{2 k-1} \\
& =1+\sum_{k=1}^{t-1} \frac{1}{2 k+1} \\
& \leq 1+\sum_{k=1}^{t-1} \frac{1}{2 k}
\end{aligned}
$$

Using the upper bound of the harmonic series,

$$
\log \left(\mathrm{E}_{t}[T]\right) \leq \frac{1}{2} \log (t-1)+\frac{3}{2}
$$

Therefore,

$$
\begin{equation*}
\mathrm{E}_{t}[T] \leq e^{\frac{3}{2}} \sqrt{t-1} \tag{13}
\end{equation*}
$$

Therefore the number of expected trials to get a successful state after $t$ applications of $\mathcal{H}$ grows sub-linearly with $t$.

## F. 2 With qubit scaling

First we show that the addition of a qubit in the $|+\rangle$ state on the least significant end decreases the probability of measuring the ancilla in the $|1\rangle$ state in the subsequent application of $\mathcal{H}$ by $1 / 2$, thereby increasing the algorithm's overall probability of success. Let $a_{n, t}(x)$ be the amplitude of the $n$ qubit basis state $|x\rangle_{n}$ after $t$ applications of $\mathcal{H}$, i.e. the state in the $n$-qubit quantum register is

$$
\begin{equation*}
\left|\psi_{t}\right\rangle_{n}=\sum_{x=0}^{t} a_{n, t}(x)|x\rangle_{n} \tag{14}
\end{equation*}
$$

Now we perform another application of $\mathcal{H}$ less the measurement, and we arrive at the state

$$
\frac{1}{2} \sum_{x=0}^{t} a_{n, t}(x)\left(|x\rangle_{n}+|x+1\rangle_{n}\right)|0\rangle+\frac{1}{2} \sum_{x=0}^{t} a_{n, t}(x)\left(|x\rangle_{n}-|x+1\rangle_{n}\right)|1\rangle
$$

therefore the probability of measuring $|1\rangle$ in the ancilla qubit is

$$
\begin{equation*}
P_{t+1}(|1\rangle)=\frac{1}{4}\left\{a_{n, t}^{2}(0)+a_{n, t}^{2}(t)+\sum_{x=0}^{t-1}\left[a_{n, t}(x)-a_{n, t}(x+1)\right]^{2}\right\} \tag{15}
\end{equation*}
$$

On the other hand, adding a qubit in the $|+\rangle$ state to Equation 14, we have

$$
\left|\phi_{t}\right\rangle_{n+1}=\frac{1}{\sqrt{2}} \sum_{x=0}^{t} a_{n, t}(x)\left(|2 x\rangle_{n+1}+|2 x+1\rangle_{n+1}\right)
$$

Now we perform another application of $\mathcal{H}$ less the measurement on $\left|\phi_{t}\right\rangle_{n+1}$, and we arrive at the state

$$
\begin{aligned}
& \frac{1}{2 \sqrt{2}} \sum_{x=0}^{t} a_{n, t}(x)\left(|2 x\rangle_{n+1}+2|2 x+1\rangle_{n+1}+|2 x+1\rangle_{n+2}\right)|0\rangle \\
& +\frac{1}{2 \sqrt{2}} \sum_{x=0}^{t} a_{n, t}(x)\left(|2 x\rangle_{n+1}-|2 x+2\rangle_{n+1}\right)|1\rangle
\end{aligned}
$$

therefore the probability of measuring $|1\rangle$ in the ancilla qubit is

$$
\begin{equation*}
P_{t+1}^{\prime}(|1\rangle)=\frac{1}{8}\left\{a_{n, t}^{2}(0)+a_{n, t}^{2}(t)+\sum_{x=0}^{t-1}\left[a_{n, t}(x)-a_{n, t}(x+1)\right]^{2}\right\} \tag{16}
\end{equation*}
$$

Comparing Equation 16 with Equation 15, we have

$$
\begin{equation*}
P_{t+1}^{\prime}(|1\rangle)=\frac{1}{2} P_{t+1}(|1\rangle) \tag{17}
\end{equation*}
$$

Therefore adding a qubit before applying $\mathcal{H}$ in the $|+\rangle$ state on the least significant end decreases the probability of measuring the ancilla in the $|1\rangle$ state in the subsequent application of $\mathcal{H}$ by $1 / 2$.

We now derive the expected number of trials in the context of qubit scaling. Suppose that we start with $n_{1}$ qubits, and gradually increase the number of qubits to $n_{m}$ in $m$ stages. The qubit count increases by 1 in each subsequent stage, and the number of applications of $\mathcal{H}$ is given by $t_{1}, t_{2}, \ldots, t_{m}$, respectively. From the analysis in Equation D we know that $t_{1}$ is determined by the scaled variance $\hat{\sigma}$ of the target distribution, and $t_{2}, \ldots, t_{m}$ are bounded by some constant number that is independent of the number of qubits in the final state. In other words, let $t_{\max }=\max t_{r}, r=2, \ldots, m$, then $t_{\max }$ is independent of $n_{m}$. With this setting, the probability of getting one successful state from the process described above would be

$$
p_{s}=\prod_{r=1}^{m} p_{s, r}\left(t_{r}\right)
$$

where $p_{s, r}\left(t_{r}\right)$ is the probability of success during stage $r$. From Equation 11 we have

$$
\begin{equation*}
p_{s, 1}\left(t_{1}\right)=\prod_{k=1}^{t_{1}}\left(1-\frac{1}{2 k}\right) . \tag{18}
\end{equation*}
$$

On the other hand, Equation 17 indicates that

$$
\begin{align*}
p_{s, r}\left(t_{r}\right) & >P_{r, 1}(|0\rangle)^{t_{r}} \\
& =\left[1-\frac{1}{2} P_{r-1, t_{r-1}+1}(|1\rangle)\right]^{t_{r}} \\
& \geq\left[1-\frac{1}{2^{r-1}} P_{1, t_{1}+1}(|1\rangle)\right]^{t_{r}} \\
& =\left[1-\frac{1}{2^{r-1}} \frac{1}{2 t_{1}}\right]^{t_{r}} \\
& \geq\left[1-\frac{1}{t_{1}} \frac{1}{2^{r}}\right]^{t_{\max }} \tag{19}
\end{align*}
$$

where $P_{r, k}(|0\rangle)$ and $P_{r, k}(|1\rangle)$ are the probabilities of measuring $|0\rangle$ and $|1\rangle$ in the ancilla qubit after the $k$-th application of $\mathcal{H}$ in stage $r$. Substituting Equation 19 into Equation F.2, we have

$$
\begin{aligned}
p_{s} & >p_{s, 1}\left(t_{1}\right) \prod_{r=2}^{m}\left[1-\frac{1}{t_{1}} \frac{1}{2^{r}}\right]^{t_{\max }} \\
& =p_{s, 1}\left(t_{1}\right)\left(1-\frac{1}{2 t_{1}}\right)^{-t_{\max }} \prod_{r=1}^{m}\left[1-\frac{1}{t_{1}} \frac{1}{2^{r}}\right]^{t_{\max }} \\
& =p_{s, 1}\left(t_{1}\right)\left[\frac{\left(\frac{1}{t_{1}} ; \frac{1}{2}\right)_{m}}{1-\frac{1}{2 t_{1}}}\right]^{t_{\max }} \\
& >p_{s, 1}\left(t_{1}\right)\left[\frac{\left(\frac{1}{t_{1}} ; \frac{1}{2}\right)_{\infty}}{1-\frac{1}{2 t_{1}}}\right]^{t_{\max }}
\end{aligned}
$$

where $(a ; q)_{n}=\prod_{k=1}^{n}\left(1-a q^{k}\right)$ is the q-Pochhammer symbol. And the expected number of trials to get one successful state is

$$
\begin{equation*}
\mathrm{E}[T]=\frac{1}{p_{s}}<\frac{1}{p_{s, 1}\left(t_{1}\right)}\left[\frac{1-\frac{1}{2 t_{1}}}{\left(\frac{1}{t_{1}} ; \frac{1}{2}\right)_{\infty}}\right]^{t_{\mathrm{max}}} . \tag{20}
\end{equation*}
$$

Without loss of generality, we may assume that $t_{1} \geq 2$, therefore

$$
p_{s}>p_{s, 1}\left(t_{1}\right)\left[\frac{\left(\frac{1}{2} ; \frac{1}{2}\right)_{\infty}}{1-\frac{1}{4}}\right]^{t_{\max }} \approx 0.385^{t_{\max }} p_{s, 1}\left(t_{1}\right)
$$

which is independent of $n_{m}$. Consequently, the expected number of trials to get one successful state would be

$$
\mathrm{E}[T]=\frac{1}{p_{s}}<\frac{1}{p_{s, 1}\left(t_{1}\right)}\left[\frac{3}{4\left(\frac{1}{2} ; \frac{1}{2}\right)_{\infty}}\right]^{t_{\max }} \approx \frac{2.597^{t_{\max }}}{p_{s, 1}\left(t_{1}\right)}
$$

Therefore we have proved that the expected number of trials to get one successful state is bounded by a number that is independent of the number of qubits in the final state.

Now we derive a lower bound for $\left(\frac{1}{t_{1}} ; \frac{1}{2}\right)_{\infty}$ in terms of $t_{1}$. To do that, we take the logarithm of the q-Pochhammer symbol, and expand each log term into Taylor series

$$
\begin{aligned}
\log \left(\frac{1}{t_{1}} ; \frac{1}{2}\right)_{\infty} & =\sum_{r=1}^{\infty} \log \left(1-\frac{1}{t_{1}} \frac{1}{2^{r}}\right) \\
& =-\sum_{r=1}^{\infty}\left[\sum_{k=1}^{\infty} \frac{1}{k}\left(\frac{1}{t_{1} 2^{r}}\right)^{k}\right] \\
& =-\sum_{k=1}^{\infty} \frac{1}{k t_{1}^{k}}\left[\sum_{r=1}^{\infty}\left(\frac{1}{2^{k}}\right)^{r}\right] \\
& =-\sum_{k=1}^{\infty} \frac{1}{k t_{1}^{k}} \frac{1}{2^{k}-1} \\
& =-\frac{1}{t_{1}}-\sum_{k=2}^{\infty} \frac{1}{k t_{1}^{k}} \frac{1}{2^{k}-1} \\
& >-\frac{1}{t_{1}}-\sum_{k=2}^{\infty} \frac{1}{k t_{1}^{k}} \frac{1}{2^{k}-2} \\
& =-\frac{1}{t_{1}}-\frac{1}{2 t_{1}} \sum_{k=1}^{\infty} \frac{1}{k t_{1}^{k}} \frac{1}{2^{k}-1} \\
& =-\frac{1}{t_{1}}-\frac{1}{2 t_{1}} \log \left(\frac{1}{t_{1}} ; \frac{1}{2}\right)_{\infty} .
\end{aligned}
$$

It immediately follows that

$$
\log \left(\frac{1}{t_{1}} ; \frac{1}{2}\right)_{\infty}>-\frac{2}{2 t_{1}-1}
$$

or equivalently,

$$
\begin{equation*}
\left(\frac{1}{t_{1}} ; \frac{1}{2}\right)_{\infty}>e^{-\frac{2}{2 t_{1}-1}} . \tag{21}
\end{equation*}
$$

Substituting Equation 21 back into Equation 20, we have

$$
\begin{equation*}
\mathrm{E}[T]<\frac{1}{p_{s, 1}\left(t_{1}\right)}\left[\left(1-\frac{1}{2 t_{1}}\right) e^{\frac{2}{2 t_{1}-1}}\right]^{t_{\max }} \tag{22}
\end{equation*}
$$

We now observe that $1-\frac{1}{2 t_{1}}<1$, and thus,

$$
\begin{equation*}
\mathrm{E}[T]<\frac{1}{p_{s, 1}\left(t_{1}\right)}\left[e^{\frac{2}{2 t_{1}-1}}\right]^{t_{\max }} \tag{23}
\end{equation*}
$$

Moreover, since we assumed that $t_{1} \geq 2$, and so $e^{\frac{2}{2 t_{1}-1}}<e^{\frac{2}{3}}<2$ (i.e. $e^{\frac{2}{2 t_{1}-1}}$ is maximized when $t_{1}$ is minimized), we get

$$
\begin{equation*}
\mathrm{E}[T]<\frac{2^{t_{\max }}}{p_{s, 1}\left(t_{1}\right)} \tag{24}
\end{equation*}
$$

From Equation 18, we observe that $p_{s, 1}\left(t_{1}\right)>\left(\frac{1}{2}\right)^{t_{1}}$, leading to the bound

$$
\begin{equation*}
\mathrm{E}[T]<2^{t_{\max }+t_{1}} \tag{25}
\end{equation*}
$$

Noting that both $t_{1}$ and $t_{\text {max }}$ are constant values independent of $n_{m}$, we have thus obtained a constant bound on the expected number of trials required for the procedure to succeed once.

## G General Proof of |1> Ancillary State Cancellation

Consider some probability amplitude function function $f: \mathbb{R} \rightarrow \mathbb{R}$, such that we obtain a superposition

$$
(f(x)|x\rangle+f(x+\delta)|x+\delta\rangle)|0\rangle+(f(x)|x\rangle-f(x+\delta)|x+\delta\rangle)|1\rangle,
$$

where $\delta$ is some small real value. The ratio of the amplitude on the $|1\rangle$ ancillary state to the total amplitude on both ancillary states may then be written as,

$$
\begin{equation*}
\lim _{\delta \rightarrow 0} \frac{f(x)-f(x+\delta)}{f(x)+f(x+\delta)+f(x)-f(x+\delta)}=\lim _{\delta \rightarrow 0} \frac{f(x)-f(x+\delta)}{2 f(x)}=\lim _{\delta \rightarrow 0}\left[\frac{1}{2}-\frac{1}{2} \frac{f(x+\delta)}{f(x)}\right] . \tag{26}
\end{equation*}
$$

Suppose that $f$ is a continuous function, meaning that $\lim _{\delta \rightarrow 0} f(x+\delta)=f(x)$. Then, Equation 26 reduces to 0 and consequently the probability of measuring the $|1\rangle$ state is 0 for a small shift $\delta$ when $f$ is a continuous function. Now suppose that $f$ is not continuous at point $x$, thus $\lim _{\delta \rightarrow 0} f(x+\delta) \neq f(x)$, then Equation 26 does not necessarily simplify to zero. Indeed, it is straightforward to construct an example in which $f(x+\delta) / f(x)$ can assume an arbitrarily large or small value (by defining $f$ to be piece-wise with the left and right hand limits at $x$ being different), and thus the probability of measuring the $|1\rangle$ could be an arbitrary zero or non-zero constant.

## H Simulation Based Approaches: Imaginary Time Evolution

Another variant of the algorithm which was explored simulates $e^{\mathcal{H} t}$ instead of $\mathcal{H}^{t}$, where $\mathcal{H}|x\rangle=|x-1\rangle+|x+1\rangle$. A number of interesting observations were made regarding simulating $e^{\mathcal{H} t}$ as opposed to $\mathcal{H}^{t}$, with some distinct advantages and disadvantages noted.

## H. 1 The Need for Imaginary Time Evolution in Simulation-Based Approaches

Originally, we considered simulating the dynamics of $\mathcal{H}$ in accordance with the Schrodinger equation,

$$
|\psi(t)\rangle=e^{-i \mathcal{H} t}|\psi(0)\rangle
$$

Of course, if such a Hamiltonian $\mathcal{H}$ could be efficiently decomposed into some basis set (such as the set of $n$-bit Pauli strings) such that $\mathcal{H}=\sum_{k} t_{k} \mathcal{H}_{k}$ where each $\mathcal{H}_{k}$ is an element of the basis set with associated weight $t_{k}$, then the time evolution under $\mathcal{H}$ could be approximated by the Trotter-Suzuki decomposition [36],

$$
e^{-i \mathcal{H} t} \approx\left(\prod_{k=1}^{k} e^{-i \tau_{k} \mathcal{H}_{k}}\right)^{P}
$$

where $\tau_{k}=\frac{t \cdot t_{k}}{P}$. However, in taking the Taylor expansion of $e^{-i \mathcal{H} t}$, we obtain,

$$
e^{-i \mathcal{H} t}=\sum_{k=0}^{\infty} \frac{(-i)^{k} t^{k} \mathcal{H}^{k}}{k!}=\sum_{k=0,4, \ldots} \frac{t^{k} \mathcal{H}^{k}}{k!}-i \sum_{k=1,5, \ldots} \frac{t^{k} \mathcal{H}^{k}}{k!}-\sum_{k=2,6, \ldots} \frac{t^{k} \mathcal{H}^{k}}{k!}+i \sum_{k=3,7, \ldots .} \frac{t^{k} \mathcal{H}^{k}}{k!},
$$

which corresponds to four sets of weighted sums of Binomial distributions, each of which is separated by differing phases. As a result, the states experience an undesirable interference
pattern which causes a distribution that is not normal to be produced. As such, we would like to simulate the object $e^{\mathcal{H} t}$, as it is intuitively similar to simulating $e^{-i \mathcal{H} t}$ and avoids the phase complications. $e^{\mathcal{H} t}$ expands as

$$
\begin{equation*}
e^{\mathcal{H} t}=\sum_{k=0}^{\infty} \frac{t^{k}}{k!} \mathcal{H}^{k}, \tag{27}
\end{equation*}
$$

and as such corresponds to an infinite sum of Binomial distributions (noting that each $\mathcal{H}^{k}$ produces a particular Binomial distribution with weight $\frac{t^{k}}{k!}$ ) which is itself an approximation to a normal distribution in accordance with the central limit theorem. In understanding how $e^{\mathcal{H} t}$ appears to produce the same normal distribution as $\mathcal{H}^{t}$, understanding the value of $k$ in terms of $t$ for which the weights $\frac{t^{k}}{k!}$ are maximum may be beneficial. Whilst a more rigorous analytical argument could likely be made by substituting the factorial with an equivalent-valued Gamma function, and then computing $\frac{\partial}{\partial k} \frac{t^{k}}{k!} \approx \frac{\partial}{\partial k} \frac{t^{k}}{\Gamma(k+1)}$, numerical simulation suggests that the coefficient is greatest when $k \approx t$. As a result, the term in Equation 27 for which $k=t, \mathcal{H}^{t}$, is also likely the term which contributes the most to the transformation produced by $e^{\mathcal{H} t}$, providing some explanation as to why $\mathcal{H}^{t}$ appears to produce the same distributions as $e^{\mathcal{H} t}$. Of course, if $\mathcal{H}=\mathcal{H}^{\dagger}$, then $e^{\mathcal{H} t}$ will not be unitary for $t \in \mathbb{R}$. Thus, we have motivated the need for imaginary time evolution in a simulation based approach.

## H. 2 Imaginary Time Evolution

In 2017, Li and Benjamin presented a quantum-classical hybrid algorithm for the purpose of simulating the Shrodinger equation [37]. This approach assumes that the state of the wave function at time $t,|\phi(t)\rangle$, may be approximated by a parameterized trial wave function $|\psi(t)\rangle \equiv|\psi(\vec{\theta}(t))\rangle$, where $\vec{\theta}(t)$ are a set of variational time-dependant parameters. In 2019, McArdle et al. expanded upon the aforementioned work to describe how this simulation procedure may be applied to the variational simulation of non-unitary objects such as $e^{\mathcal{H t}}$. In summary, the non-unitary trajectory of a given starting state is projected onto a unitary evolution through the parameters of the variational quantum circuit implementing the trial wave function.

As a result, a simulation based approach utilizing ITE for the presented distribution generating algorithm has a number of advantages and one significant challenge. The advantages are (1) the parameterized circuit captures degrees of freedom enabling shallower NISQ-friendly circuits to be obtained, (2) allows for continuous values of $t$ (as opposed to the main approach presented which requires integer values for $t$ ), (3) the alternating even and odd state cancellation which motivated the redefinition of the Hamiltonian as $\mathcal{H}|x\rangle=|x\rangle+|x+1\rangle$ is circumvented, allowing the original definition of $\mathcal{H}|x\rangle=|x-1\rangle+|x+1\rangle$ to be directly used, and (4) the definition of the Hamiltonian could be modified to easily enable the creation of a potentially broad class of distributions. The main challenge in using the ITE approach is that it requires the creation of a variational circuit capturing the necessary degrees of freedom in its parameters. Moreover, it is necessary to do so without introducing an exponential number of parameters (and thus number of gates) so as to remain competitive with the primary Galton simulation approaches described in this work. As such, a variational circuit should be constructed which effectively captures knowledge of the problem at hand, but how to do so in this particular instance remains a problem for future investigation.


[^0]:    ${ }^{1}$ Note that this principle also applies if the transitions maps $j$ to $j-1$ or $j+1$. However, we choose the mapping $j \rightarrow j$ and $j \rightarrow j+1$ as it avoids a rectifiable problem where alternating states have zero amplitude, and moreover avoids the use of an additional ancillary qubit.

[^1]:    ${ }^{2}$ I.e. as qubit $n_{i+1}$, where $n_{i}$ refers to the number of qubits at the start of the $i^{t h}$ step of the procedure. Thus, $n_{i+1}=n_{i}+1$.
    ${ }^{3}$ Similarly, $t_{i}$ refers to the number of iterations of $\mathcal{H}$ applied on the $i^{\text {th }}$ step of the procedure.

