A Threshold for Quantum Advantage in Derivative Pricing

Shouvanik Chakrabarti^{1,2}, Rajiv Krishnakumar¹, Guglielmo Mazzola³,
Nikitas Stamatopoulos¹, Stefan Woerner³, and William J. Zeng¹

¹Goldman, Sachs & Co., New York, NY

²University of Maryland, College Park, MD

³IBM Quantum, IBM Research – Zurich

We give an upper bound on the resources required for valuable quantum advantage in pricing derivatives. To do so, we give the first complete resource estimates for useful quantum derivative pricing, using autocallable and Target Accrual Redemption Forward (TARF) derivatives as benchmark use cases. We uncover blocking challenges in known approaches and introduce a new method for quantum derivative pricing - the *re-parameterization method* - that avoids them. This method combines pre-trained variational circuits with fault-tolerant quantum computing to dramatically reduce resource requirements. We find that the benchmark use cases we examine require 8k logical qubits and a T-depth of 54 million. We estimate that quantum advantage would require executing this program at the order of a second. While the resource requirements given here are out of reach of current systems, we hope they will provide a roadmap for further improvements in algorithms, implementations, and planned hardware architectures.

1 Introduction

A derivative contract is a financial asset whose value is based on (or *derived* from) the price of one or more underlying assets. Examples of these underlying assets include stocks, currencies, commodities, etc. A derivative contract is typically issued between an issuer and a holder, and is valid until its *expiration date*. Each derivative defines a payoff that quantifies what the holder stands to gain. Generically, payoffs depend on the value of the underlying asset(s) across the duration of the contract. Derivative contracts are ubiquitous in finance with various uses from hedging risk to speculation, and currently have an estimated global gross market value in the tens of trillions of dollars [1]. A more detailed introduction to derivatives with descriptions of some of the common derivatives used by financial institutions is given in Appendix A.

The goal of derivative *pricing* is to determine the value of entering a derivative contract today, given uncertainty about future values of the underlying assets and consequently the ultimate payoff. In many cases, the pricing of derivative contracts uses Monte Carlo methods which consume significant computational resources for financial institutions. Therefore, finding a quantum advantage for this application would be very valuable to the financial sector as a whole. This work gives the first detailed resource estimates of the required conditions for quantum advantage in derivative pricing. In doing so, it also introduces new methods for loading stochastic processes into a quantum computer.

The rest of the paper is structured as follows: **Section 2** gives a brief background on classical Monte Carlo methods for pricing derivatives and summarizes our results: the resources required for the quantum algorithms for pricing some commonly traded derivatives. Then, in **Section 3**, we present our core approach to the quantum algorithm and discuss its error analysis. **Section 4** covers subroutines to load a stochastic model of the underlying assets into a quantum state along with the resources required for these methods. Here, we introduce the *re-parametrization method*, a novel method that plays a crucial role in the first feasible end-to-end path to quantum advantage

in derivative pricing. In **Section 5** we discuss the subroutine that applies the payoff operator to the stored stochastic process. Finally, in **Section 6**, we end with a discussion on the implications of this work and potential future paths.

2 Derivative Pricing and Summarized Results

The price of the underlying asset(s) of a derivative contract is typically modeled as a stochastic process under assumptions like no-arbitrage.¹ A common model, that we will use in this work, is that the underlying assets evolve under geometric Brownian motion [2]. Let $\vec{S}^t \in \mathbb{R}_+^d$ be a vector of values for d underlyings at time t . Let $(\vec{S}^0, \dots, \vec{S}^T) = \bar{\omega} \in \bar{\Omega}$ be a path of a discrete time multivariate stochastic process describing the values of those assets. We use both notations for a path in the text. The corresponding probability density function is denoted by $\bar{p}(\bar{\omega})$. Let $f(\bar{\omega}) = f(\vec{S}^0, \dots, \vec{S}^T) \in \mathbb{R}$ be the discounted payoff of some derivative on those assets. To price the derivative we calculate

$$\mathbb{E}(f) = \int_{\bar{\omega} \in \bar{\Omega}} \bar{p}(\bar{\omega}) f(\bar{\omega}) d\bar{\omega}. \quad (1)$$

The reason for having a discounted payoff is to take into account the opportunity cost of investing in a risk-free asset. For the rest of the paper all payoffs will be implicitly regarded as discounted, except in **Section 5** and Appendix A where we will be explicit about whether payoffs are discounted or not. More details on the concept of discounted payoffs can be found in Appendix A.3.

If the underlying stochastic processes are modeled with geometric Brownian motion then they have transition probabilities

$$P(\vec{S}^t | \vec{S}^{t-1}) = \frac{\exp\left(-\frac{1}{2}(\ln \vec{S}^t - \vec{\mu}^{t-1})^\top \Sigma^{-1}(\ln \vec{S}^t - \vec{\mu}^{t-1})\right)}{(2\pi)^{d/2}(\det \Sigma)^{1/2} \prod_{j=1}^d S_j^t}. \quad (2)$$

where

$$\begin{aligned} \ln \vec{S}^t &= (\ln S_1^t, \ln S_2^t, \dots, \ln S_d^t)^\top \\ \vec{\mu}^{t-1} &= (\mu_1^{t-1}, \mu_2^{t-1}, \dots, \mu_d^{t-1})^\top \\ \mu_j^{t-1} &= (r - 0.5\sigma_j^2) \Delta t + \ln S_j^{t-1}. \end{aligned} \quad (3)$$

Note that Eq. (2) at time t has a dependency on the asset vector at time $t-1$ via $\ln S_j^{t-1}$ in μ_j^{t-1} . The parameters r and σ_j are the risk-free rate² and the volatility of the j -th asset respectively, Δt is the time duration between steps of the stochastic process, and Σ is the $d \times d$ positive-definite covariance matrix of the d underlyings

$$\Sigma = \Delta t \begin{bmatrix} \sigma_1^2 & \rho_{12}\sigma_1\sigma_2 & \dots & \rho_{1d}\sigma_1\sigma_d \\ \rho_{21}\sigma_2\sigma_1 & \sigma_2^2 & \dots & \rho_{2d}\sigma_2\sigma_d \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{d1}\sigma_d\sigma_1 & \dots & \dots & \sigma_d^2 \end{bmatrix} \quad (4)$$

where $-1 \leq \rho_{ij} \leq 1$ is the correlation between assets i and j . The probability of any particular path $\bar{\omega} \in \bar{\Omega}$ is then

$$\bar{p}(\bar{\omega}) = \prod_{t=1}^T P(\vec{S}^t | \vec{S}^{t-1}). \quad (5)$$

¹No-arbitrage is the assumption that no specific asset is priced differently in different marketplaces such that one can never buy an asset from one marketplace and immediately sell it at another for a profit.

² The *risk-free* rate is the rate of return of investing in a risk-free asset. Although such an asset is purely theoretical, we typically use treasury bonds to represent such an asset and approximate r as the yield of the treasury bond minus the current inflation rate.

Classically, some simple derivatives under this model are easy to price, such as European call options that can be priced analytically using the Black-Scholes equation [2]. Easy to price derivatives are often *path independent*, i.e. where the payoff is only a function of the final prices at exercise time $f(\bar{\omega}) = f(\bar{S}^T)$. This contrasts with *path dependent* derivatives that are more difficult to price. Path dependent derivatives are often priced in practice with classical Monte Carlo methods. More examples of standard and more complex derivatives are given in Appendix A.

When using classical Monte Carlo, the accuracy of derivative pricing converges as $O(1/\sqrt{M})$, where M is the number of paths that are sampled. In general cases, Montanaro [3] showed that quantum algorithms based on amplitude estimation [4] can be used to improve this to $O(1/M)$. Recent work has considered how to specialize this advantage to options pricing (options are a subcategory of derivatives) [5–7] and risk analysis [8, 9].

As this is only a quadratic speedup, it is important to focus on derivatives that are complicated enough to require a large M in practice. In this work we give end to end quantum resource estimates for two examples of such derivatives (autocallable options and TARFs) that are both computationally expensive, path-dependent derivatives. In doing so, we detail and optimize the loading into quantum states of the underlying distribution over asset paths. This loading step was left open in previous work [5, 6], and we give the first account of the resources required for it. Although autocallables and TARFs are usually not well known derivatives to those outside the financial sector, they are very commonly traded derivatives, particularly among financial institutions. It is for this reason, in addition to their complexity, that we have chosen them as suitable examples for the end-to-end quantum resource estimates. A more detailed description of autocallables and TARFs can be seen in Appendix A.4 and Appendix A.5 respectively.

In addition to estimating the resources needed for path loading using known methods (an extension of [7] that we call the Riemann Sum method), we introduce several optimizations, including intentional shifts from price space to return space calculations and the new re-parameterization method. These methods reduce the required resources significantly and are summarized in Table 1. In this table, we quantify resource requirements in the fault tolerant setting where the number of T -gates, called the T -count dominates the computational requirements. The T -depth is the sequential depth that dominates the runtime. *Logical qubits* consist of physical qubits in a quantum error correcting code of sufficient distance to support the required number of operations.

Method	(d, T)		Error		T-count		T-depth		# Logical Qubits	
	Auto	TARF	Auto	TARF	Auto	TARF	Auto	TARF	Auto	TARF
Riemann Sum					$\geq 10^{43}$	$\geq 10^{18}$	$\geq 10^{43}$	$\geq 10^{18}$	-	-
Riemann Sum (no-norm)	(3, 20)	(1, 26)	2×10^{-3}		1.6×10^{11}	5.5×10^{10}	1.5×10^8	1.6×10^8	23k	17k
Re-parameterization					1.2×10^{10}	9.8×10^9	5.4×10^7	8.2×10^7	8k	11.5k

Table 1: Resources estimated in this work for pricing derivatives using different methods for a target error of 2×10^{-3} . As representative use cases of business interest with non-trivial complexity, we consider a basket autocallable (Auto) with 3 underlyings, 5 payment dates and a knock-in put option with 20 barrier dates, and a TARF with one underlying and 26 payment dates. Detailed definitions of these contracts and their parameters can be found in Appendix A.4. We find that Grover-Rudolph methods [10] are not applicable in practice (details in Appendix B) and that Riemann summation methods require normalization assumptions to avoid errors that grow exponentially in T . Even if those normalization issues were avoided, as detailed in the Riemann Sum (no-norm) row, the re-parameterization method still performs best. See Section 4.1 for a discussion of the Riemann summation normalization. The detailed resource estimation is discussed in Sections 4.1.2 and 4.2.3.

2.1 Discretized Derivative Pricing

In order to map our derivative pricing problem into quantum states, we must discretize the values $\bar{S}^t$. Classically, this is not that important as high precision is available, but, in order to study the minimal qubit requirements, we need to consider discretization explicitly in the quantum case.

Let each value $\bar{S}^t$ be discretized into a different n -qubit register, i.e., mapped to a regular grid. We then define $\omega \in \Omega$ as the discrete space of paths. The price expectation is now a sum

$$\mathbb{E}(f) = \sum_{\omega \in \Omega} p(\omega) f(\omega), \quad (6)$$

where the probability $p(\omega)$ can be defined in multiple ways. For instance, one can take the mid-points of the grid cells so that

$$p(\omega) = \prod_{t=1}^T P(\vec{S}^t | \vec{S}^{t-1}), \quad (7)$$

where the $\vec{S}^t$ are restricted to discrete midpoints. Or $p(\omega)$ can be defined as an integral over the discrete cells. These representations are the same in the limit of fine grids and in the following we will choose the midpoint method.

2.2 Price Space vs. Return Space

In Section 2, Eq. (2) introduces geometric Brownian motion to model the price on underlying assets. We call this the *price space* description of the underlying stochastic process. In price space, transition probabilities are given by a multivariate log-normal distribution.

An alternative, but equivalent representation, is to consider the stochastic process on the log-returns of the underlying assets, and perform all calculations in *return space*. When asset prices obey a log-normal distribution, then the log-returns are distributed normally. We define a vector of underlying log-returns for d assets at time t as $\vec{R}^t = (R_1^t, R_2^t, \dots, R_d^t)$. At any time t' we can calculate the price of asset j from return space using

$$S_j^{t'} = S_j^0 \prod_{t=1}^{t'} e^{R_j^t}. \quad (8)$$

The transition probabilities are then given by a multivariate normal distribution

$$P(\vec{R}^t) = \frac{\exp\left(-\frac{1}{2}(\vec{R}^t - \vec{\mu})^\top \Sigma^{-1}(\vec{R}^t - \vec{\mu})\right)}{(2\pi)^{d/2}(\det \Sigma)^{1/2}}, \quad (9)$$

where,

$$\mu = (\mu_1, \mu_2, \dots, \mu_d)^\top, \quad (10)$$

$$\mu_j = (r - 0.5\sigma_j^2) \Delta t, \quad (11)$$

and σ , Δt , Σ and r are the same Brownian motion parameters as in price space. Note that this is no longer conditioned on the value at the previous time step. In fact, the path distribution in return space consists of dT independent Gaussians.

Note that we have overloaded the notation from the price space formulation as these representations are interchangeable via Eq. (8). This calculation is needed when the stochastic process has been modeled in return space but the payoff is defined in terms of asset prices. In the following sections, it will be made clear from the context which space we are operating in.

Switching between price space and return space changes from log-normal distribution loading to normal distribution loading. In general, the loading of normals is easier since their underlying stochastic evolution is independent of the price at a previous time step which can be seen by comparing Eq. (3) and (11). As such, the probability distribution $P(\vec{R}^1, \vec{R}^2, \dots, \vec{R}^T)$ across all T timesteps of the stochastic process can be computed simultaneously with

$$P(\vec{R}) \equiv P(\vec{R}^1, \vec{R}^2, \dots, \vec{R}^T) = \frac{\exp\left(\sum_{t=1}^T -\frac{1}{2}(\vec{R}^t - \vec{\mu})^\top \Sigma^{-1}(\vec{R}^t - \vec{\mu})\right)}{(2\pi)^{dT/2}(\det \Sigma)^{T/2}}. \quad (12)$$

This advantage can compensate for the quantum arithmetic needed to evaluate the exponentials in Eq. (8). We will leverage this advantage with the re-parameterization method in Section 4.2. Additionally, when working with derivatives that have payoffs defined in terms of log-returns directly and are independent of individual asset prices, this is another reason to work in return space.

Algorithm 3.1 Core approach to derivative pricing

Require: Parameters n , d , and T that are all positive integers.

Require: An operator $\mathcal{P}$ for loading a probabilistically weighted superposition of paths onto a register of ndT -qubits.

1. Apply operator $\mathcal{P}$ to prepare the quantum state

$$\mathcal{P}|0\rangle = \sum_{\omega} \sqrt{p(\omega)} |\omega\rangle. \quad (14)$$

2. Calculate $\delta(\omega) = \arcsin \sqrt{\tilde{f}(\omega)}$ into a quantum register

$$\sum_{\omega} \sqrt{p(\omega)} |\omega\rangle |\delta(\omega)\rangle. \quad (15)$$

3. Introduce an ancilla qubit and rotate the value of the $\tilde{f}(\omega)$ register into its amplitude:

$$\sum_{\omega} \sqrt{p(\omega)(1 - \tilde{f}(\omega))} |\omega\rangle |0\rangle + \sum_{\omega} \sqrt{p(\omega)\tilde{f}(\omega)} |\omega\rangle |1\rangle. \quad (16)$$

4. Use amplitude estimation to extract the probability of the ancilla being $|1\rangle$.

Output: The (discretized) expected payoff $\mathbb{E}(\tilde{f}) = \sum_{\omega} p(\omega)\tilde{f}(\omega)$. We rescale this to obtain $\mathbb{E}(f) = (f_{\max} - f_{\min})\mathbb{E}(\tilde{f}) + f_{\min}$.

3 Core Approach

Our approach to derivative pricing extends the quantum mean estimation method from [3]. In this section we review this approach and introduce an error analysis for the discrete case of option pricing.

Let the normalized discounted payoff of any path ω be given by

$$\tilde{f}(\omega) = \frac{f(\omega) - f_{\min}}{f_{\max} - f_{\min}} \in [0, 1]. \quad (13)$$

where $f_{\max}$ and $f_{\min}$ are the maximum and minimum possible payoffs respectively across all paths. The algorithm proceeds in four key phases. First, a probability distribution is loaded in form of a superposition over all possible paths. Second, payoffs for all possible paths are calculated in quantum parallel. Third, the expected payoff is stored in the amplitude of a marked state. Fourth, amplitude estimation is used to read out the amplitude using $\mathcal{O}(1/\epsilon)$ queries for a given target accuracy $\epsilon > 0$. This approach is detailed in Algorithm 3.1.

Note that Steps 1-3 in the Algorithm 3.1 load the exact answer after a single execution. Were it possible to read out an amplitude directly, then we could compute the expectation over all paths in a constant number of queries. This is, unfortunately, not possible, and so amplitude estimation introduces a linear overhead to extract an answer to a given precision. This is a key conceptual difference from classical Monte Carlo where samples from paths are taken. In Algorithm 3.1, we compute all possible paths and take (amplitude estimated) samples of the expected payoff.

Another important distinguishing feature of the quantum approach is that we must normalize the payoff in order to store it in the amplitude of a state. This normalization must be rescaled at the end and can have a critical impact on error scaling, as errors are also scaled up. In the Riemann summation method, discussed in Section 4.1, a version of this normalization rescaling can rapidly accumulate errors.

3.1 Amplitude Estimation for Derivative Pricing

Typically, path-dependent derivatives like autocallables and TARFs are priced using Monte Carlo or quasi-Monte Carlo methods. Paths $\omega = (\vec{S}^0, \vec{S}^1, \dots, \vec{S}^T)$ are generated by modeling the underlying stochastic process and then the expected payoff is calculated using the estimator

$$\mathbb{E}(f) \approx \frac{1}{M} \sum_{\omega=1}^M f(\omega). \quad (17)$$

This estimator converges to the true expected value with error $\epsilon = O(M^{-1/2})$ by the Central Limit Theorem [11].

This convergence can be quadratically accelerated to $\epsilon = O(M^{-1})$ using quantum amplitude estimation [4] for Monte Carlo [3, 5, 6]. Amplitude estimation takes as input a unitary operator $\mathcal{A}$ on $n + 1$ qubits such that

$$\mathcal{A}|0\rangle_{n+1} = \sqrt{1-a}|\psi_0\rangle_n|0\rangle + \sqrt{a}|\psi_1\rangle_n|1\rangle, \quad (18)$$

where the parameter a is unknown. Here, the final qubit acts as a label to distinguish $|\psi_0\rangle$ states from $|\psi_1\rangle$ states.

Amplitude estimation determines a by repeated applications of the operator³ $\mathcal{Q} = \mathcal{A}S_0\mathcal{A}^\dagger S_{\psi_0}$, where $S_0 = \mathbb{I} - 2|0\rangle_{n+1}\langle 0|_{n+1}$ and $S_{\psi_0} = \mathbb{I} - 2|\psi_0\rangle_n\langle 0|_0\langle 0|_0\langle \psi_0|_n$ are reflection operators. By using phase estimation and the quantum Fourier transform a can be determined with accuracy $O(M^{-1})$ [4]. Unfortunately, the required depth of the resulting quantum circuits scales as $O(1/\epsilon)$ and requires the use of a resource expensive quantum Fourier transform. Recent developments have introduced other approaches [12–17] that aim to reduce the resource requirements needed for amplitude estimation and can remove quantum phase estimation.

The most efficient variant of amplitude estimation known to date is Iterative Quantum Amplitude Estimation (IQAE) introduced in [14]. It has been shown empirically that IQAE outperforms the other known variants. Although it omits quantum phase estimation [18], it achieves a four times better performance than the canonical phase estimation approach. Further, it has been shown that for practical considerations, the following bound holds:

$$N_{\text{oracle}}^{\text{wc}} \leq \frac{1.4}{\epsilon} \log \left(\frac{2}{\alpha} \log_2 \left(\frac{\pi}{4\epsilon} \right) \right), \quad (19)$$

where $N_{\text{oracle}}^{\text{wc}}$ denotes the worst-case number of oracle calls, i.e., applications of $\mathcal{Q}$, to achieve an estimation error of $\epsilon > 0$ with confidence level $1 - \alpha$, $\alpha \in (0, 1)$.

We use the performance estimates from Eq. (19) for IQAE for our resource estimates in this work. This approach gives a full quadratic speedup, however it requires a quantum processor to successfully execute programs of oracle depth $\mathcal{O}(1/\epsilon)$. This large depth is a demanding requirement on QPU performance and is a dominant contributor to the required resources. Recent work [17, 19] has shown that it is possible to use shorter depth quantum programs $\mathcal{O}(1/\epsilon^{1-\beta})$ in exchange for less quantum advantage in total oracle calls $N_{\text{oracle}} = \mathcal{O}(1/\epsilon^{1+\beta})$ for $\beta \in (0, 1]$. Using shorter depth quantum programs means more tolerance to error and as such may result in less needed overhead for error correction. While there may be settings where this tradeoff is advantageous overall, we leave this analysis to future work.

3.2 Path Distribution Loading

In order for Algorithm 3.1 to achieve a practical quantum advantage, the resources needed for path loading and payoff calculation need to be taken into account. In some cases, there is an analytic form that can simplify path loading. For example in the case of path-independent derivatives, a distribution over paths is not needed. All that is needed is a distribution over final underlying prices $\vec{S}^T$, such as the log-normal distribution given by the Black-Scholes model. This means that the distribution could be analytically computed and then loaded either variationally or explicitly

³This is often called the Grover operator.

into quantum states. Unfortunately for quantum advantage, the analytic form for this distribution means that these derivatives are typically easy to price classically. Thus it is critical to focus on path dependent derivatives where a superposition over paths needs to be computed.

While the loading of general distributions is exponentially hard [20], several methods have been proposed. If the distribution is efficiently integrable, then there does exist an efficient quantum algorithm for loading, the Grover-Rudolph method [10]. However, the algorithm has limited applicability in practice for derivative pricing, because the relevant probability distributions still require Monte Carlo integration (albeit quantumly) which is precisely what we are trying to avoid by using Amplitude Estimation. More details on the insufficiency of this method are detailed in Appendix B.

An alternative method for loading the path distribution, using a qGAN [21], was proposed in [6]. While this has appeal for lower overhead loading, it is not yet clear how to anticipate what the overhead from training a given qGAN will be in practice.

3.3 Error Analysis

In this section, we investigate the various elements that contribute to the overall error in the quantum approach to derivative pricing. There are three main components that introduce error in Algorithm 3.1. Let $f_\delta = f_{\max}^{\text{disc}} - f_{\min}^{\text{disc}}$.

Truncation Error The price of a derivative is determined by an integral over all the possible values of the underlying price (or return). Given finite quantum memory, we cannot feasibly compute an integral over an infinite domain, and thus we restrict the domain of integration as follows: the prices/log-returns are restricted to a range $[B_l, B_u]$. This restriction of the domain leaves out a probability mass of α . Given an upper bound of $P_{\max}$ on the density functions at each step and an upper bound f_δ on the discounted payoff, we incur a truncation error which we denote by $\epsilon_{\text{trunc}} = P_{\max}^T f_\delta \alpha$.

Discretization Error This error (denoted by ϵ_{disc}) arises from the use of a Riemann Sum over a finite grid of points to approximate the integral. This error can be reduced by increasing the number of qubits (n) to approximate the sum.

Amplitude Estimation Error Amplitude estimation incurs an error of ϵ_{amp} when using $1/\epsilon_{\text{amp}}$ repetitions of the state preparation procedure and price computation.

We now analyze the truncation and discretization errors in more detail.

3.3.1 Truncation Error

We present the truncation error in return space as it then extends to price space straightforwardly. Denote the maximum eigenvalue of the covariance matrix Σ by $\sigma_{\max}$. Using Chernoff tail bounds on Gaussians, the probability that the log-returns for asset i lie outside of the interval $[\mu_i - w\sigma_{\max}, \mu_i + w\sigma_{\max}]$ is upper bounded by $2e^{-w^2/2}$. By the union bound the probability that any log-return (d assets over T time steps) lies outside the interval $[B_l = (r - 0.5\sigma_{\max}^2)\Delta t - w\sigma_{\max}, B_u = (r - 0.5\sigma_{\max}^2)\Delta t + w\sigma_{\max}]$ is upper bounded by $2dT e^{-w^2/2}$. Let the initial asset prices lie in the range $[S_{\min}^0, S_{\max}^0]$. Then the corresponding interval in price space is given by $[S_{\min}^0 e^{B_l T}, S_{\max}^0 e^{B_u T}]$.

We can then define the truncated window of values for our dT different n -qubit registers that are w standard deviations around the mean for each time step. The truncation error of the integral already normalized by $P_{\max}^T f_\delta$ is then given by

$$\epsilon_{\text{trunc}} \leq 2dT e^{-w^2/2}. \quad (20)$$

3.3.2 Discretization Error

The final output of the amplitude estimation algorithm represents a Riemann Sum that approximates the truncated multidimensional integral. The integral is over dT variables corresponding to d assets over T time steps. We assume that each underlying asset/return is restricted to an interval $[B_l, B_u]$. To compute the discretization error, we apply a multidimensional variation of

the midpoint rule as follows: let there be n qubits used to represent each underlying asset, the domain is divided into 2^{ndT} cells, and corresponding to each value of the register we associate the value of the integrand at the midpoint of the corresponding cell. Assume that β provides an upper bound on the second derivatives of the integrand (this can be restated as saying that the deviation from linearity over a range of length l is bounded by $\beta l^2/2$).

We consider the error from discretization that accumulates over a single cell. Each cell has side length $(B_u - B_l)/2^n$ and is a hypercube of dimension l . Note by symmetry that the linear component of the deviation from the value at the center of the cell integrates to 0 over the cell. The error in each cell can thus be bounded by the integral of the term $\beta x^2/2$ over a dT -hypercube of side length $l = (B_u - B_l)/2^n$ centered at the origin.

$$\underbrace{\int_{l/2}^{l/2} \cdots \int_{l/2}^{l/2}}_{dT} \beta x^2/2 = l^{dT-1} \beta [\beta x^3/6]_{l/2}^{l/2} = \frac{\beta l^{dT+2}}{24} = \frac{\beta(B_u - B_l)^{dT+2}}{24 \cdot 2^{n(dT+2)}}. \quad (21)$$

Aggregating the error over all the cells, we have

$$\epsilon_{\text{disc}} = \frac{\beta(B_u - B_l)^{dT+2}}{24 \cdot 2^{2n}}. \quad (22)$$

In terms of the number of standard deviations used in the discretization and the largest eigenvalue of the covariance matrix σ_{max} , the total discretization error is bounded by

$$\epsilon_{\text{disc}} \leq \frac{\beta(2w\sigma_{\text{max}})^{dT+2}}{24 \cdot 2^{2n}}. \quad (23)$$

For a target discretization error, Eq. (23) also gives us the total number of qubits required to represent d assets for T timesteps, given by

$$ndT = dT \left\lceil \frac{1}{2} (\log_2(\beta/24) - \log_2(\epsilon_{\text{disc}}) + (dT + 2) \log_2(2w\sigma_{\text{max}})) \right\rceil. \quad (24)$$

The truncation and discretization errors apply in general to the methods we introduce, though each method has additional method-specific error sources which are discussed separately.

4 Methods for Advantage in Quantum Derivative Pricing

In the following sections we introduce two approaches that can work effectively for quantum derivative pricing in practice: Riemann summation and re-parameterization. Riemann summation was introduced in [7], and we present the first resource analysis for its application for quantum advantage. This analysis uncovers limitations in error scaling due to normalization. We then introduce a new method called *re-parameterization* that avoids the downsides of other methods and offers the first end-to-end path to quantum advantage in practice.

4.1 Riemann Summation

The Riemann summation method [7] gives an approach to construct the $\mathcal{P}$ path loading operator in Algorithm 3.1. Let $N = 2^{ndT}$ be the size of the Hilbert space that contains all possible paths. Let $\tilde{P}_{\text{max}}$ be the maximum value of the d -asset multivariate transition probabilities from Eq. (2). Then $\tilde{P}(\tilde{S}^t | \tilde{S}^{t-1}) = P(\tilde{S}^t | \tilde{S}^{t-1})/\tilde{P}_{\text{max}} \in [0, 1]$ are the normalized transition probabilities over all choices of $\tilde{S}^t$ and $\tilde{S}^{t-1}$. Let the asset price for each asset at each timestep be discretized in the interval $[0, S_{\text{max}}]$. The steps of the method summarized in Algorithm 4.1 calculate the price of the derivative with a normalization factor $1/P_{\text{max}}^T$, with $P_{\text{max}} = \tilde{P}_{\text{max}} S_{\text{max}}^d$. Critically, we note that the normalization factor in the final step scales exponentially in T . If $P_{\text{max}} < 1$ no normalization is needed, but this factor can be used to improve the performance. However, if $P_{\text{max}} > 1$, the error increases exponentially, which renders this approach impractical.

The normalization factor P_{max} is easier to handle in return space where the probability density function is given by Eq. (9). If we discretize the log-returns at each timestep for each asset to $\pm w$ times the asset's volatility σ_j , we have

Algorithm 4.1 Riemann summation pricing

Require: Parameters n , d , and T that are all positive integers.

Require: Access to operators $\mathcal{W}_t, t = 1, \dots, T$ that apply the transition probabilities of the stochastic process into an ancilla via

$$\mathcal{W}_t |\vec{S}^t\rangle_n |\vec{S}^{t-1}\rangle_n |0\rangle \mapsto |\vec{S}^t\rangle_n |\vec{S}^{t-1}\rangle_n \left[\sqrt{1 - \tilde{P}(\vec{S}^t | \vec{S}^{t-1} | 0)} |0\rangle + \sqrt{\tilde{P}(\vec{S}^t | \vec{S}^{t-1} | 1)} |1\rangle \right] \quad (25)$$

1. Apply Hadamards to ndT qubits to prepare an equal superposition of all paths.
2. Load the initial prices $\vec{S}^0$ into the zero-th nd -qubit register.
3. Apply each of the T transition operators $\mathcal{W}_t$ to construct

$$\begin{aligned} \frac{1}{\sqrt{N}} \sum_{\omega} |\vec{S}^0 \dots \vec{S}^T\rangle & \left[\dots + \sqrt{\prod_{t=1}^T \tilde{P}(\vec{S}^t | \vec{S}^{t-1} | 1 \dots 1)} |1 \dots 1\rangle_T \right] \\ & = \frac{1}{\sqrt{P_{\max}^T}} \frac{1}{\sqrt{N}} \sum_{\omega} |\vec{S}^0 \dots \vec{S}^T\rangle \left[\dots + \sqrt{p(\omega)} |1 \dots 1\rangle_T \right], \end{aligned} \quad (26)$$

where $N = 2^{ndT}$.

4. Calculate $\delta(\omega) = \arcsin \sqrt{\tilde{f}(\omega)}$ into a quantum register, obtaining

$$\frac{1}{\sqrt{P_{\max}^T}} \frac{1}{\sqrt{N}} \sum_{\omega} |\vec{S}^0 \dots \vec{S}^T\rangle \left[\dots + \sqrt{p(\omega)} |1 \dots 1\rangle_T \right] |\delta(\omega)\rangle. \quad (27)$$

5. Introduce an ancilla qubit and rotate the value of the $\tilde{f}(\omega)$ register into its amplitude:

$$\dots + \frac{1}{\sqrt{P_{\max}^T}} \frac{1}{\sqrt{N}} \sum_{\omega} \sqrt{p(\omega) \tilde{f}(\omega)} |\omega\rangle |1 \dots 1\rangle_T |1\rangle. \quad (28)$$

6. Use amplitude estimation to extract the probability of the ancilla being $|1\rangle$.

Output: The (discretized) expected payoff $\mathbb{E}(\tilde{f}(\omega)/P_{\max}^T) = 1/(P_{\max}^T N) \sum_{\omega} p(\omega) \tilde{f}(\omega)$. We rescale this to obtain $\mathbb{E}(f) = P_{\max}^T (f_{\delta} \mathbb{E}(\tilde{f}) + f_{\min})$.

$$P_{\max} = \frac{(2w)^d \prod_{j=1}^d \sigma_j}{(2\pi)^{d/2} (\det \Sigma)^{1/2}}. \quad (29)$$

When the d assets are uncorrelated, we have

$$P_{\max} = \left(\frac{2w}{\sqrt{2\pi}} \right)^d, \quad (30)$$

and therefore we need to choose $w \leq \pi/\sqrt{2}$ for $P_{\max} \leq 1$. However, choosing a small discretization window w increases the truncation error discussed in Section 3.3.1, and for $w \leq \pi/\sqrt{2}$ we have $\epsilon_{\text{trunc}} \geq 2e^{-\pi^2/4} \sim 0.17$, which increases proportionally to the number of assets and timesteps in the computation.

4.1.1 Riemann Summation Error Analysis

In addition to the truncation and discretization errors from Section 3.3, the Riemann summation approach includes errors due to scaling considerations and quantum arithmetic.

When working in return space, we only need one transition operator which computes Eq. (12) and performs the amplitude encoding of $\sqrt{p(\omega)}$ in Eq. (26). Assuming the transition operator introduces a maximum additive error ϵ_{dens} and the payoff operator computing Eq. (27) and Eq. (28) introduces payoff error ϵ_f , the total arithmetic error of the quantity we estimate using amplitude estimation is

$$\epsilon_{\text{arith}} = \frac{1}{N} \sum_{\omega} [(p(\omega) + \epsilon_{\text{dens}})(f(\omega) + \epsilon_f) - p(\omega)f(\omega)]. \quad (31)$$

Ignoring quadratic error terms, we have

$$\epsilon_{\text{arith}} \approx \frac{1}{N} \sum_{\omega} p(\omega)\epsilon_f + \frac{1}{N} \sum_{\omega} f(\omega)\epsilon_{\text{dens}} \leq \frac{\epsilon_f}{(2w\sigma_{\max})dT} + \epsilon_{\text{dens}}, \quad (32)$$

where we assume the payoff has been normalized to lie in $[0, 1]$ and the log-returns for each asset and each timestep have been constructed to discretize the domain $[-w\sigma_{\max}, w\sigma_{\max}]$.

The probability density error ϵ_{dens} arises from the computation of $|\arcsin \sqrt{P(\vec{R})}\rangle$ with $P(\vec{R})$ given by Eq. (12), and the ancilla rotation in Eq. (26). The term inside the exponential in Eq. (12) can be written as

$$-\frac{1}{2} \sum_{t=1}^T (\vec{R}^t - \vec{\mu})^\top \Sigma^{-1} (\vec{R}^t - \vec{\mu}) = -\frac{1}{2} \sum_{t=1}^T \sum_{i=1}^d \sum_{j=i}^d C_{ij} \bar{R}_i^t \bar{R}_j^t, \quad (33)$$

where $\bar{R} = R - \mu$ and C_{ij} are classical variables containing volatility and correlation parameters from the correlation matrix Σ . In Eq. (33), each calculation of $\bar{R}$ thus incurs an error of ϵ_A and there are $(d + \binom{d}{2}) \cdot T$ multiplications in total. Each $\bar{R}$ term is bounded by $|w|\sigma_{\max}$ by construction, where each quantum register representing a log-return R is constructed to represent values in the window $[-w\sigma_{\max}, w\sigma_{\max}]$. Using the error analysis for addition and multiplication in Appendix C.2, the total error in computing Eq. (33) is

$$\epsilon_{\text{sum}} = \left(\frac{2w\sigma_{\max} + n}{2^{n-p}} + \frac{1}{4^{n-p}} \right) \left(d + \binom{d}{2} \right) \cdot T. \quad (34)$$

Then using the error propagation analysis in Appendix C.2 for computing the exponential, square root, arcsine and sine functions on quantum registers which already contain arithmetic errors, we can bound ϵ_{dens} by

$$\epsilon_{\text{dens}} \leq \epsilon_{\text{sin}} + \epsilon_{\text{arcsin}} - \arcsin(0.5 - (\epsilon_{\text{sq}} + \sqrt{\epsilon_{\text{exp}} + \epsilon_{\text{sum}}})) + \arcsin(0.5). \quad (35)$$

Each rescaling we perform to the input variables introduces a corresponding rescaling error. In addition to the $P_{\max}$ rescaling discussed in the previous section, we also need to scale the

payoff by $1/f_\delta$ to lie in $[0, 1]$. The final answer thus needs to be multiplied by $P_{\max}^T f_\delta$ to account for these rescalings, and the error in the estimate of the *truncated* integral by amplitude estimation is therefore scaled by $P_{\max}^T f_\delta$. We then can bound the error in the Riemann Summation approach

$$\epsilon_{\text{total}} \leq P_{\max}^T f_\delta (\epsilon_{\text{trunc}} + \epsilon_{\text{disc}} + \epsilon_{\text{arith}} + \epsilon_{\text{amp}}), \quad (36)$$

where ϵ_{trunc} , ϵ_{disc} , and ϵ_{amp} are defined as in Section 3.3.

4.1.2 Resource Estimates

As an example, we consider a basket autocallable with 5 autocall dates and parameters $T = 20, d = 3$, and target an error of $\epsilon_{\text{total}}/f_\delta \leq 2 \times 10^{-3}$. We need to choose $w \sim 5$ for the truncation error in Eq. (20) to be within the total target error, and Eq. (30) gives $P_{\max} \approx 4^3$. This makes the scaling factor prohibitively large with $P_{\max}^T \approx 10^{40}$. However, there may be some methods to deal with this normalization issue, such as a method inspired by importance sampling and discussed in Appendix D.2. For the sake of argument, we continue the resource analysis assuming that some method could be invented to deal with the normalization, and set $P_{\max} = 1$.

Then, using resource calculations as discussed in Appendix D.1, we can bound $\epsilon_{\text{arith}} \leq 2 \times 10^{-3}$ with $n = 34$ and $p = 2$. Here p is the integer part of the fixed point representation as defined in Eq. (64). The $\mathcal{Q}$ operator in this case requires 23k qubits and has a T-depth of 26k, including the resources required to compute prices from log-returns using Eq. (8). For a choice of $\Delta t = 1/20$ and $\sigma_{\min} = 0.1$ we compute that $\beta \approx 17$. Choose $\sigma_{\max} = 0.4$ and $w = 5$. Thus for the choice of n , $\epsilon_{\text{disc}} \approx f_\delta 10^{-5}$ and $\epsilon_{\text{trunc}} \leq f_\delta \cdot 10^{-4}$. When the derivative is priced classically with Monte Carlo, we define the pricing error as the standard deviation of the calculated derivative prices over all simulated paths. We therefore pick $1 - \alpha = 0.68$ in Eq. (19) to calculate the number of oracle calls needed for a given target error at the same confidence level as classical Monte Carlo. Choosing a target ϵ_{amp} for the amplitude estimation of 10^{-3} , we then obtain $N_{\text{oracle}}^{\text{wc}} \leq 6k$. This means that the total T-depth is about 1.5×10^8 .

Using the same analysis, for a TARF contract (see Section 5.2) with $d = 1, T = 26$ and $\Delta t = 1/26$, assuming the underlying has annualized volatility $\sigma = 0.4$, a target error of $\epsilon_{\text{total}}/f_\delta \leq 2 \times 10^{-3}$ can be achieved with a total T-depth of 1.6×10^8 and 17k qubits. These resource estimates are summarized in Table 1.

4.2 Re-parameterization Method

The limitations of normalization in Riemann summation motivate the need for a new method for loading stochastic processes. In the re-parameterization method, we shift to modeling assets in return space. As described in Section 2.2, in return space underlying assets consist of uncorrelated normal distributions. We recognize that these different distributions can be loaded by preparing, in parallel, many standard normals and then applying affine transformations to obtain the required means and standard deviations. This approach extracts a specific subroutine - loading a standard normal into a quantum state - and uses it as a resource to load the full distribution of underlying paths. The normal loading subroutine itself can then be precomputed and optimized using variational methods. This is an advantageous combination of fault-tolerant quantum computing with variational compilation and will be discussed in Section 4.2.1. Overall the re-parameterization method avoids the normalization issues in Riemann summation and reduces the computational requirements.

The steps in re-parameterization pricing are described in Algorithm 4.2. We note that a path $\omega_R \in \Omega_R$ in this context refers to a series of log-returns $\vec{R}^1, \dots, \vec{R}^T$. The re-parameterization method removes the problematic dependence on $P_{\max}$, and the operators $\mathcal{G}_j$ can be implemented with relatively few resources using variationally trained circuits as discussed in the following.

4.2.1 Variationally Trained Gaussian Loaders

The standard Gaussian loading operator $\mathcal{G}$ can be pre-computed because in the re-parameterization method it is problem independent. In this section, we describe an approach to variationally optimize this operator. Let us consider the problem of preparing a standard normal distribution $g(x_i)$

Algorithm 4.2 Re-parameterization method pricing

Require: Parameters n , d , and T that are all positive integers.

Require: Access to an operator $\mathcal{G}$ that loads a standard Gaussian distribution $\sum_i \sqrt{g_i} |i\rangle$ into an n -qubit register. Let g_i be the value of the probability mass function for a standard Gaussian distribution discretized into 2^n -bins.

1. Apply dT Gaussian operators $\mathcal{G}$, to ndT qubits. This constructs

$$\bigotimes_{t=1}^T \bigotimes_{j=1}^d \mathcal{G} |0\rangle_n = \sum_{\omega_{\bar{R}}} \sqrt{p(\omega_{\bar{R}})} |\omega_{\bar{R}}\rangle_{ndT}, \quad (37)$$

where $\omega_{\bar{R}}$ runs over all 2^{ndT} different realizations of this multivariate standard Gaussian, and $p(\omega_{\bar{R}})$ denotes the corresponding probabilities.

2. Let $\Sigma = LL^\top$ be the Cholesky decomposition of the covariance matrix. Perform affine transformations $\vec{R}^t = \vec{\mu}^t + L^\top \vec{R}^t$ to adjust the center and volatility of each Gaussian. We denote the corresponding return paths and probabilities by ω_R and $p(\omega_R)$, respectively.
3. If the payoff can be computed directly from the log-returns, then we directly calculate $\delta(\omega_R) = \arcsin \sqrt{\tilde{f}(\omega_R)}$ into a quantum register

$$\sum_{\omega_R} \sqrt{p(\omega_R)} |\omega_R\rangle |\delta(\omega_R)\rangle. \quad (38)$$

If the payoff is defined in terms of prices and not just log-returns, then we first compute the price space path ω for each asset using $\vec{S}^t = \vec{S}^0 e^{\sum_{j=1}^t \vec{R}^j}$. This calculation can be done in parallel for each asset.

4. Introduce an ancilla qubit and rotate the value of the $\tilde{f}(\omega_R)$ register into its amplitude:

$$\sum_{\omega_R} \sqrt{p(\omega_R)(1 - \tilde{f}(\omega_R))} |\omega_R\rangle |0\rangle + \sum_{\omega_R} \sqrt{p(\omega_R)\tilde{f}(\omega_R)} |\omega_R\rangle |1\rangle. \quad (39)$$

5. Use amplitude estimation to extract the probability of the ancilla being $|1\rangle$.

Output: The (discretized) expected payoff $\mathbb{E}(\tilde{f}) = \sum_{\omega_R} p(\omega_R) \tilde{f}(\omega_R)$. We rescale this to obtain $\mathbb{E}(f) = f_\delta \mathbb{E}(\tilde{f}) + f_{\min}$.

defined on a discretized mesh of points $x_i = -w + i \Delta x$, with $i = 0, \dots, 2^n - 1$, and $\Delta x = 2w/2^n$. In the following example we fix the domain to $w = 5$, so that the full range of values considered is $2w = 10$. This choice leaves outside the domain a probability mass of $\sim 5 \times 10^{-7}$. We take into account the different metrics used for normalizing a function in real space and a wavefunction in a quantum register (which is normalized in such a way that the sum of its squared elements is one), and therefore the distribution we aim to load in the quantum register is

$$g(x_i) \times \Delta x, \quad \sum_i g(x_i) \times \Delta x = 1. \quad (40)$$

Notice that due to the finite truncation domain, the target distribution is normalized to $1 - \alpha$. In principle we can re-normalize the distribution to one in the chosen interval of width $2w$. Either way this choice provides a negligible difference when compared with the error we observe in the training.

The variational ansatz of choice is represented by a so-called *Ry*-CNOT ansatz, with linear connectivity (see Appendix F). In this work we introduce a novel strategy to optimize the circuit in this context, which relies on an energy-based method [22], and is also detailed in Appendix F. In short, the target cost function is the energy of the associate quantum harmonic oscillator problem, whose ground state is naturally Gaussian [23]. Here, we must be careful because the squared modulus of the solution of the discretized quantum harmonic oscillator matches a normal distribution only in the limit of $\Delta x \rightarrow 0$. To fix this we perform a subsequent re-optimization targeting directly the infinity-norm between the two distributions.

$$L_\infty = \max_i |g(x_i) \times \Delta x - \tilde{g}(x_i)|, \quad (41)$$

where the quantum state encoded in the register is defined by coefficients $\sqrt{\tilde{g}(x_i)}$.

We notice that training directly with Eq. (41) as a cost function is not sufficient, and pre-training using the energy-based approach is needed to produce accurate results. We indeed observe how the L_∞ cost-function displays a much more corrugated landscape in the circuit parameter space compared to the energy of the associate quantum harmonic oscillator problem.

It is important to note that the circuits needed to encode these Gaussian states for different choices of the register size n can be pre-trained and used for any derivative pricing problem, and therefore the training cost is not included in our overall resource estimations. We show in Fig. 1 results for different register sizes and depths of the circuit ansatz. More details are provided in Appendix F.

Figure 1: L_∞ errors from training variational Ry-CNOT circuits to approximate $\mathcal{G}$ for different register sizes n .

This numerical study shows that the state we can prepare variationally, approaches the target exponentially fast in the depth, and hence in the number of gate operations. This observation is in

good agreement with the expected behavior from the Solovay-Kitaev theorem [24], that provides an upper bound for the number of gates required to achieve a desired accuracy for a cost function. Indeed, for any target operation $U \in SU(2^n)$, there is a sequence of operators $S = U_{s_1} U_{s_2} \dots U_{s_D}$ in a dense subset of $SU(2^n)$, such that error in the energy ϵ decreases exponentially with the depth $D = \mathcal{O}(\log^c(1/\epsilon))$. Although the subset of $SU(2^n)$ operations generated by the entangler blocks in our circuit does not generate a dense subset of $SU(2^n)$ arbitrarily close to the exact unitary U (the generator of the target state), we can numerically observe that the exponential decrease of the error with the number of gates still holds.

We end this section investigating the portability of these results in the fault-tolerant regime, which is necessary for the applicability of the derivative pricing algorithm. While our numerical results provide evidence for a rather efficient Gaussian state preparation in terms of circuit depth for an *Ry*-CNOT ansatz, an additional step has to be made in view of a fault-tolerant implementation of such circuits. In this new-framework, the continuous rotation *Ry* gate needs to be expanded as a finite product of discrete operations. Following again the Solovay-Kitaev theorem, or more specialized results [25], it is possible to also have an efficient representation of any $SU(2)$ operator with a sequence of Clifford + T gates that scale logarithmically with the threshold error ϵ . We investigate how the results obtained before can be transferred in this regime where rotation angles can only take discretized values. We therefore assume that each parameter $\vartheta_{q_j}^k$ can only be represented in the format $j * 2\pi/M_{\text{digit}}$, where j is an integer. We numerically show in Appendix F that the error introduced by such digitization decreases systematically with the mesh size as $\mathcal{O}(1/M_{\text{digit}})$.

4.2.2 Error Analysis

The total error in the re-parameterization approach is

$$\frac{\epsilon_{\text{total}}}{f_\delta} \leq \epsilon_{\text{trunc}} + \epsilon_{\text{disc}} + \epsilon_{\text{arith}} + \epsilon_{\text{amp}}, \quad (42)$$

where ϵ_{trunc} , ϵ_{disc} , and ϵ_{amp} are the truncation, discretization, and amplitude estimation error bounded in Section 3.3. Here, the term ϵ_{arith} arises from the individual errors we introduce during the preparation of the Gaussians and the calculation of the payoff. Assuming that each Gaussian $g(x_i)$ we prepare has L_∞ error ϵ_{dens} and the payoff calculation introduces a max error of ϵ_f , the total error will be

$$\epsilon_{\text{arith}} = \underbrace{\sum_{x_1=-w}^w \dots \sum_{x_{dT}=-w}^w}_{dT} \left[\prod_{i=1}^{dT} (g(x_i) + \epsilon_{\text{dens}}) (f(\mathbf{x}) + \epsilon_f) - \prod_{i=1}^{dT} g(x_i) f(\mathbf{x}) \right], \quad (43)$$

where $\mathbf{x} = (x_1, x_2, \dots, x_{dT})$. Expanding the integrand and keeping only the linear error terms, we get

$$\epsilon_{\text{arith}} \leq 2wdT\epsilon_{\text{dens}} + \epsilon_f, \quad (44)$$

where we use that $\sum_{-w}^w g(x) \leq 1$ due to truncation of the probability mass function.

4.2.3 Resource Estimates

We calculate the resources required for the same basket autocallable as in Section 4.1.2, where $d = 3$, $T = 20$, $\Delta t = 1/20$, $\sigma_{\text{max}} = 0.4$, $\sigma_{\text{min}} = 0.1$ and $w = 5$, and the contract has 5 autocall dates. We further assume that each Gaussian is prepared with $n = 5$ qubits, such that $\epsilon_{\text{dens}} = 2 \times 10^{-6}$, $\epsilon_{\text{amp}} = \epsilon_f = 10^{-4}$, which gives us a total error of $\epsilon_{\text{total}}/f_\delta \approx 2 \times 10^{-3}$. From Fig. 1 we observe that we can prepare Gaussian states with $L_\infty \sim 2 \times 10^{-6}$ using 5 qubits and circuit depth 6, requiring 7 layers of *Ry* gates. With these inputs and using the resource calculations described in Appendix E, constructing the $\mathcal{Q}$ operator using re-parameterization requires 8k qubits and has a T-depth of 9.5k, which includes the computation of prices from log-returns, Eq. (8). For a target confidence level of $1 - \alpha = 0.68$, the total T-depth is 5.4×10^7 . With the re-parameterization method, pricing the TARF of Section 4.1.2 with $d = 1$, $T = 26$, $\Delta t = 1/26$ and $\sigma = 0.4$ to accuracy $\epsilon_{\text{total}}/f_\delta \approx 2 \times 10^{-3}$ requires total T-depth of 8.2×10^7 and 11.5k qubits. These resource estimates are summarized in Table 1.

5 Payoffs

The previous sections analyzed methods for performing steps 1-3 in Algorithm 4.2. This results in a quantum state representing a superposition of all possible paths. In this section, we apply payoff functions to these superpositions (step 4) so that the normalized expected discounted payoff is stored in an amplitude of an accumulator qubit. This allows amplitude estimation to extract this amplitude to complete the pricing algorithm. We also analyze the additional errors introduced by the payoff function. We cover two example derivative cases: autocallables and TARFs. In addition, throughout this section (unlike in the other sections), we will assume that payoffs are *not* discounted unless explicitly stated.

5.1 Autocallables

An *autocallable* contract is typically defined in terms of asset returns relative to predefined reference levels, and includes a *notional* value which is used to calculate the dollar value of the contract. For a single underlying, an autocallable consists of:

- a set of m binary options $\{(K_i, t_i, f_i)\}_{i=1\dots m}$ each with strike returns K_i , payment dates t_i , and binary payoffs f_i . Assume these are sorted so that $t_i < t_{i+1}$.
- a short knock-in put with strike K_{put} , barrier b and notional value k , and
- the condition that if any binary option has a non-zero payoff then all subsequent options at later times including the put option are knocked out.

The payoff f_i of the binary options is equal to a fixed number p_i when $\tilde{R}_c^{t_i} \geq K$ and zero otherwise, where $\tilde{R}_c^{t_i}$ is the cumulative return of the underlying asset at timestep t_i . The payoff of the put option is

$$f_{put} = \begin{cases} k(\tilde{R}_c^T - K_{put}) & \text{if } \tilde{R}_c^T < K_{put} \text{ and } \tilde{R}_c^i < b, \forall i \in \{0, \dots, T\} \\ 0 & \text{otherwise.} \end{cases} \quad (45)$$

A more detailed explanation of autocallable options and all the parameters mentioned above can be found in Appendix A.4.

Note that the minimum payoff of the option occurs when the return of the underlying asset falls to zero. Therefore, we can compute the normalized discounted payoff $\tilde{f}_i$ from Eq (13) as

$$\tilde{f}_i = \frac{e^{-rt_i} f_i + e^{-rt_m} k K_{put}}{f_{\max}^{\text{disc}} + e^{-rt_m} k K_{put}}, \quad (46)$$

where $f_{\max}^{\text{disc}}$ is the maximum possible discounted payoff among the possible discounted payoffs across all the binary options and r is the risk-free rate. We now give a sketch of the algorithm used to implement the payoff for an autocallable shown in Algorithm 5.1. In steps 1-2, we compute the strike and put qubits that indicate which of the payoffs $f_1, f_2, \dots, f_m, f_{put}$ are non-zero. In step 3, we control on the strike qubits to apply the appropriate rotation on the accumulator qubit corresponding to the normalized discounted payoff from the binary options. In parallel, we execute step 4 which stores the normalized discounted payoff from the put option into another register. Finally, in step 5, we use the previously computed register to apply the appropriate rotation on the accumulator qubit corresponding to the normalized discounted payoff from the put options. This procedure is illustrated in Fig. 2.

We elaborate on a few subtleties in Algorithm 5.1:

1. Throughout the algorithm, we can assume that we have access to $|\tilde{R}_c^i\rangle$ for all i . This is because these registers are set when computing the prices S^i of the underlying when loading in the paths onto the quantum state (Appendix E).
2. Strictly speaking, the output register in step 4 is equal to $|\arcsin(\sqrt{\tilde{f}_{put}})\rangle$ only if the part of the $|put\rangle$ qubit that is entangled to this final register is $|1\rangle$. However, in the next step, the R_y

Figure 2: A block diagram of the subroutine presented in Algorithm 5.1. $|\tilde{R}_c^i\rangle_T$ represents the set of all cumulative return vectors for all timesteps, $|\tilde{R}_c^{t_i} \geq K_i\rangle_m$ and $|put\rangle$ are the strike and put qubits respectively, $|\arcsin \sqrt{\tilde{f}_{put}}\rangle$ is the register containing the value of the arcsine of the normalized discounted payoff of the put option, $|\psi_b\rangle$ is the accumulator qubit with just the rotations from the binary options applied to it and $|\psi\rangle$ is the accumulator qubit with the full expected discounted payoff rotations applied to it. Junk qubits were omitted from the diagram for clarity.

rotations that are controlled on this register are also always controlled on the $|put\rangle$ qubit, so the instances when the register has an unknown value have no effect on the rotations applied to the accumulator qubit.

3. The autocallable can also be defined on a basket assets instead of just one. Typical examples include *BestOf* and *WorstOf*, where the return of the contract is based on the return of the best or the worst performing asset in the basket respectively. The only difference is that, in step 1, we would apply the comparators on all the assets and then compare the return of each asset $|\tilde{R}_{c,j}^i\rangle_{j=1\dots d}$ to find the largest or smallest as necessary. We can then use the fact that if the worst performing asset is above the strike price, then so are all the other assets. Conversely, if the best performing asset is below the strike price, then so are all the other assets.
4. If we want to price an autocallable option that does not have a knock-in put, then, not only can skip steps 4-5 of the algorithm, but we can also perform the algorithm by just using the sum of log-returns instead of the sum of returns. Then it would be sufficient to have access to the registers $|\sum_{j=1}^{t_i} R^j\rangle$ instead of $|\tilde{R}_c^{t_i}\rangle$ which would reduce the resources required in the loading of the paths in superposition into the quantum state.

We now discuss the error arising from Algorithm 5.1. Steps 1-2 are performed with logical operation circuits (Comparator, AND, OR) which introduce no error, while steps 3 and 5 require controlled- R_y rotations whose decomposition into T-gates is a function of an additive error ϵ , which we can choose depending on the desired accuracy of the calculation.

Step 4 is the most resource heavy component of the payoff circuit, which requires the computation of the quantum register $|\tilde{R}_c^T - K_{put}\rangle$, the division of that register by the classical constant in the denominator of Eq. (46), as well as the computation of the square root and arcsine of the register. We describe in detail the resource requirements for all the above circuit components in Appendix C.1, and the corresponding arithmetic and gate synthesis error in Appendix C.2.

Consider again the autocallable contract from Section 4.1.2 and Section 4.2.3 with 5 autocall dates, defined on $d = 3$ assets and simulated using $T = 20$ timesteps. We target a total additive payoff error ϵ_f which when distributed across the operations of steps 3, 5, 6 in Algorithm 5.1 determines the resources required by each component. For $\epsilon_f = 10^{-4}$, the circuit computing the autocallable payoff requires 1.6k qubits and a T-depth of 3.2k, assuming we can parallelize computations wherever possible. These resources are included in the end-to-end summary estimates of Table 1.

5.2 Target Accrual Redemption Forwards

In this section, we consider the payoff implementation for the second example derivative: TARFs. To simplify the discussion, we describe the TARF implementation for a single underlying in price space. TARFs are usually contingent on the price of the underlying asset rather than the return.

A *TARF* is:

Algorithm 5.1 Autocallable payoff implementation

Require: An autocallable with parameters $\{(K_i, t_i, f_i)\}_{i=1\dots m}$, K_{put} , b and k .

1. We apply in parallel a set of T comparators to obtain the register $|\tilde{R}_c^i < b\rangle_T$, a set of m comparators to obtain the strike register $|\tilde{R}_c^{t_i} \geq K_i\rangle_m$ and a single comparator to obtain the qubit $|\tilde{R}_c^T < K_{put}\rangle$.
 2. We then apply the necessary AND and OR operations on all the registers from the previous step to obtain the $|put\rangle$ qubit which is set to $|1\rangle$ if all of the conditions for the put option were fulfilled i.e. none of the binary options payed off, the put option was knocked in and $\tilde{R}_c^T < K_{put}$.
 3. Let $\theta_i = \arcsin(\sqrt{\tilde{f}_i})$. Serially, for each bit of the strike register we apply a controlled rotation of θ_i on the accumulator qubit conditioned on all previous bits having been zero. This is illustrated in Figure 3.
 4. We use the register $|\tilde{R}_c^T\rangle$ to compute the arcsine of the the normalized expected payoff using quantum arithmetic and obtain the register $|\arcsin(\sqrt{\tilde{f}_{put}})\rangle$ where $\tilde{f}_{put}$ is defined in Eq. (46).
 5. Finally we apply a rotation of $\theta_{put} = \arcsin(\sqrt{\tilde{f}_{put}})$ on the accumulator qubit on the condition that the put option is activated. This is done using a series of $R_y(2^i)$ rotations where each rotation is controlled on the i th qubit of the $|\arcsin(\sqrt{\tilde{f}_{put}})\rangle$ register and the $|put\rangle$ qubit.
-

- A forward price F , payment dates denoted chronologically by timesteps $t = 1, 2, \dots, T$, two strike prices $K_{upper} > F$ and $K_{lower} \leq F$, a knock-out price b , a constant α and an accrual cap $C \in \mathbb{R}_+$.
- At each timestep t the TARF has a payoff

$$f_t = \begin{cases} S^t - F & \text{if } S^t > K_{upper} \\ \alpha(S^t - F) & \text{if } S^t < K_{lower} \\ 0 & \text{otherwise.} \end{cases} \quad \begin{array}{l} \text{and the contract is not knocked out ('upper condition')} \\ \text{and the contract is not knocked out ('lower condition')} \\ \end{array} \quad (47)$$

- a knock-out condition that if at any time t the price is above $b > F$ all payoffs that haven't been paid yet (including the one for the current time t) are knocked out. We will call this the 'barrier condition'.
- an accrual cap condition such that if the total gain accumulated by any payment date exceeds C the contract holder only receives a payoff such that the total gains equals C and the rest of the forward contracts are knocked out. We will call this the 'cap condition'.

A more detailed explanation of TARFs and all the parameters mentioned above can be found in Appendix A.5.

The minimum TARF payoff occurs when the price of the underlying asset falls to zero, and the maximum payoff occurs when the payoff at every payment date is $b - F$ until the accrual cap is reached. Thus, we can compute the normalized discounted payoff $\tilde{f}(\omega)$ using Eq. (13) to be

$$\tilde{f}(\omega) = \frac{f^{\text{disc}}(\omega) + \sum_{j=1}^T e^{-rt_j} 2TF}{\sum_{j=1}^{C/(b-F)} e^{-rt_j} (b - F) + \sum_{t_j=1}^T e^{-rt_j} 2TF}, \quad (48)$$

Figure 3: Example circuit used in step 4 of Algorithm 5.1 to accumulate binary option discounted payoffs in an autocallable with 6 binary options. Here the qubits $s_0, \dots, s_5$ represent the boolean comparisons for the 6 strike values K_i . The normalized discounted payoff $\hat{f}_i$ at time t_i (given by a particular phase in the RY rotation) is only non-zero if all payoffs at previous times $t_{j < i}$ are zero. The overall normalized discounted payoff is loaded into the amplitude of qubit e_0 .

where $f^{\text{disc}}(\omega)$ is the discounted payoff of the TARF for the path ω and r is the risk-free rate.

Algorithm 5.2 details an implementation for a TARF payoff. In steps 1-3, we compute the ‘partial conditional payoffs’ f_t^{partial} at each timestep t , i.e. the payoff given that we ignore the cap condition. In steps 4-8, we compute whether the cap condition is fulfilled at each timestep and correct the partial conditional payoffs to take into account the cap condition (giving us the actual payoffs). Finally, in steps 9-11 we discount the payoffs, sum them up, normalize the result according to Eq. (48) and apply the appropriate rotation on the accumulator qubit. This procedure is illustrated in Fig 4.

Figure 4: A block diagram of the subroutine presented in Algorithm 5.2. $|S^t\rangle_T$ represents the set of all prices for all timesteps, $|f_t^{\text{partial}}\rangle_T$ represent the partial conditional payoffs for all timesteps, $|\sum_{j=1}^t f_j^{\text{partial}} \geq C\rangle_T$ represent the sum of partial conditional payoffs over each timestep, $|f_t\rangle_T$ represent the actual payoffs for all timesteps and $|\psi\rangle$ is the accumulator qubit with the full expected discounted payoff rotations applied to it. Large amounts of junk qubits were omitted from the diagram to avoid dangling wires.

An interesting note is that were we to ignore the discount factor, this algorithm would be much simpler: we would add an extra step after step 5 in which we would set the payoffs for all the paths in which the cap condition was fulfilled equal to C and then to skip ahead to steps 10 and 11.

An error analysis for the TARF payoff is very similar to that described in the previous section for autocallables. There are 2 main differences. First, when computing the payoff f_t^{partial} for the lower condition, we require a multiplication. We set the error on this multiplication to be 10 times lower than our final desired error to make it negligible. Second, when discounting the payoffs, we have to ensure an error of $\frac{1}{\sqrt{T}}$ times smaller than in the auto-callable case because we are adding the payoffs after discounting them, causing the errors to add in quadrature. For $\epsilon_f = 10^{-4}$, the circuit computing the TARF payoff requires 9k qubits and a T-depth of 6k, assuming we can parallelize computations wherever possible.

Algorithm 5.2 TARF payoff implementation

Require: A TARF with parameters $(F, T, K_{\text{upper}}, K_{\text{lower}}, b, \alpha, C)$.

1. We apply in parallel a set of comparators to obtain the registers $|S^t < K_{\text{lower}}\rangle_T$, $|S^t > K_{\text{upper}}\rangle_T$ and $|S^t > b\rangle_T$.
 2. We then apply the necessary AND and OR operations on all the registers from the previous step to obtain registers $|upper\rangle_T$ and $|lower\rangle_T$ where the t th qubit in each register represents whether the upper and lower conditions were partially fulfilled (not taking into account the cap condition).
 3. We use $|S^t\rangle$, $|upper\rangle_T$ and $|lower\rangle_T$ to create $|f_t^{\text{partial}}\rangle$ registers in parallel using quantum arithmetic and control-copies.
 4. For each t , we compute $\sum_{j=1}^t f_j^{\text{partial}}$ in series, storing each result in a separate register.
 5. We compute the $|\sum_{j=1}^t f_j^{\text{partial}} \geq C\rangle$ qubits in parallel for all t .
 6. Then applying AND and OR gates in series to the $|\sum_{j=1}^t f_j^{\text{partial}} \geq C\rangle$ qubits, we compute $|cap\rangle_T$ which has all $T - 1$ $|0\rangle$ qubits a single $|1\rangle$ qubit in the c th position where c is the timestep at which the cap condition occurred. We now have
 - (a) the value of f_t^{partial} stored in the register $|f_t^{\text{partial}}\rangle$
 - (b) the register $|\sum_{j=1}^t f_j^{\text{partial}} \geq C\rangle$ for every timestep t indicating whether the cap condition has been fulfilled at the current or previous timestep
 - (c) the register $|cap\rangle_T$ indicating which timestep the cap condition was fulfilled.
 7. We control-copy all the registers containing f_t^{partial} in parallel controlled on the cap conditions to give us $T - 1$ registers of $|f_{t \neq c}\rangle$ and an additional register $|\vec{0}\rangle$. We now have a register for each timestep with the correctly encoded payoff at that timestep, except for time c which is the timestep at which the cap condition was fulfilled.
 8. We then take care of timestep c by adding $C - \sum_{j=1}^{t-1} f_j$ to each register t controlled on the t th qubit of $|cap\rangle_T$ such that nothing changes except for the $|\vec{0}\rangle$ register in step 7 which now turns into $|f_c\rangle$.
 9. We apply the discount factor to each f_t in parallel to obtain $|f_t^{\text{disc}}\rangle$.
 10. We then add up all the discounted payoffs, normalize the sum and take the arcsine to obtain $|\arcsin(\sqrt{\sum_{j=1}^T \tilde{f}_j})\rangle$.
 11. Finally we apply a rotation of $\theta = \arcsin(\sqrt{\sum_{j=1}^T \tilde{f}_j})$ on the accumulator qubit. This is done using a series of $R_y(2^i)$ rotations where each rotation is controlled on the i th qubit of the $|\arcsin(\sqrt{\sum_{j=1}^T \tilde{f}_j})\rangle$.
-

6 Discussion

We provide a thorough resource and error analysis to price financial derivatives using quantum computers. In particular, we investigate autocallables and TARFs which are two important path-dependent options that are relevant in practice and computationally expensive to price classically. To achieve this we introduce the re-parameterization method: a new method to load stochastic processes that overcomes the limitations of existing approaches. Although we limit our analysis to geometric Brownian motion, our approach can be straightforwardly extended to other models e.g. to stochastic or local volatility methods by loading multiple independent stochastic processes and introducing a conditional or non-stationary re-parametrization.

Our resource estimates give a target performance threshold for quantum computers capable of demonstrating advantage in derivative pricing. Assuming a target of 1 second for pricing an autocallable option, the quantum processor would need to execute T-gates at a rate of 10MHz at a code distance that can support 10^{10} logical operations. Further improvements in reducing the T-depth for this algorithm would linearly lessen this requirement.

The resource estimates in this work concur with recent work [26] that emphasizes the importance of going beyond complexity scaling in order to understand thresholds for quantum advantage. In particular, the quadratic speedup available in amplitude estimation-based algorithms could be lost in the constant factor overheads of error correction.

Although current estimates target logical clock rates around 10kHz [27] (i.e. orders of magnitudes slower than our requirement), we are optimistic that future work on algorithms, circuit optimization, error correction, and hardware will continue to improve the required resource estimates and runtimes. For example, in the case of Shor’s algorithm, the estimated resource requirements have reduced by almost three orders of magnitude through careful analysis across several publications [28]. This work represents the first milestone on the journey towards quantum advantage for pricing financial derivatives and we are looking forward to future enhancements.

Further, we emphasize that the resource estimation approach here can be fruitfully applied to analyze thresholds in other financially relevant applications. As summarized in two recent reviews [19, 29] there are many potential areas for quantum advantage in finance where advantage thresholds would provide useful targets for both industry and the research community.

Acknowledgments

We thank Paul Burchard for guidance on the derivative pricing problem domain and Thomas Häner for useful discussions regarding quantum arithmetic. We thank Graham Griffiths, Alex Hurst, Dunstan Marris and Elmer Tan for their technical and business insights on derivative products. We thank Ryan Babbush for suggesting clarifications to the manuscript. SC contributed to this work during his internship at Goldman Sachs.

References

- [1] A. Prabha, S. Savard, and H. Wickramarachi, *Deriving the Economic Impact of Derivatives*, Tech. Rep. (Milken Institute, 2013).
- [2] F. Black and M. Scholes, “The pricing of options and corporate liabilities,” *Journal of Political Economy* **81**, 637 (1973).
- [3] A. Montanaro, “Quantum speedup of monte carlo methods,” *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **471** (2015), 10.1098/rspa.2015.0301.
- [4] G. Brassard, P. Hoyer, M. Mosca, and A. Tapp, “Quantum Amplitude Amplification and Estimation,” *Contemporary Mathematics* **305** (2002), 10.1090/conm/305/05215.
- [5] P. Rebentrost, B. Gupt, and T. R. Bromley, “Quantum computational finance: Monte carlo pricing of financial derivatives,” *Phys. Rev. A* **98**, 022321 (2018).
- [6] N. Stamatopoulos, D. J. Egger, Y. Sun, C. Zoufal, R. Iten, N. Shen, and S. Woerner, “Option pricing using quantum computers,” *Quantum* **4**, 291 (2020).

- [7] A. Carrera Vazquez and S. Woerner, “Efficient state preparation for quantum amplitude estimation,” *Physical Review Applied* **15** (2021), 10.1103/physrevapplied.15.034027.
- [8] D. J. Egger, R. G. Gutierrez, J. C. Mestre, and S. Woerner, “Credit risk analysis using quantum computers,” *IEEE Transactions on Computers* (2020), 10.1109/TC.2020.3038063.
- [9] S. Woerner and D. J. Egger, “Quantum risk analysis,” *npj Quantum Information* **5** (2019), 10.1038/s41534-019-0130-6.
- [10] L. Grover and T. Rudolph, “Creating superpositions that correspond to efficiently integrable probability distributions,” (2002), [arXiv:quant-ph/0208112](#).
- [11] R. Y. Rubinstein, *Simulation and the Monte Carlo Method*, Wiley Series in Probability and Statistics (Wiley, 1981).
- [12] Y. Suzuki, S. Uno, R. Raymond, T. Tanaka, T. Onodera, and N. Yamamoto, “Amplitude estimation without phase estimation,” *Quantum Information Processing* **19**, 75 (2020).
- [13] S. Aaronson and P. Rall, “Quantum approximate counting, simplified,” *Symposium on Simplicity in Algorithms*, 24–32 (2020).
- [14] D. Grinko, J. Gacon, C. Zoufal, and S. Woerner, “Iterative quantum amplitude estimation,” *npj Quantum Information* **7** (2021), 10.1038/s41534-021-00379-1.
- [15] K. Nakaji, “Faster Amplitude Estimation,” (2020), [arXiv:2003.02417 \[quant-ph\]](#).
- [16] T. Tanaka, Y. Suzuki, S. Uno, R. Raymond, T. Onodera, and N. Yamamoto, “Amplitude estimation via maximum likelihood on noisy quantum computer,” (2020), [arXiv:2006.16223 \[quant-ph\]](#).
- [17] T. Giurgica-Tiron, I. Kerenidis, F. Labib, A. Prakash, and W. Zeng, “Low depth algorithms for quantum amplitude estimation,” (2020), [arXiv:2012.03348 \[quant-ph\]](#).
- [18] M. A. Nielsen and I. L. Chuang, *Cambridge University Press* (2010) p. 702.
- [19] A. Bouland, W. van Dam, H. Joorati, I. Kerenidis, and A. Prakash, “Prospects and challenges of quantum finance,” (2020), [arXiv:2011.06492 \[q-fin.CP\]](#).
- [20] M. Plesch and Č. Brukner, “Quantum-state preparation with universal gate decompositions,” *Physical Review A* **83**, 10.1103/physreva.83.032302.
- [21] C. Zoufal, A. Lucchi, and S. Woerner, “Quantum generative adversarial networks for learning and loading random distributions,” *npj Quantum Information* **5**, 1 (2019).
- [22] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, “A variational eigenvalue solver on a photonic quantum processor,” *Nature Communications* **5** (2014), 10.1038/ncomms5213.
- [23] P. J. Ollitrault, G. Mazzola, and I. Tavernelli, “Nonadiabatic molecular quantum dynamics with quantum computers,” *Physical Review Letters* **125** (2020), 10.1103/physrevlett.125.260511.
- [24] C. M. Dawson and M. A. Nielsen, “The solovay-kitaev algorithm,” *Quantum Info. Comput.* **6**, 81–95 (2006), [arXiv:quant-ph/0505030](#).
- [25] P. Selinger, “Efficient clifford+t approximation of single-qubit operators,” *Quantum Info. Comput.* **15**, 159–180 (2015), [arXiv:1212.6253 \[quant-ph\]](#).
- [26] R. Babbush, J. R. McClean, M. Newman, C. Gidney, S. Boixo, and H. Neven, “Focus beyond quadratic speedups for error-corrected quantum advantage,” *PRX Quantum* **2** (2021), 10.1103/prxquantum.2.010103.
- [27] A. G. Fowler and C. Gidney, “Low overhead quantum computation using lattice surgery,” [arXiv:1808.06709](#) (2018).
- [28] C. Gidney and M. Ekerå, “How to factor 2048 bit rsa integers in 8 hours using 20 million noisy qubits,” *Quantum* **5**, 433 (2021).
- [29] D. J. Egger, C. Gambella, J. Marecek, S. McFaddin, M. Mevissen, R. Raymond, A. Simonetto, S. Woerner, and E. Yndurain, “Quantum computing for finance: State-of-the-art and future prospects,” *IEEE Transactions on Quantum Engineering* **1**, 1–24 (2020).
- [30] S. Herbert, “The problem with grover-rudolph state preparation for quantum monte-carlo,” (2021), [arXiv:2101.02240 \[quant-ph\]](#).
- [31] K. Kaneko, K. Miyamoto, N. Takeda, and K. Yoshino, “Quantum pricing with a smile: Implementation of local volatility model on quantum computer,” (2020), [arXiv:2007.01467 \[quant-ph\]](#).
- [32] P. Selinger, “Quantum circuits of t-depth one,” *Physical Review A* **87** (2013), 10.1103/physreva.87.042302.

- [33] T. Häner, M. Roetteler, and K. M. Svore, “Optimizing quantum circuits for arithmetic,” (2018), [arXiv:1805.12445 \[quant-ph\]](#) .
- [34] T. G. Draper, S. A. Kutin, E. M. Rains, and K. M. Svore, “A logarithmic-depth quantum carry-lookahead adder,” *Quantum Information and Computation* **6**, 351 (2006), [arXiv:quant-ph/0406142](#) .
- [35] D. Maslov and M. Saeedi, “Reversible circuit optimization via leaving the boolean domain,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **30**, 806 (2011).
- [36] Y. Takahashi, S. Tani, and N. Kunihiro, “Quantum addition circuits and unbounded fan-out,” (2009), [arXiv:0910.2530 \[quant-ph\]](#) .
- [37] E. Muñoz Coreas and H. Thapliyal, “T-count and qubit optimized quantum circuit design of the non-restoring square root algorithm,” *J. Emerg. Technol. Comput. Syst.* **14** (2018), [10.1145/3264816](#).
- [38] N. J. Ross and P. Selinger, “Optimal ancilla-free clifford+t approximation of z-rotations,” *Quantum Info. Comput.* **16**, 901 (2016), [arXiv:1403.2975 \[quant-ph\]](#) .
- [39] A. Bocharov, M. Roetteler, and K. M. Svore, “Efficient synthesis of universal repeat-until-success quantum circuits,” *Physical Review Letters* **114** (2015), [10.1103/physrevlett.114.080502](#).
- [40] T. Kim and B. Choi, “Efficient decomposition methods for controlled-r n using a single ancillary qubit,” *Scientific Reports* **8** (2018), [10.1038/s41598-018-23764-x](#).
- [41] C. Zalka, “Simulating quantum systems on a quantum computer,” *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* **454**, 313 (1998).
- [42] S. Wiesner, “Simulations of many-body quantum systems by a quantum computer,” (1996), [arXiv:quant-ph/9603028 \[quant-ph\]](#) .
- [43] P. K. Barkoutsos, J. F. Gonthier, I. Sokolov, N. Moll, G. Salis, A. Fuhrer, M. Ganzhorn, D. J. Egger, M. Troyer, A. Mezzacapo, S. Filipp, and I. Tavernelli, “Quantum algorithms for electronic structure calculations: Particle-hole hamiltonian and optimized wave-function expansions,” *Phys. Rev. A* **98**, 022322 (2018).
- [44] J. Stokes, J. Izaac, N. Killoran, and G. Carleo, “Quantum natural gradient,” *Quantum* **4**, 269 (2020).
- [45] S. McArdle, T. Jones, S. Endo, Y. Li, S. C. Benjamin, and X. Yuan, “Variational ansatz-based quantum simulation of imaginary time evolution,” *npj Quantum Information* **5** (2019), [10.1038/s41534-019-0187-2](#).

A Background on Derivatives

This section presents some examples of commonly used derivatives in the financial sector. Unlike in most other sections of this paper where all payoffs are assumed to be discounted payoffs, in this section they are by default *not* discounted unless explicitly stated.

A.1 Forwards

An example of a derivative is a forward contract, often simply called a forward. Here, the holder promises to buy or sell a certain asset to the issuer on a specified date in the future at a fixed price F known as the *forward price*. A simple path-independent example is where the holder promises to buy x amount of an asset at F dollars per asset m months from now. Forwards are typically *settled in cash* i.e. instead of the money and asset exchanging hands on the expiration date, a payoff is determined based on the value of the asset and there is only an exchange of money determined by this payoff. For example, if the price at the expiration date T of the asset is S^T , the payoff is given by $f(S^T) = x(S^T - F)$, where if $f(S^T) > 0$, the contract holder makes a profit (and the issuer a loss) and the opposite if $f(S^T) < 0$.

A.2 Options

Another example of a derivative is an option. Options can be viewed as conditional forwards. With an option contract, the holder has *the option* to buy or sell a certain asset to the issuer on some future date at a pre-determined price (unlike the forward where the issuer is obliged to buy or sell the asset). If the holder chooses to buy or sell the asset, we say that they have *exercised the option*. Similarly to the forwards, option contracts are usually settled in cash based on the value of the asset on the exercise date. An example of a path-independent option with a single underlying asset is a *European call option*, where the issuer has the option of buying an asset at a strike price K on expiration date. The payoff on expiration date can then be written as $f(S^T) = \max(S^T - K, 0)$. A European *put* option is where the issuer has the option of *selling* an asset at a strike price K on expiration date, which gives a payoff of $f(S^T) = \max(K - S^T, 0)$. Another example of a path-independent option is a *binary option* which has a fixed payoff if the underlying asset is above (or below) the strike at time T .

A.3 Path-dependence and Discounted Payoffs

An example of a *path-dependent* derivative is a *knock-out* European call option. This is the same as a European call option, but with an additional knock-out price (or *barrier*) b . If at any time from 0 to T the underlying asset goes above this value, then the contract is worth nothing. This path-dependent payoff function has the form

$$f(S^0, S^1, \dots, S^T) = \begin{cases} S^T - K & \text{if } S^T > K \text{ and } S^i < b, \forall i \in \{0, \dots, T\} \\ 0 & \text{otherwise.} \end{cases} \quad (49)$$

The inclusion of the value of the underlying at times other than T is what introduces path dependence. Another example is a *knock-in* put option which has payoff

$$f(S^0, S^1, \dots, S^T) = \begin{cases} K - S^T & \text{if } S^T < K \text{ and } S^i < b, \forall i \in \{0, \dots, T\} \\ 0 & \text{otherwise.} \end{cases} \quad (50)$$

Here the contract is knock-in because it only has non-zero payoff if the asset goes below some value b .

In the examples discussed so far, there has only been one *payment* date where an exchange takes place between the contract issuer and holder, at time T . It is possible (as we will see later) for some path-dependent options to have several payment dates, i.e. where several payments are made at different times throughout the course of the contract duration.

We now introduce the notion of a discounted payoff. As expected, the price today for any derivative is related to its expected payoff in the future. However we also want to take into account the time delay for the payoff to account for the opportunity cost of investing in a risk-free asset with interest rate r . If a contract has a payoff f_i at time t_i from today, we define the discounted payoff as $e^{-rt_i} f_i$.

The price of a derivatives contract is given by the expected value of the discounted payoff under the stochastic process for the underlying assets. In practice, path-dependent derivatives are much more difficult to price computationally and are often priced using Monte Carlo simulations of the paths. This is in contrast to some models for path-independent derivatives that can even have analytic solutions, such as the Black-Scholes model for European call options [2]. Path-dependent options present an opportunity to use quantum speedups for Monte Carlo to gain advantage. In this work, we will consider two specific examples of path-dependent derivatives: autocallables and target accrual redemption forwards (TARFs).

A.4 Auto-callable Options

A typical example of an auto-callable ('automatically callable') option is a set of binary options, each of which pays different binary payouts at different payment dates and then knocks out the whole product (i.e. voids all future payoffs) if it makes a non-zero payout at any of the payment dates. Autocallables are typically contingent on the returns of the underlying asset (as opposed to

directly on the price). More formally, let (K_i, t_i, f_i) be a binary option that has payoff f_i defined as

$$f_i = \begin{cases} p_i & \text{if } \tilde{R}_c^{t_i} > K_i \\ 0 & \text{otherwise.} \end{cases} \quad (51)$$

where p_i is a fixed dollar value and $\tilde{R}_c^{t_i}$ is the cumulative return of the underlying asset at time t_i defined as the product of the returns $\tilde{R}^{t_j}$ for $j = \{1, 2, \dots, i\}$. We have used the notation $\tilde{R}$ to represent the return to differentiate from R which we have used previously to represent the log return. An autocallable is then a set

$$\{(K_1, t_1, f_1), (K_2, t_2, f_2), \dots, (K_3, t_m, f_m)\}, \quad (52)$$

where $\{t_i\}$ and $\{p_i\}$ typically increase linearly. If any of the binary options (K_i, t_i, f_i) pays out a non-zero dollar amount (i.e. is *in the money*), then all subsequent options $\{(K_j, t_j, f_j)\}_{j>i}$ are knocked out i.e. voided.

In practice, these binary options are often bundled with a *short* knock-in put option i.e. a knock-in put option given to the *issuer* by the holder, which mitigates risk for the issuer and decreases the price for the holder. This put option is also typically contingent on the return space of the underlying asset. More formally, the payoff (to the holder) from the put option is defined as

$$f_{put} = \begin{cases} k(\tilde{R}_c^T - K_{put}) & \text{if } \tilde{R}_c^T < K_{put} \text{ and } \tilde{R}_c^i < b, \forall i \in \{0, \dots, T\} \\ 0 & \text{otherwise.} \end{cases} \quad (53)$$

where K_{put} and b are the dimensionless put strike and barrier parameters respectively and k is a constant notional value. We note that in the case where $\tilde{R}_c^T < K_{put}$, the payoff is negative, implying that the contract *holder* has to pay the contract *issuer*. As with the set of binary options, this put option is also knocked out if any of the binary options (K_i, t_i, f_i) is in the money. An example of the full payoff structure for an auto-callable option with a single underlying and 3 payment dates is given in Algorithm A.1.

Algorithm A.1 Auto-callable example

Require: The following displays the payoff scheme of a 3-year auto-callable with yearly payment dates, binary option strike return of 1.1, a knock-in barrier of 0.7, a put option strike return of 1 and a notional value of \$18. The timestep values i represents the elapsed time in years and the non-zero payoffs p_i at each year from the binary option are \$2*i*. Finally the cumulative return of the underlying asset at timestep i is denoted by $\tilde{R}_c^i$.

1. For i in $[1, 2, 3]$:
 - if $\tilde{R}_c^i \geq 1.1$: the contract holder receives \$2*i* from the issuer and the contract is immediately ended.
 2. After 3 years, if the contract was not previously ended and at any point in time in the last 3 years the cumulative return of the asset was less than 0.7:
 - if $\tilde{R}_c^3 \leq 1$: the contract holder pays the issuer \$18(1 - $\tilde{R}_c^3$)
-

It is possible (and simpler) to describe the autocallable option with a single underlying to be contingent directly on the price of the underlying asset as opposed to the return. However defining it as such does not allow us to trivially generalize it to the case of multiple underlying assets. This is because it is typical to tie the overall option payoff to the best or worst performing asset, where performance is defined in terms of returns. In principle the different underlying assets could have independent put strike returns but this is not common.

The contingent payoffs and the knock-in put mean that autocallables have a payoff that is strongly path dependent. This means that they are computationally expensive to price in practice,

sometimes taking five to ten seconds using classical Monte Carlo methods with at least forty thousand paths.

A.5 Target Accrual Redemption Forwards

A target accrual redemption note (TARN) is any derivative whose payoff is capped at a specified target amount.⁴ For the purposes of this paper, we will focus on a commonly used TARN called a target accrual redemption forward (TARF). A TARF is a set of forwards with a couple of knock-out conditions. Specifically, it is a derivative with a single underlying with several (typically 20-60) payment dates and a forward price F . Throughout the contract, we have two fixed strike prices $K_{\text{upper}} \geq F$ and $K_{\text{lower}} < F$. At each payment date t , the payoff is defined as:

$$f_t = \begin{cases} S^t - F & \text{if } S^t > K_{\text{upper}} \\ 0 & \text{if } K_{\text{lower}} \leq S^t \leq K_{\text{upper}} \\ \alpha(S^t - F) & \text{if } S^t < K_{\text{lower}} \end{cases} \quad (54)$$

where S^t is the price of the underlying at the payment date t and α is a positive constant. We note that when $S^t < K_{\text{lower}}$, the payoff is negative and hence the holder of derivative makes a loss. The constant α makes this loss asymmetric if it happens and is often one or two.

In addition, a TARF will have two knock-out conditions based on a knock-out threshold b and accrual cap C . The first condition states that if at any payment date the price of the underlying is greater or equal to b , the derivative contract is immediately knocked out (without payment for that date). The second condition is if at any payment date t the total gains of the holder are going to exceed the accrual cap C due to the payoff f_t , the contract holder instead only receives the amount such that their total gains sum up to C and the contract is then knocked out. An example of the full payoff structure for TARF with 52 payment dates is given in Algorithm A.2.

Algorithm A.2 TARF example

Require: The following displays the payoff scheme of a 1-year TARF with weekly payment dates, a forward price of \$20, strike prices of $K_{\text{upper}} = \$20$ and $K_{\text{lower}} = \$15$, an $\alpha = 2$, a knock-out threshold of \$30 and an accrual cap of \$5. The timestep values i represents the elapsed time weeks.. Finally the price of the underlying asset at timestep t is denoted by S^t .

1. Total := \$0
 2. For t in $[1, 2, \dots, 52]$:
 - if $S^t \geq \$30$: the contract is immediately ended.
 - if $\$20 \leq S^t < \30 : set $f_t = S^t - \$20$.
 - if $S^t < \$15$: set $f_t = 2(S^t - \$20)$ (this will be negative number).
 - if Total + $f_t \geq \$5$:
 - set $f_t = \$5 - \text{Total}$
 - the contract holder receives f_t from the issuer and the contract is ended.
 - else:
 - the contract holder receives f_t from the issuer (if f_t is negative then the holder effectively gives money to the issuer)
 - Total := Total + f_t
-

⁴The term historically referred only to notes (hence the name) but has now come to include any derivative with an accrual cap

B Insufficiency of Grover-Rudolph Loading

The Grover-Rudolph algorithm [10] is often cited as a method to efficiently create quantum superpositions that correspond to classical distributions. For a given probability distribution $\{p_i\}$ of a random variable x , the algorithm creates a quantum superposition of the form

$$|\psi(\{p_i\})\rangle = \sum_i \sqrt{p_i} |i\rangle. \quad (55)$$

The algorithm is inductive in nature and starts by assuming that there is a way to divide the probability distributions into some number 2^m of regions in the domain of interest and create the state

$$|\psi_m\rangle = \sum_{i=0}^{2^m-1} \sqrt{p_i^{(m)}} |i\rangle, \quad (56)$$

where $p_i^{(m)}$ is the probability for the random variable to lie in region i . Then it aims to add one qubit to the state of Eq. (56), to further subdivide the 2^m regions into a 2^{m+1} discretization of the probability distribution with an evolution of the form

$$\sqrt{p_i^{(m)}} |i\rangle \rightarrow \sqrt{\alpha_i} |i\rangle |0\rangle + \sqrt{\beta_i} |i\rangle |1\rangle, \quad (57)$$

where α_i (β_i) is the probability for the random variable to lie in the left (right) half of region i . Letting x_L^i and x_R^i denote the left and right boundaries of region i , the function

$$f(i) = \frac{\int_{x_L^i}^{\frac{x_R^i + x_L^i}{2}} p(x) dx}{\int_{x_L^i}^{x_R^i} p(x) dx} \quad (58)$$

is the probability that, given x lies in region i , it also lies in the left half of the region. If we can construct a circuit which performs the computation

$$\sqrt{p_i^{(m)}} |i\rangle |0 \cdots 0\rangle \rightarrow \sqrt{p_i^{(m)}} |i\rangle |\theta_i\rangle, \quad (59)$$

with $\theta_i = \arccos \sqrt{f(i)}$, then a controlled rotation of angle θ_i on the $m+1$ th qubit yields

$$\sqrt{p_i^{(m)}} |i\rangle |\theta_i\rangle |0\rangle \rightarrow \sqrt{p_i^{(m)}} |i\rangle |\theta_i\rangle (\cos \theta_i |0\rangle + \sin \theta_i |1\rangle). \quad (60)$$

After uncomputing $|\theta_i\rangle$, we are left with

$$|\psi_{m+1}\rangle = \sum_{i=0}^{2^{m+1}-1} \sqrt{p_i^{(m+1)}} |i\rangle, \quad (61)$$

which is the extension of the state in Eq. (56) to one extra qubit. Performing this iteration $n = \log_2 N$ times, we will have a discretization of the distribution over N total number of points across n qubits.

In practice, the efficiency of the Grover-Rudolph method relies on the ability to perform the integrals in Eq. (58) in superposition. The argument in the original formulation in [10] is that probability distributions that can be integrated efficiently classically using probabilistic methods (e.g using Monte Carlo) can be equivalently efficiently integrated quantumly. However, since the ultimate goal in quantum derivative pricing is to provide a faster alternative to Monte Carlo integration over a probability distribution, performing this integral as part of our initial state preparation without any corresponding quantum speedup, nullifies the advantage offered by amplitude estimation as an alternative to Monte Carlo. While efficient from a complexity point of view, this means that Grover-Rudolph is insufficient as a method for quantum advantage in derivative pricing.⁵

⁵During revisions of the manuscript, a new pre-print [30] rigorously demonstrated the insufficiency of the Grover-Rudolph method for quantum-accelerated Monte Carlo.

More recently, an approximate method to implement the Grover-Rudolph algorithm for standard normal probability distributions was presented in [31], where the authors suggest the expression in Eq. (58), written as

$$g(x, \delta) = \frac{\int_x^{x+\delta/2} p(x) dx}{\int_x^{x+\delta} p(x) dx}, \quad (62)$$

can be approximated as

$$g(x, \delta) \approx \frac{1}{2} + \frac{1}{8}\delta x + \mathcal{O}(\delta^2), \quad (63)$$

for small δ . As the δ parameter decreases with each iteration of the Grover-Rudolph algorithm adding a qubit to the discretization, the authors highlight that for $m \geq 7$ the approximation in Eq. (63) becomes sufficiently accurate. However, because the Grover-Rudolph construction is iterative, the $m < 7$ terms need to be computed before the above approximation becomes possible. As such, the integrals in Eq. (58) are computed classically and then loaded into the corresponding quantum registers. While this approximation allows the simplification of the general Grover-Rudolph algorithm for standard normal distributions after a certain point in the iteration, it does not change the fact that it requires computing integrals over the entire domain of the probability distribution, thus making it practically infeasible for the same reason as the original Grover-Rudolph method.

C Fixed-point Quantum Arithmetic Resources

This section reviews preliminaries for common quantum arithmetic operations and the synthesis of arbitrary rotations. These operations are used in resource estimation and error analysis. Quantum arithmetic is required for path loading using the Riemann summation method (Section 4.1) and the re-parameterization method (Section 4.2), as well as the payoff calculation described in Section 5. For the Riemann sum method, we need to perform all the arithmetic operations involved in Eq. (12) as well as compute the arcsine and square root of a quantum register for the payoff calculation in Eq. (15). We identify algorithms for performing individual arithmetic operations efficiently, where resources are usually reported as a number of Toffoli gates or T-gates. In cases where we employ arithmetic algorithms from previous work in the literature, we report the gate cost in terms of the gate set reported by the authors.

As we are working in the fault-tolerant setting, we estimate the T-depth of the circuits in a Clifford + T gate set decomposition and assume Toffoli gates can be constructed with a T-depth of one using ancilla qubits [32]. For each operation we assume that we can parallelize the resulting circuits wherever possible.

C.1 Resource Estimation

We perform all calculations in fixed-point arithmetic similarly to [33], which allows us to use the quantum algorithms for reversible function evaluation described therein. An n -bit representation of a number x is

$$x = \underbrace{x_{n-1} \cdots x_{n-p}}_p \cdot \underbrace{x_{n-p-1} \cdots x_0}_{n-p}, \quad (64)$$

where $x_i \in 0, 1$ denotes the i -th bit of the binary representation of x and p denotes the number of bits to the left of the binary decimal point. The choice of n and p controls the error that we allow in each calculation as well as the resources required to perform arithmetic on the registers. Once we choose the values of (n, p) so that the overall arithmetic error is acceptable for the problem under consideration, we keep them constant throughout the analysis. It is possible that we can tailor these values for different components of the circuit and reduce the overall resources required, but for simplicity in this paper we ignore this potential optimization.

Let TF_f and T_f denote the number of Toffoli gates and the T-depth required to compute an arithmetic function or logical operation f . The estimates for the operations are functions of the

fixed-point register size (n, p) that will be used to represent the underlying quantum states involved in the computations.

Addition/Subtraction Using the algorithm described in [34], we can perform addition of two n -qubit registers in place with a Toffoli cost of $10n - 3w(n) - 3w(n-1) - 3\log_2 n - 3\log_2(n-1) - 7$ where $w(n)$ denotes the number of ones in the binary expansion of n , and a Toffoli depth of $\lceil \log_2(n) \rceil + \lceil \log_2(n-1) \rceil + \lceil \log_2(\frac{n}{3}) \rceil + \lceil \log_2(\frac{n-1}{3}) \rceil + 8$. Note that subtraction is given by $a - b = \sim(\sim a + b)$ and so can be implemented as an addition with $2n$ extra X gates, which does not change the Toffoli count.

We can turn an addition gate into a controlled addition gate by using the method shown in Figure 3 in [35]. This requires an additional n -qubit ancilla register, along with two sets of n parallel controlled swap gates. Each individual controlled swap gate is comprised of 3 Toffoli gates in series.

Multiplication For multiplication we follow the method from [33], which uses the controlled addition circuit in [36] and requires a Toffoli count of

$$\text{TF}_{\text{mul}}(n, p) = \frac{3}{2}n^2 + 3np + \frac{3}{2}n - 3p^2 + 3p. \quad (65)$$

This method can also be used for division of a quantum register by a classical value, which we do by inverting the classical value and employing the multiplication algorithm.

The controlled additions in the fixed-point multiplication method from [33] require ancilla qubits proportional to the register size, but the circuits include uncomputing the ancillas, meaning that they can be reused for each subsequent addition that is not done in parallel. Because we parallelize the computations across the d assets and T timesteps, we include an additional $T * d * n$ qubits when we count the total to account for these required ancilla qubits.

We can additionally parallelize each multiplication circuit, by considering the register of one factor as $z \geq 1$ independent registers of size n/z , and each controlled addition can happen in parallel for the z subregisters. This requires $n \cdot (z - 1)$ extra qubits and $z - 1$ additions to accumulate the z sub-results into the final result. $z = 1$ denotes that no extra parallelization is employed. If we can parallelize the pairwise accumulation additions as well, we arrive at a total T-depth cost of parallelized fixed-point multiplication given by

$$\text{T}_{\text{mul}}(n, z) = \lceil \frac{n}{z} \rceil \cdot (\text{T}_{\text{add}} + 6) + \lceil \log_2 z \rceil \cdot \text{T}_{\text{add}}. \quad (66)$$

$(\text{T}_{\text{add}} + 6)$ is the T-depth of a controlled addition discussed in the Addition/Subtraction section.

Square Root We employ the square root algorithm described in [37], which we extend for quantum registers in fixed-point representation. For an (n, p) -sized number x , we can compute $\sqrt{x}$ by treating x as an n -digit integer, and then shifting the result to the right $(n - p)/2$ times. This amounts to performing

$$\sqrt{x} \mapsto \sqrt{\frac{x * 2^{n-p}}{2^{n-p}}}. \quad (67)$$

The Toffoli count of this square root algorithm is [37]

$$\text{TF}_{\text{sq}}(n, p) = \frac{n^2}{2} + 3n - 4. \quad (68)$$

The T-depth of this algorithm as reported by the authors is given $\text{T}_{\text{sq}}(n) = 5n + 3$ and requires $2n + 1$ qubits.

Logical Operations For comparisons between quantum registers or between a quantum register and a constant, we use the logarithmic comparator from [34] with Toffoli/T-depth of $2\lceil \log_2(n - 1) \rceil + 5$, which includes uncomputing the intermediate ancillas. The logical *OR* operation for a 2-qubit input can be performed with a Toffoli/T-depth of one [6].

Exponential In [33], the authors introduce a generic quantum algorithm to calculate smooth classical functions using a parallel piecewise polynomial approximation. We apply this to estimate the resources of computing exponentials. The algorithm takes parameters k and M , which control the polynomial degree chosen for the piecewise approximations and the number of domain subintervals respectively. The total number of Toffolis is given by

$$\text{TF}_{\text{exp}}(n, p, k, M) = \frac{3}{2}n^2k + 3npk + \frac{7}{2}nk - 3p^2d + 3pk - d + 2Md(4\lceil \log_2 M \rceil - 8) + 4Mn. \quad (69)$$

This algorithm, which we also use to compute the arcsine function, requires k iterations of a multiplication and an addition, where k -degree polynomials are used for the approximation. Additionally, for M chosen subintervals, it requires M comparison circuits between the n -qubit input register and a classical value. Using the comparator from [34] with T-depth of $2\lceil \log_2(n-1) \rceil + 5$, the T-depth of a parallel polynomial evaluation circuit is

$$\text{T}_{\text{pp}}(n, z) = k(\text{T}_{\text{mul}}(n, z) + \text{T}_{\text{add}}) + M(2\lceil \log_2(n-1) \rceil + 5), \quad (70)$$

where z is the optional parallelization factor introduced in the resource estimation above for the multiplication circuit.

The qubit count for the parallel polynomial evaluation scheme for choices of the polynomial degree k and number of subintervals M is given by [33]

$$q_{\text{pp}}(n, k, M) = n(d+1) + \lceil \log_2 M \rceil + 1. \quad (71)$$

Arcsine To calculate the arcsine we employ the algorithm from [33] just as we do for the exponential. However, because the derivative

$$\frac{d \arcsin(x)}{dx} = \frac{1}{\sqrt{1-x^2}} \quad (72)$$

diverges near ± 1 , the authors use the transformation

$$\arcsin(x) = \frac{\pi}{2} - 2 \arcsin\left(\sqrt{\frac{1-x}{2}}\right) \quad (73)$$

to handle the interval $x \in [0.5, 1]$. Since the computation of the arcsine requires a conditional square root evaluation of the argument and, whenever we need to calculate an arcsine, we have to calculate the square root as well (e.g Eq. (15)), we can instead use the transformation

$$\arcsin(\sqrt{x}) = \frac{\pi}{2} - \arcsin(\sqrt{1-x}). \quad (74)$$

The resource estimation considerations then follow similarly to those in Appendix D.1/D.2 of [33]. We need:

- A comparator to check if $x < 0.25$ ($\sqrt{x} < 0.5$) that indicates whether we need to apply the above transformation, which would require two Toffoli gates assuming the value in the quantum register is normalized.
- A conditional subtraction and conditional copy depending on the comparator value above to either prepare $\sqrt{x}$ or $\sqrt{1-x}$. A conditional copy requires n Toffolis, a conditional subtraction requires $\text{TF}_{\text{add}} + n$ Toffoli gates [33].
- TF_{sq} Toffoli gates for the square root computation [37].
- The Toffoli gates required for the polynomial evaluation from [33] to compute the arcsine.
- A conditional copy and conditional subtraction depending again on the comparator result from the first step, to get either $\arcsin(\sqrt{x})$ for $x < 0.25$ or $\pi/2 - \arcsin(\sqrt{1-x})$ otherwise.

With the above considerations and the Toffoli count for the polynomial approximation of $\arcsin(x)$ from [33], the total Toffoli count for computing $|\arcsin \sqrt{x}\rangle$ is

$$\text{TF}_{\text{arcsq}}(n, p, k, M) = k \left(\frac{3}{2}n^2 + n(3p + \frac{7}{2}) - 3(p-1)p - 1 \right) + \frac{n^2}{2} + 11n + 2Md(4\lceil \log_2 M \rceil - 8) + 4Mn - 2. \quad (75)$$

The T-depth for computing $\arcsin(\sqrt{x})$ of a number x represented in a register of size (n, p) , calculated similarly to the exponential is

$$\text{T}_{\text{arcsq}}(n, p, z) = \text{T}_{\text{sq}}(n) + \text{T}_{\text{pp}}(n, z) + 8n + 6, \quad (76)$$

where $\text{T}_{\text{sq}}(n) = 5n + 3$ is the T-depth for the square root algorithm from [37].

The operation will require q_{arcsq} qubits, where the qubit requirements for the arcsine will be given by Eq. (71) for a choice of k and M , and $2n + 1$ for the square root operation

$$q_{\text{arcsq}}(n, k, M) = q_{\text{pp}}(n, k, M) + 2n + 1 \quad (77)$$

R_y We use $R_y(\theta)$ rotations in the variational preparation of Gaussians discussed in Sec. 4.2.1 and controlled- R_y rotations to encode the payoff into the amplitude of an ancilla in Eq. (16) as well as the transition probabilities in the Riemann summation method in Eq. (28). Using the method described in [38], an arbitrary single-qubit unitary can be performed within precision ϵ with a T-depth of approximately $3 \log_2(1/\epsilon)$.⁶

When the angle θ we wish to rotate is stored in a separate register $|\theta\rangle$, we require a series of $R_y(\theta_k)$ rotations, each controlled on the k th qubit of $|\theta\rangle$ where

$$\theta_k = \frac{2^k}{2^{n-p}}. \quad (78)$$

A single controlled- R_n can be performed with an R_n -depth of one, R_n -count of 3 and with a single ancilla qubit using the decomposition from [40]. However each rotation contributes an error ϵ so if $|\theta\rangle$ is an n -qubit register (with p bits to the left of the binary point), the end-to-end operation can be performed to precision ϵ with T-depth of at most $3n \log_2(n/\epsilon)$. We can reduce this depth slightly by noticing that the amplitude increase due to any controlled- R_n rotation where $\theta_k < \arcsin(\epsilon)$ is less than ϵ and hence is unnecessary. Therefore using that observation and Eq. (78), we compute the total number of rotations required to be $n - \max(\lfloor \log_2(\arcsin(\epsilon)) \rfloor + (n - p), 0)$. This gives us a final T-depth for a controlled- $R_y(\theta)$ operation of

$$\text{T}_{R_y}(n, p, \epsilon) = 3\tilde{n} \log_2(\tilde{n}/\epsilon) \quad (79)$$

where $\tilde{n} = n - \max(\lfloor \log_2(\arcsin(\epsilon)) \rfloor + (n - p), 0)$.

C.2 Error Analysis

Given the fixed-point representation of Eq. (64), each arithmetic operation involving registers results in some approximation error, depending on the specific method used. Here we outline the arithmetic error associated with each of the operations described in the previous section.

Addition/Multiplication We use the fixed-point addition and multiplication methods described in [33], where the addition of two (n, p) -sized registers introduces an error bounded by $\epsilon_A = \frac{1}{2^{n-p}}$, and the error associated with multiplication is at most

$$\epsilon_M(n, p) = \frac{n}{2^{n-p}}. \quad (80)$$

For (n, p) -sized registers X and Y , where each register already contains additive errors ϵ_X, ϵ_Y and each factor X, Y is bounded above by b , the error in the computation of $X \cdot Y$ is given by

$$\epsilon_{\text{mul}} = b * (\epsilon_X + \epsilon_Y) + \epsilon_X \epsilon_Y + \epsilon_M(n, p) \quad (81)$$

⁶A possible optimization could be to use the Repeat-Until-Success method described in [39]. This method shows that an arbitrary single-qubit unitary can be performed within precision ϵ with a T-depth of approximately $1.15 \log_2(1/\epsilon)$ using one ancilla qubit and measurement. However the method includes a probability of failure that complicates the analysis. Thus we leave the inclusion of this method to future work.

Exponential We employ the parallel polynomial evaluation methods from [33] to estimate the resources and associated error in computing exponentials. The error associated with the algorithm depends on choices for the degree of the polynomial approximation and the number of subintervals chosen, but the authors provide explicit error estimates and corresponding required resources in Table II for errors ranging from 10^{-5} to 10^{-9} . We use these in our overall error estimate. In our case, we compute the exponential of a register that itself contains arithmetic error ξ . Denoting the error in computing the exponential of a register ϵ_{exp} , the total arithmetic error in computing the exponential of a register can be approximated to first order in ξ in the Taylor expansion of $\exp(-x + \xi)$ as

$$\bar{\epsilon}_{\text{exp}} \lesssim \epsilon_{\text{exp}} + \xi. \quad (82)$$

Square root As discussed in the previous section, for square root computations we consider the square root algorithm described in [37], extended for quantum registers in fixed-point representation. The mapping in Eq. (67) introduces a maximum error of

$$\epsilon_{\text{sq}} = \frac{1}{2^{(n-p)/2}}. \quad (83)$$

When computing the square root of a register x which already contains (positive) additive error ξ , the total additive error from the square root operation is bounded by $\epsilon_{\text{sq}} + \sqrt{\xi}$. This is easily seen by observing that if we have a square root algorithm which gives us an estimate $\hat{x}$ with $|\sqrt{x} - \hat{x}| \leq \epsilon_{\text{sq}}$, then

$$\begin{aligned} |\sqrt{x + \xi} - \hat{x}| &\leq |\sqrt{x} - \hat{x}| + \sqrt{\xi} \\ &\leq \epsilon_{\text{sq}} + \sqrt{\xi} \end{aligned}$$

where the first inequality follows from $(\sqrt{x} + \sqrt{\xi})^2 = x + \xi + 2\sqrt{x\xi} \geq x + \xi$, for positive x and ξ .

Arcsine For the arcsine calculation we again use the polynomial evaluation method from [33], where the authors give sample resource estimates for error rates ranging from 10^{-5} to 10^{-9} . We want to bound the error from the computation of arcsine on a register containing an arithmetic error ξ to begin with. As discussed in Appendix C.1 we only need to compute $\arcsin(x)$ for $x \leq 0.5$. In addition, whenever we are computing the function $\arcsin(x)$ in our algorithms presented in the paper, we are only doing it for $x \geq 0$. This gives us a domain of $0 \leq x \leq 0.5$ for our $\arcsin(x)$ error calculation. Given this domain, we notice that the slope of $\arcsin(x)$ is always monotonically increasing with a maximum at $x = 0.5$. Therefore computing the error when $x = 0.5$ gives us the upper bound:

$$\bar{\epsilon}_{\arcsin} \leq |\arcsin(0.5) - \arcsin(0.5 - \xi)| + \epsilon_{\arcsin}, \quad (84)$$

where $\epsilon_{\arcsin}$ is the error from the computation of the arcsine from [33], given a choice of polynomial degree and number of subintervals.

Sine As discussed in the previous section, we compute the $\sin(\theta)$ function with a series of controlled- R_y rotations controlled on qubits from a register containing the angle θ . We can bound the error from the computation of $\sin(\theta)$ when the register that is supposed to represent θ is actually representing $\theta + \xi$ due to an arithmetic error. To quantify the upper bound, we notice that in the domain of $0 \leq \theta \leq \pi/2$, the slope of $\sin(\theta)$ is monotonically decreasing, and therefore has a maximum slope at $\theta = 0$. Therefore computing the error when $\theta = 0$ gives us the upper bound:

$$\begin{aligned} \bar{\epsilon}_{\sin} &\leq |\sin(0 + \xi) - \sin(0)| + \epsilon_{\sin} \\ &\leq \xi + \epsilon_{\sin} \end{aligned} \quad (85)$$

where we have used the inequality $\sin(a + b) \leq \sin(a) + b$ for $b \geq 0$ and where $\epsilon_{\sin}$ is the error arising from the gate decomposition of the R_y operator discussed in Appendix C.1.

D Riemann Summation

D.1 Riemann Summation Path Loading Resource Estimates

In this section, we examine the T-depth and qubit count required to compute Eq. (12) in a quantum register, and encode that value into the amplitude of an ancilla qubit as described in Algorithm 4.1. The calculation is done in log-return space (see Sec. 2.2) and it involves the resource estimates for the operations that we introduced in Appendix C.1.

Let T_f and q_f denote the T-depth and qubit count required for an operation f respectively. Assuming we can parallelize the computation across the d assets and T timesteps wherever possible, the contributions to the resources for computing $|\arcsin \sqrt{P(\vec{R})}\rangle$ with $P(\vec{R})$ given by Eq. (12) are

- T_{add} for computing the terms $(R - \mu)$ which can be done in parallel for d assets and T timesteps, where $T * d * n$ qubits are used to hold the log-returns R for all assets and timesteps.
- T_{mul} for all R^2 terms in the expansion of Eq. (33) (in parallel for all d and T), requiring $T * d * n$ additional qubits.
- $T_{\text{mul}} * \binom{d}{2} / (d/2)$ for all $R_i R_j$ terms in the expansion of Eq. (33) (parallel in d, T) and $T * \binom{d}{2} * n$ qubits.
- $T_{\text{add}} * \lceil \log \left(\binom{d}{2} + d \right) \rceil$ to sum all the terms in Eq. (33) in parallel. The qubits from the previous step can be reused here.
- T_{exp} to calculate the exponential in Eq. (9), requiring q_{exp} extra qubits with q_{exp} given by Eq. (71) for a choice of parameter values determined by the desired approximation accuracy.
- T_{arcsq} to calculate the arcsin and square root in $|\arcsin \sqrt{P(\vec{R})}\rangle$, with qubit resources given by Eq. (77).
- $T_{\text{add}} * (T - 1)$ and $(T - 1) * d * n$ qubits to calculate all the sums $R_j^{t=1} + R_j^{t=2} + \dots + R_j^{t=t'}$ for $t' \in [2, T]$ in Eq. (8).
- T_{exp} to calculate the prices across all assets and all timesteps in Eq. (8) in parallel, using $q_{\text{exp}} * d * T$ more qubits.
- A T-depth of $3n \log_2(n/\epsilon)$ to perform the ancilla rotation in Eq. (28) to precision ϵ , controlled on the register where $|\arcsin \sqrt{P(\vec{R})}\rangle$ is computed. This requires n ancilla qubits using the controlled- R_y decomposition from [40].

Moreover, we will need an additional register of size $T * d * n$ to implement the addition circuit used in [37] with constant T-depth and $(z - 1) * T * d$ extra qubits if we use the parallel multiplication scheme described in Appendix C.1 during the calculation of prices across assets and timesteps, where $z \geq 1$ is the optional parallelization factor we choose. Note that we have not included extra qubit counts to compute the $(R - \mu)$ terms and the sum in Eq. (33) because we can do these in place using the existing registers we have to hold each R_i . This is possible because after we compute the sums and exponentials in Eq. (8) (which we can do before computing the sums) we do not need the values of R_i again.

The total T-depth of the Riemann summation path loading process to precision ϵ for d assets and T timesteps using registers of size (n, p) is then

$$T_{\text{RS}}(n, p, d, T, \epsilon) = n^2 + 2n^2 \binom{d}{2} / d + 10 \left(\binom{d}{2} + d \right) + 10T + 9n + 5 + 3n \log_2(n/\epsilon) + 2T_{\text{exp}}(n, p, \epsilon) + T_{\text{arcsin}}(n, p, \epsilon), \quad (86)$$

where the dependency of T_{exp} and T_{arcsin} on ϵ denotes that the polynomial approximation parameters k and M in Eq. (70) for each function will depend on the target accuracy of the process. The total number of qubits required is

$$q_{\text{RS}}(n, p, d, T, \epsilon) = Tn \left(4d + \binom{d}{2} \right) + 3n + 1 + q_{\text{exp}}(n, p, \epsilon)(1 + dT) + q_{\text{arcsin}}(n, p, \epsilon). \quad (87)$$

D.2 Importance Sampling for Normalization in Riemann Summation

Within this section, we introduce a technique closely related to classical importance sampling to overcome the problem of the exponentially increasing scaling shown in 4.1. The main idea is to approximate the target distribution by another distribution that can be loaded efficiently and then use quantum arithmetic only to adjust for the (multiplicative) error.

Suppose a univariate probability density function $f : [0, 1] \rightarrow [0, P]$, with $P > 1$ and $\int_{x=0}^1 f(x)dx = 1$ and a payoff function $g : [0, 1] \rightarrow [0, 1]^7$. As introduced before, we can consider the scaled function $f(x)/P$ and a corresponding operator $\mathcal{F}$, as well as a corresponding operator $\mathcal{G}$ to prepare a state on $n + 2$ qubits given by

$$\frac{1}{\sqrt{N}} \sum_{i=0}^{N-1} |i\rangle_n \left(\sqrt{1 - f(x_i)/P} |0\rangle + \sqrt{f(x_i)/P} |1\rangle \right) \left(\sqrt{1 - g(x_i)} |0\rangle + \sqrt{g(x_i)} |1\rangle \right), \quad (88)$$

where we set $x_i = i/N$. Then, the probability of measuring $|11\rangle$ in the last two qubits is given by

$$\frac{1}{PN} \sum_{i=0}^{N-1} f(x_i)g(x_i), \quad (89)$$

and when multiplied with P corresponds to the Riemann sum approximating $\int_{x=0}^1 f(x)g(x)dx = \mathbb{E}[g(X)]$ for $X \sim f$.

Further, let us consider a probability distribution $h(x_i) \in [0, 1]$ that can be efficiently loaded into a quantum state, i.e., where we know how to efficiently construct a quantum operator $\mathcal{H}$ such that

$$\mathcal{H} |0\rangle_n = \sum_{i=0}^{N-1} \sqrt{h(x_i)} |i\rangle_n. \quad (90)$$

Suppose now that we have h such that $f(x)/(h(x)N) \in [0, 1]$ for all x , then we can construct a new operator $\mathcal{F}_h$ defined as

$$\mathcal{F}_h : |i\rangle_n |0\rangle \mapsto |i\rangle_n \left(\sqrt{1 - f(x_i)/(h(x_i)N)} |0\rangle + \sqrt{f(x_i)/(h(x_i)N)} |1\rangle \right). \quad (91)$$

Combining $\mathcal{H}$ and $\mathcal{F}_h$ leads to

$$\mathcal{F}_h \mathcal{H} |0\rangle_n |0\rangle = \sum_{i=0}^{N-1} \sqrt{h(x_i)} |i\rangle_n \left(\dots + \sqrt{f(x_i)/(h(x_i)N)} |1\rangle \right) \left(\dots + \sqrt{g(x_i)} |1\rangle \right), \quad (92)$$

which implies a probability of measuring $|11\rangle$ in the last two qubits given by

$$\frac{1}{N} \sum_{i=0}^{N-1} f(x_i)g(x_i), \quad (93)$$

i.e., the Riemann sum approximating $\int_{x=0}^1 f(x)g(x)dx = \mathbb{E}[g(X)]$ for $X \sim f$. Thus, if we can find such a probability distribution h , we can construct a state that directly corresponds to $\mathbb{E}[g(X)]$ without the need to rescale by multiplying P . It can be easily seen that for $P \leq 1$ we can set $h(x) = 1/N$ to recover the original approach without importance sampling.

In case of multivariate probability density functions, we distinguish three cases. First, separable functions that we can write as a product of univariate functions f_t for $t = 0, \dots, T$. In this case, the univariate approach can be applied directly and we need to find a corresponding h_t for each f_t . Second, non-separable multivariate probability density functions $f : [0, 1]^d \rightarrow [0, P]$, with $P > 1$ and $\int_{x \in [0, 1]^d} f(x)dx = 1$. Suppose we discretize each dimension using n qubits, i.e., we have in total N^d grid points. Then, we need to find a probability distribution h such that

⁷In the considered context, g will be applied only once. Thus, we can assume it to take values in $[0, 1]$ without changing the overall complexity of our approach.

$f(x)/(h(x)N^d) \in [0, 1]$ for all x , and the analysis is analog to the univariate case. Last, we consider the case of a multivariate probability density function coming from a stochastic process and given by

$$f(x_0, \dots, x_T) = f_0(x_0) \prod_{t=1}^T f_t(x_t | x_{t-1}), \quad (94)$$

where $x_t \in [0, 1]^d$ and $f_0(x_0), f_t(x_t | x_{t-1}) \in [0, P]$ for $t = 0, \dots, T$. Suppose a separable probability distribution

$$h(x_0, \dots, x_T) = \prod_{t=0}^T h_t(x_t), \quad (95)$$

that can be loaded efficiently as well as a corresponding decomposition $h_t(x_t) = h_t^t(x_t)h_t^{t+1}(x_t)$, with $h_T^{T+1}(x) = 1$. Then, we can write

$$\frac{f(x_0, \dots, x_T)}{N^{(T+1)d}} = \frac{f_0(x_0)}{h_0^0(x_0)N^d} h_0^1(x_0) \prod_{t=1}^T \frac{f_t(x_t | x_{t-1})}{h_{t-1}^t(x_{t-1})h_t^t(x_t)N^d} h_t^t(x_t)h_t^{t+1}(x_t). \quad (96)$$

Thus, if we find h such that the individual h_t can be efficiently loaded and

$$\frac{f_0(x_0)}{h_0^0(x_0)N^d} \in [0, 1], \quad \forall x_0 \quad (97)$$

$$\frac{f_t(x_t | x_{t-1})}{h_{t-1}^t(x_{t-1})h_t^t(x_t)N^d} \in [0, 1], \quad \forall x_{t-1}, x_t, \quad t = 1, \dots, T, \quad (98)$$

then, we can efficiently load the stochastic processes without the exponential scaling overhead P^{T+1} . Again, it can be easily seen that for $P \leq 1$ we can set $h_t(x_t) = h_t^t(x_t) = 1/N^d$ and $h_t^{t+1}(x_t) = 1$ to recover the original approach without importance sampling.

Note that even though we may not always find an h that satisfies all requirements, this approach can still help to lower the overhead coming from scaling.

E Re-parameterization Path Loading Resource Estimates

To prepare the standard normal distributions that we require in the re-parameterization loading approach, we can employ the variational method described in Sec. 4.2.1 and the corresponding gate/qubit cost depending on the desired accuracy of the approximation. In addition to that, we will also have to incur the cost of computing the affine transformation $\vec{R}^t = \vec{\mu}^t + L^\top \vec{R}^t$ as described in Algorithm 4.2. Note that the affine transformation is required when we need to calculate the asset prices from the log-returns, which for asset j at time t' will be

$$S_j^{t'} = S_j^{t=0} e^{\mu_j t' + \sum_{i=0}^d L_{ji}^\top \vec{R}_i^{t'}} = e^{\ln S_j^{t=0} + \mu_j t' + \sum_{i=0}^{d-1} L_{ji}^\top \vec{R}_i^{t'}}, \quad (99)$$

where $\vec{R}_i^{t'}$ is the i th component of the sum $(\vec{R}(t=1) + \vec{R}(t=2) + \dots + \vec{R}(t=t'))$. The graphical representation of the circuit that performs this calculation is shown in Fig. 5. One complication in Eq. (99) is that we cannot compute each asset price fully in parallel across the d assets, because the log-returns of any correlated assets will contribute to the computation of each other's price. In the case where all assets are pairwise correlated, we will need to compute the contributions to each asset's price from the log-returns of all d assets at that timestep, requiring in total d^2 additions to compute all asset prices per timestep. We can however perform d additions in parallel where the contribution of asset j 's return to the price of asset $(j+i)\%d$ is computed for a choice of $i \in [0, d-1]$, since all d such operations have distinct source and target registers. Then d rounds of additions will compute the term $\sum_{i=0}^{d-1} L_{ji}^\top \vec{R}_i^{t'}$ for all assets, and if we compute $(\vec{R}(t=1) + \vec{R}(t=2) + \dots + \vec{R}(t=t'))$ in a separate register for each t' and each asset, the above calculation can be also parallelized across all timesteps. This procedure is illustrated in Fig. 6.

The arithmetic error in computing Eq. (99) can be minimized by increasing the qubit register sizes to accommodate the largest values possible for the sums over the timesteps T and assets d . If each gaussian prepared in Eq. (37) is discretized using n qubits, then $n + \lceil \log_2 T \rceil$ qubits will be enough to hold the largest value of the sum represented by $\bar{R}_i^{t'}$. An additional $\lceil \log_2 d \rceil$ qubits will achieve the same for $\sum_{i=0}^{d-1} L_{ji}^\top \bar{R}_i^{t'}$, assuming the coefficients $|L_{ji}^\top| \leq 1$ for all i, j . This condition is not hard to satisfy for typical situations of practical interest, which we can argue by looking at the elements of the covariance matrix $\Sigma_{ij} = \Delta t \rho_{ij} \sigma_i \sigma_j$ (where by definition $|\rho_{ij}| \leq 1$). Typically, annualized volatilities are smaller than 100% (i.e. $\sigma_i < 1$) and the timestep usually satisfies $\Delta t < 1$, meaning the price of the underlying assets needs to be sampled more frequently than just yearly. If neither condition is satisfied however, we can choose a smaller Δt to ensure $|\Sigma_{ij}| < 1$, at the cost of increasing the number of timesteps in the calculation.

The contributions to the T-depth and qubit count for loading the paths and computing the asset prices in the re-parameterization approach for a derivative defined on d assets T timesteps are

- $T_{R_y}(n) \cdot (L + 1)$ T-depth for loading the gaussian states in Eq. (37) using the variational method from Sec. 4.2.1, where each Gaussian is prepared in parallel and the variational ansatz has depth L . This step requires $T * d * n$ qubits where n qubits are used to prepare each Gaussian state.
- $T_{\text{add}} * (T - 1)$ for calculating all the component-wise sums $\left(\vec{R}(t = 1) + \vec{R}(t = 2) + \dots + \vec{R}(t = t') \right)$ for $t' \in [2, T]$ in Eq. (99), requiring an extra $T * d * (n + \lceil \log_2 T \rceil)$ qubits (see Fig. 5).
- $T_{\text{mul}} * d$ to compute all contributions to $\sum_{i=0}^{d-1} L_{ji}^\top \bar{R}_i^{t'}$ in Eq. (99) and $T * d * \lceil \log_2 d \rceil$ more qubits.
- T_{add} to compute the $\mu_j t' + \ln S_j^{t=0}$ contribution in Eq. (99) across assets and timesteps.
- T_{exp} to compute the exponential in Eq. (99) across assets and timesteps, and $q_{\text{exp}} * d * T$ additional qubits with q_{exp} given by Eq. (71).

All in all, the total T-depth for path loading using the re-parameterization method to precision ϵ for d assets and T timesteps is

$$T_{RP}(n, d, T, L, \epsilon) = 3n \log_2(n/\epsilon)(L + 1) + 10T + d\bar{n}^2 + T_{\text{exp}}(\bar{n}, \epsilon), \quad (100)$$

with qubit count

$$q_{RP}(n, d, T) = (n + \bar{n} + q_{\text{exp}}(\bar{n}, \epsilon)) dT, \quad (101)$$

where $\bar{n} = n + \lceil \log_2 T \rceil + \lceil \log_2 d \rceil$.

F Method for Gaussian Loader Training

In this section we illustrate an approximate method to initialize the quantum register using the Variational Quantum Eigensolver (VQE) approach [22]. This algorithm features a parametrized circuit which in turn produces a parametrized state $|\psi(\{\theta\})\rangle$ that approximately represents the target state $|\phi_0\rangle$ and updates its parameters $\{\theta\}$ to optimize the expectation value of a suitable cost function. Here we show that the choice of the cost function to optimize is crucial for the success of the training.

Energy based training As a first method, we adopt a physics-based approach and define an operator H , such that its expectation value E assumes its lowest possible value, E_0 , when evaluated on the target state,

$$E_0 = \langle \psi_0 | H | \psi_0 \rangle. \quad (102)$$

In physics applications, the operator H is usually called the Hamiltonian, E the energy, and $|\phi_0\rangle$ the ground state. It is well known the Gaussian function

$$\phi_0(x) = \left(\frac{m}{\pi} \right)^{1/4} e^{-m(x-x_0)^2} \quad (103)$$

Figure 5: Circuit that computes T registers containing the cumulative log-returns $\bar{R}_i^{t'}$ from Eq. (99) for one asset at each timestep $t' \in [1, T]$. We first apply T $\mathcal{G}$ operators (see Sec. 4.2.1) to generate states corresponding to standard Gaussian probability distributions for each timestep, and then serially apply ADD operators which perform $|x\rangle|y\rangle|0\rangle \rightarrow |x\rangle|y\rangle|x+y\rangle$. The ADD operator is discussed in more detail in Appendix C. The circuit has $\mathcal{G}$ -depth of 1 and T_{add} -depth of $T - 1$ and can be applied in parallel for each asset in the derivative pricing calculation.

is the ground state of the quantum harmonic oscillator Hamiltonian

$$H = \frac{P^2}{2m} + \frac{m(X - x_0)^2}{2}, \quad (104)$$

where X is the position operator in real space, and $P = -i\frac{d}{dx}$ is the momentum operator [41, 42]. m is a parameter that determines the *variance* of the desired Gaussian distribution, and x_0 is the center of gaussian distribution. In this case, as we seek to find a state $\phi_0(x)$ such that $\phi_0^2(x) = \mathcal{N}(x_0, \sigma)$, we have to set $m = 1/(2\sigma^2)$. We notice that it is always possible to find a generating Hamiltonian function such that its ground state is the square root of the smooth distribution function that we aim to load.

To translate these considerations into an operational workflow we just have to define a way to compute the expectation value of Eq. (104) using a quantum computer. To this end we observe that the operator X^2 is diagonal in the computational basis, so it can be measured directly from the bit-string histogram counts $N_{\text{counts}}(j)$ generated by the repeated wavefunction collapses. The operator P^2 is diagonal in the momentum basis. This implies the addition of a centered Quantum Fourier Transform (QFT) circuit after the state preparation block. We use the centered Fourier transform to allow for negative momenta [23]. As introduced in the main text, we work in discrete position space $x_i = -w + i \Delta x$, with $i = 0, \dots, 2^n - 1$, and $\Delta x = 2w/2^n$. Without loss of generality we choose the domain to be centered at zero. The energy, $E = E_{X^2} + E_{P^2}$, can be computed in the following way,

$$E_{X^2} = \frac{1}{N_{\text{shots}}} \sum_{j=0}^{\mathcal{N}} \frac{m}{2} N_{\text{counts}}(j) (j \times \Delta x - x_0)^2 \quad (105)$$

$$E_{P^2} = \frac{1}{N_{\text{shots}}} \sum_{j=0}^{\mathcal{N}} \frac{1}{2m} N_{\text{counts}}(j) (j \times \Delta p)^2 \quad (106)$$

where N_{shots} is the total number circuit repetitions for the spacial and momentum basis. $N_{\text{counts}}(j)$ (with $0 \leq N_{\text{counts}}(j) \leq N_{\text{shots}}$, $\sum_j N_{\text{counts}}(j) = N_{\text{shots}}$) is the number of measurements that collapsed onto the qubit basis state corresponding to the binary representation of integer j . This strategy bypasses the need to obtain a Pauli representation of Eq. (104), which would include an exponentially increasing number of Pauli strings to be measured with the qubit register size.

The first step of our program is to verify numerically the possibility to prepare a state that systematically converges to Eq. (103), using a quantum circuit. Adopting a variational approach

$$\begin{aligned}
 &|\sum_{i=0}^{d-1} L_{i0}^T \bar{R}_0^t\rangle \xrightarrow{ADD} \xrightarrow{EXP} |S_0^t\rangle \\
 &|\sum_{i=0}^{d-1} L_{i1}^T \bar{R}_1^t\rangle \xrightarrow{ADD} \xrightarrow{EXP} |S_1^t\rangle \\
 &\vdots \\
 &|\sum_{i=0}^{d-1} L_{id-1}^T \bar{R}_{d-1}^t\rangle \xrightarrow{ADD} \xrightarrow{EXP} |S_{d-1}^t\rangle
 \end{aligned}$$

Figure 6: Circuits that compute asset prices $|S_i^t\rangle$ in separate quantum registers for d assets at timestep t using Eq. (99). The top figure shows the circuit which takes the cumulative log-returns for each asset created by the circuit in Fig. 5 as input states and computes $|\sum_{i=0}^{d-1} L_{ji}^T \bar{R}_i^t\rangle$ for each asset $j \in [0, d-1]$. There are d layers of multiplications where each MUL operator performs $|x\rangle|y\rangle|z\rangle \rightarrow |x\rangle|y\rangle|z+xy\rangle$. Note that the MUL operators in each box can be performed in parallel and therefore the entire circuit has depth of $T_{mul} * d$. Then, each computed register $|\sum_{i=0}^{d-1} L_{ji}^T \bar{R}_i^t\rangle$ requires one extra addition and exponentiation to compute the asset price $|S_i^t\rangle$ (bottom figure), which can also be applied in parallel for each asset. Both circuits can be applied in parallel across all timesteps T of the calculation. For more details on the ADD/MUL/EXP operators, see Appendix C.

will circumvent the need of costly quantum arithmetic operations at the expense of introducing sources of error which are always present in numerical variational approaches. The most trivial one concerns the possibility of getting trapped in local minima during the (classical) optimization procedure. The second, and more profound one, is linked with the representational power of trial states produced by the (shallow) quantum circuits.

Our main choice for the ansatz is the so-called R_y -CNOT circuit [43]. The initial state, defined on an n -qubit register which we set to $|0\rangle^{\otimes n}$, is evolved under the action of a unitary $U(\vec{\theta})$ to give the trial wave function $|\psi(\vec{\theta})\rangle$.

The circuit is made of a series of L blocks built from single-qubit rotations $U_R(\vec{\theta}^k)$, followed by an entangler U_{ENT} , that spans the required length of the qubit register. In our tests, we used the simplest choice of a ladder of CNOT gates with linear connectivity, such that qubit q_i is target of qubit q_{i-1} and controls qubit q_{i+1} , with $i = 1, \dots, n-2$. One additional layer of U_R gates is applied at the end, such that the number of variational parameters is $n \times (L+1)$.

Since the single-qubit rotations are all local operations, $U_R(\vec{\theta}^k)$ can be written as a tensor product of rotations of a single qubit:

$$U_R(\vec{\theta}^k) = \bigotimes_{i=0}^{n-1} R_Y(\vartheta_{q_i}^k), \quad (107)$$

where $R_Y(\vartheta_{q_i}^k)$ is a rotation on the Y-axis on the Bloch sphere of qubit q_i , and $k = 1, \dots, L+1$. The full unitary circuit operation is described by

$$U(\vec{\theta}) = U_R(\vec{\theta}^{L+1}) \overbrace{U_{\text{ENT}} U_R(\vec{\theta}^L) \dots U_{\text{ENT}} U_R(\vec{\theta}^1)}^{\text{L-times}}, \quad (108)$$

and the parametrized state is

$$|\psi(\vec{\theta})\rangle = U(\vec{\theta}) |0\rangle^{\otimes n}. \quad (109)$$

Note that the unitary $U(\vec{\theta})$ describes the full circuit, but not the pre-measurement change of basis required to collapse the wavefunction in momentum space as explained above.

For each value of parameters n and L , we repeat the optimization runs eight times in order to gather sufficient statistics, as it may happen that the optimizations remain stuck in suboptimal minima. Since we use classical emulation of the quantum circuits the only source of error in the optimizations originates from the classical optimizer. In our runs we first perform a warm up run with the COBYLA optimizer, followed by a longer run using the BFGS optimizer. To enhance the efficiency of the optimizations, the starting point for the VQE run at depth L , uses the optimal parameters found at previous optimization at depth of $L-2$ or $L-1$ when available. We notice that the part of the algorithm that concerns the classical optimization feedback can be greatly improved, for example using gradient based methods [44] or imaginary-time inspired update schemes [45].

L_∞ training refinements As discussed in the main text we use pre-optimized circuits obtained using the energy optimization method as a starting guess, and then re-optimize using the L_∞ as the cost function. In Fig. 7 we show indeed how the direct L_∞ optimization consistently fails to provide accurate results.

We show the complete outcome of the optimizations in Fig. 8. This careful numerical study shows that the convergence to the exact ground state is exponential in the depth, and therefore the number of gate operations.

Failure of the L_∞ norm direct optimization We provide an empirical explanation concerning the observed failure of the direct norm optimization technique. To this end we probe the cost function landscape for both methods, the energy-based and the direct L_∞ optimization. We start from an optimized parameter configuration $\vec{\theta}_0$ and we perform a *cut* in the parameter space, using

$$\vec{\theta} = \vec{\theta}_0 + \lambda \vec{\eta} \quad (110)$$

where $\vec{\eta}$ is a vector containing uniformly distributed random numbers in the range $[-1, 1]$, and $\lambda \in [-\pi, \pi]$ is a scalar which parametrizes the deviation from the optimal solution.

Figure 7: Optimization runs obtained with the energy-based method (green), the direct L_∞ optimization (red) and the mixed strategy where the energy based optimization is further refined using the L_∞ optimization (blue). We run eight independent runs for different values of the parameters n (number of qubits) and L (ansatz depth).

Figure 8: Left figure: L_∞ norm difference between the prepared and the target distribution as a function of the circuit depth L for different qubit register sizes n . We plot here the best of the eight independent optimizations for each parameter. Empty symbols correspond to optimizations performed using the energy of the quantum harmonic oscillator as a cost function, while solid symbols denote the refined optimizations using the L_∞ as cost function. Right figure: we plot the difference in energy of the associated quantum harmonic oscillator model. As expected the refinement targeting the L_∞ does not improve this quantity.

Figure 9: Cost function landscapes for the energy E (above), and the L_∞ norm (below) computed using three different cuts (red, blue and green colors) along the parameter space, and for two different setup $n = 5, 7$ and depths $L = 6, 10$ respectively.

In Fig. 9 we probe the cost function landscape for three different cut direction (e.g. three different realizations of the vector $\vec{\eta}$). We observe indeed that the cost function defined by the L_∞ norm is much more corrugated than the one defined by the energy E of the associate quantum mechanical toy problem, which instead displays a smoother surface. Crucially the basins of attraction of the energy cost-function and the L_∞ cost-function are overlapping (this happens because the ground state of the physical problem is very close to the Gaussian function we want to achieve), therefore the second optimization with the L_∞ norm does not remain stuck in high-cost local minima outside such basin.

Variational parameters digitization. While our numerical results provide evidence for a rather efficient Gaussian state preparation in terms of circuit depths for a parametrized circuit, an additional step has to be made in view of a fault-tolerant implementation of such circuits. In this new-framework, the continuous rotation R_y gate needs to be expanded as a finite product of discrete operations. Following again the Solovay-Kitaev theorem, or more specialized results [25], it is possible to also have an efficient representation of any $SU(2)$ operator with a sequence of Clifford + T gates that scale logarithmically with the threshold error ϵ . We investigate how the results obtained before can be transferred in this regime where rotation angles can only take discretized values. We therefore assume that each parameter $\vartheta_{q_i}^k$ can only be represented in the format $i * 2\pi/M_{digit}$, where i is an integer.

We adopt a simple protocol to optimize the parameters on an a grid. First we project the original continuous parameter values on the grid, choosing the closest grid point for each parameter. Subsequently, we perform a local search on the grid to find a better combination of the digitized parameters which minimize the L_∞ norm difference compared to the target distribution. We numerically show that the error introduced by such digitization decreases systematically with the mesh size. Interestingly, if we consider the error in the L_∞ norm difference introduced by this digitization, it decreases as $O(1/M_{digit})$. We observe that in all cases, we are able to obtain values comparable, or better, with the continuous solution, when the mesh size reaches $M_{digit} \sim 10^5$, which is equivalent to discretizing the space using $2\pi/M_{digit} \approx 0.0001$ rad.

Figure 10: L_∞ norm difference between the prepared and the target distribution as a function of the digitization mesh size M_{digit} for two different circuit depths L (blue and red), for the $n = 4$ qubit case. For each M_{digit} we *digitize* the eight parameter sets obtained by the previous independent optimizations (which were performed considering a continuous domain for the values of the rotation angles). Empty symbols refer to the full dataset, while the solid symbols highlight only minimum values in the set. Colored horizontal lines denote the best values obtained in the previous optimizations with a continuous domain of rotation angles for each L parameter. Interestingly, in some cases the digitization helps in escaping local minima and achieve slightly better solutions. Black diagonal lines are a guide-to-the-eye and represent the functions $1/M_{digit}$ and $0.1/M_{digit}$.

Figure 11: Same as Fig. 10 but with $n = 5$

Figure 12: Same as Fig. 10 but with $n = 6$