

An introduction to financial option pricing on a qudit-based quantum computer

Nicholas Bornman*

Superconducting Quantum Materials and Systems Center (SQMS), Fermi National Accelerator Laboratory, Batavia, IL 60510, USA

Date: October 25, 2023

The financial sector is anticipated to be one of the first industries to benefit from the increased computational power of quantum computers, in areas such as portfolio optimisation and risk management to financial derivative pricing. Financial mathematics, and derivative pricing in particular, are not areas quantum physicists are traditionally trained in despite the fact that they often have the raw technical skills needed to understand such topics. On the other hand, most quantum algorithms have largely focused on qubits, which are comprised of two discrete states, as the information carriers. However, discrete higher-dimensional qudits, in addition to possibly possessing increased noise robustness and allowing for novel error correction protocols in certain hardware implementations, also have logarithmically greater information storage and processing capacity. In the current NISQ era of quantum computing, a wide array of hardware paradigms are still being studied and any potential advantage a platform offers is worth exploring. Here we introduce the basic concepts behind financial derivatives for the unfamiliar enthusiast as well as outline in great detail the quantum algorithm routines needed to price a European option, the simplest derivative. This is done within the context of a quantum computer comprised of qudits and employing the natural higher-dimensional analogue of a qubit-based pricing algorithm with its various subroutines. From these pieces, one should relatively easily be able to tailor the scheme to more complex, realistic financial derivatives. Finally, the entire stack is numerically simulated with the results demonstrating how the qudit-based scheme's payoff quickly approaches that of both a similarly-resourced classical computer as well as the true payoff, within error, for a modest increase in qudit dimension.

1 Financial derivatives: a primer's primer

Financial assets can take a variety of forms, such as company stocks, government bonds, owning the debt of mortgage holders, as well as raw commodities such as precious metals and important foodstuffs. In reality, many of these assets are not themselves exchanged between a buyer and seller on the market. Rather, financial derivatives, which are simply contracts that are based on one or more of these underlying assets, are entered into and are the items that are in fact traded. The contracts themselves have values 'derived' from the underlying assets, and given that the assets' values are subject to various - largely indeterministic - market forces and other processes, the prime focus of derivative pricing in financial mathematics is to determine the fair value of such contracts.

The simplest kind of derivative is called a *futures* contract. A future is a publicly-traded contract which obliges the buyer to purchase a particular asset from the seller at a pre-agreed upon price at a later date. The utility of such a contract is that it allows investors either to make money by speculating (which is a risky investment strategy focused on taking advantage of price fluctuations to earn a profit), or to hedge (i.e. reduce the risk of) their betting position against market fluctuations. For example, if one speculates that the price of an asset is likely to increase at a later date, one could, today, purchase a futures contract in which one agrees to purchase the asset at the later date at a price lower than the speculated future price. On the other hand, to hedge their market position, an investor may purchase a futures contract since it allows them to lock in the purchase price of the asset, thereby creating a degree of certainty and controlling risk. However, buying or selling such contracts is never without risk: if a speculative investor purchases a future in the belief that an asset's price will increase, but the asset's price in fact decreases at the contract's maturation date, the investor will lose money as he will be forced to purchase the cheaper asset for more

*Email: nbornman@fnal.gov

than its actual value. As such, the decision of whether to buy or sell any derivative should take into account the expected return of investing in the most ‘risk-free’ investment - typically US government treasury bonds - available on the market.

Financial *options* contracts, on the other hand, are similar to futures, but with one key difference: they give the buyer the prerogative, not the obligation, to either buy or sell the asset. In particular, the simplest option, a European call (put) option, gives the buyer the right to purchase (sell) the asset (which we take to have a value of S_t at time t , with today being $t = 0$), at the strike price K once the contract is exercised at time $t = T$. If the asset’s price at the expiration date $t = T$ is S_T , the payoff (denoted by f) for the buyer of the call option would be

$$f(S_T) = \max(0, S_T - K). \tag{1}$$

This formula is intuitive: if it ends up being that $S_T > K$ at expiration, it makes sense for the buyer to exercise the contract as he’ll be acquiring an asset with a value of S_T for the cheaper price of K , and hence making a profit of $S_T - K$. However, if $S_T < K$, the buyer should not exercise the right to buy the asset since its mandated purchase price of K would be more than its actual market value.

If one somehow knew, today, that the asset’s price would be S_T at time T , the fair price of a European call option based on that asset would be the discounted future payoff, namely $e^{-rT} f(S_T)$. Here, r is the continuously-compounded return rate of risk-free government bonds, and the e^{-rT} discount factor accounts for the opportunity cost of not having invested in these bonds instead of the asset. However, asset prices do not evolve deterministically, but stochastically (see Fig. 1). Indeed, the value of a single asset, such as the price of a barrel of oil or a share of a publicly-listed company, is influenced by a number of factors such as reserve bank interest rates, traditional supply and demand economics, as well as global geopolitics and human psychology [1]. Furthermore, the global market is comprised of countless assets, all influencing one another in a terribly interconnected, unpredictable fashion. This indeterminism complicates matters so much that one cannot compute the exact future value of an asset, let alone an exotic option based on it. Instead, we use the discounted, *expected* payoff, namely $e^{-rT} \mathbb{E}_P[f(S_T)]$. This expectation value is taken with respect to a particular probability distribution P of S_t (we’ll elaborate on the distribution shortly).




Figure 1: A sample of stochastic asset price paths (modelled using geometric Brownian motion with drift 0.05 and volatility 0.2, c.f. Eqn. 2). The market’s inherent indeterminism means that we don’t know which path the asset’s price will follow over time; we can only compute the expected payoff, which is ‘averaged’ over all possible paths under the chosen stochastic model.

During early attempts to understand the complicated problem of option pricing, one of the first - and most important - option pricing models proposed was the Black-Scholes-Merton (BSM) model [2]. It takes the asset’s price to be primarily steered by the most common stochastic processes: one-dimensional Brownian

motion (see [3] for an intuitive primer on Brownian motion). If W_t denotes such a Brownian process and S_t the price of the asset at time $t \leq T$, the BSM model assumes that S_t follows this stochastic differential equation (known as geometric Brownian motion)

$$dS_t = \alpha S_t dt + \sigma S_t dW_t. \quad (2)$$

Here, the asset's evolution comprises two parts: the first term models a deterministic trend, or drift, in the 'mean' price of the asset (with α the percentage drift), whereas the second term (with σ the percentage volatility and dW_t a Brownian increment) models the unpredictable changes in the asset's price during each time increment. In the BSM model, the drift and volatility are assumed to be constants, independent of the asset price S_t and time (although more realistic stochastic volatility models, in which σ itself is driven by a separate Brownian process, exist [4]). Finally, it should be noted that in a formal treatment of stochastic calculus, a process such as W_t would need to be considered with respect to a probability measure. There are a number of such measures with respect to which W_t models standard Brownian motion. However, such a formal treatment is beyond the scope of this introduction.

W_t is stochastic and has a quadratic variance proportional to the time increment t , and ordinary calculus cannot solve Eq. 2 for S_t as a function of t and W_t . However, using the famed Itô's lemma (see [4]), we can solve for S_t as,

$$S_t = S_0 e^{\sigma W_t + (\alpha - \frac{\sigma^2}{2})t}, \quad (3)$$

where S_0 is the asset's price at $t = 0$. Fig. 1 represents samples of this function.

Eq. 3 gives, in an informal sense, a relation between S_t and one particular *realisation* of the Brownian motion W_t . One cannot definitively say what value S_t will assume at a later stage because one cannot definitively say what value W_t will assume: W_t has independent, random increments. Even so, these increments have a well-defined distribution: one property characterising W_t is that $W_t - W_0$ is normally distribution with mean 0 and variance t . This suggests that S_t itself potentially has a well-defined distribution. Indeed, re-writing Eq. 2 as $\frac{dS_t}{S_t} = \alpha dt + \sigma dW_t$, the process $\frac{dS_t}{S_t}$ is an Itô drift-diffusion process comprising a fixed drift term and a random volatile term. Given this, the Fokker-Planck equation [5, 6] is a partial differential equation describing the evolution of the probability density of a process subject to these two exact kinds of forces. Solving it (see [7]) for the probability density $p(S_t; t, \alpha, \sigma, S_0)$ gives

$$p(S_t; t, \alpha, \sigma, S_0) = \frac{1}{S_t \sigma \sqrt{2\pi t}} \exp\left(-\frac{\left(\ln\left(\frac{S_t}{S_0}\right) - (\alpha - \frac{1}{2}\sigma^2)t\right)^2}{2\sigma^2 t}\right), \quad (4)$$

i.e. the asset's price evolution follows a log-normal distribution.

We almost have all of the ingredients needed to compute the realistic discounted expected payoff for the European call option of Eq. 1. However, we need to hit all of this mathematics with a dose of real-world financial economics. Suppose that we start with the assumption that the market is rational (despite much evidence to the contrary [8] - as noted economist John Maynard Keynes opined: "The market can stay irrational longer than you can stay solvent"). One characteristic of a rational market is the assumption that, for a given capital investment, one cannot make a profit larger than the profit from having instead invested in a risk-free asset, without taking on risk. Such a market is called *arbitrage-free* and is based on the presumption that if an arbitrage opportunity were present in reality, market forces would speedily move to eliminate the opportunity¹. In this way, the market moves towards an arbitrage-free equilibrium.

With this in mind, the first fundamental theorem of asset pricing states that a market is arbitrage-free if and only if there exists a so-called *risk-neutral* measure [4]. Intuitively, such a measure P is one with respect to which the spot price (i.e. today's price) of an asset is exactly equal to the expected future price of the asset discounted to today, namely $S_0 \equiv e^{-rt} E_P[S_t]$, for all t . While we will not delve into such concepts further, a formal, in-depth definition of derivative pricing and measures is best understood in terms

¹For example, if it were possible to purchase a commodity at a lower price on one exchange and immediately sell it on another at a higher price (with negligible fees), many people would likely exploit this opportunity and earn profit, with no risk. The lower-priced market would likely quickly become aware of the disparity and consequently increasing the commodity's price (so as not to make a loss itself), thereby eliminating the arbitrage opportunity.

of measure-theoretic probability spaces [9]. However, in the result, choosing a risk-free neutral measure in the current context is equivalent to setting the drift α of the underlying asset to identically equal r , the risk-free interest rate². With $p(S_t; t, \alpha = r, \sigma, S_0)$ from Eq. 4 and the payoff function f from Eq. 1, we can now compute the European call option’s fair value, denoted ν , according to

$$\nu = e^{-rT} \mathbb{E}_{\mathbb{P}}[f(S_T)]. \quad (5)$$

An analytic expression for the above quantity exists (see [10]), given the relative simplicity of pricing a single asset whose payoff depends on only a single time point, namely the contract expiration date T . Unsurprisingly, a whole host of options far more complex than European options exist. Examples include American options (whereby one can exercise the option at *any* time prior to the expiration date of the contract) and Asian options (in which the payoff f doesn’t depend merely on the asset’s price at the contract’s expiration date, but rather upon the asset’s average price \bar{S} sampled at regular intervals over the contract length). Worth mentioning, too, are barrier options. In the case of ‘knock-out’ barrier options, the payoff f at the maturation date $t = T$ is 0 if the asset’s price S_t (with $t < T$) ever crosses a pre-defined ‘barrier’ value B at any time before the contract matures (and if it doesn’t, the payoff follows Eq. 1); for ‘knock-in’ barrier options, the payoff is 0 unless the asset’s price crosses the barrier at least once. Finally, one can string many of these options and concepts together in increasing complexity. For example, options contracts with a payoff function with multiple barriers are commonplace [11], as are options consisting of baskets of assets rather than only one. While we only discuss the mathematical details of European options in this text, it is clear that things can quickly get out of hand when trying to price real-world financial derivatives.

2 Qudit option pricing algorithm

Many real-world options depend not only on the asset’s final price, S_T , but rather on the asset’s value S_t at a number of time points, $0 < t \leq T$. In other words, the *path* the asset’s price follows, not merely the path’s endpoint. Such path-dependent option pricing calculations are expensive and investment firms typically carry out such computations overnight, on high-performance classical computers using Monte Carlo integration techniques³. It is *path-dependent* derivative pricing that is expected to gain the most from quantum computers [11]. However, our goal here is not to discuss the pricing of such complex options on a qudit-based quantum computer. Our goal is first to give an introduction to some of the concepts behind financial option pricing and work through the mathematical details, from start to finish, of how a simple option would be priced on a qubit-based quantum computer. As such, we will primarily restrict ourselves to a European call option and outline the protocol in great detail. Simultaneously, our second goal is to simulate the analogous protocol on a quantum computer with an information register comprised of qudits. While we don’t expect to see any advantage beyond the expected logarithmic increase in information storage for a qudit register in the completely analogous qudit protocol, qudit-based quantum computers offer a number of advantages over their qubit counterparts [13, 14] and any future implementation of derivative pricing on such a machine would need to begin with just such an analysis as this.

The traditional quantum derivative pricing algorithm is based on quantum Monte Carlo integration (see, for example, [15]), which is an extension of the quantum amplitude estimation (QAE) algorithm [16] to a weighted combination of amplitudes. The scheme is relatively simple. First, one prepares a ‘probability’ register of qubits into a weighted superposition of states. Each basis state in the superposition represents a potential asset price S_T , and the weighted coefficient represents the corresponding probability of realising

²Indeed, if this weren’t the case, an arbitrage opportunity would exist: if $\alpha < r$, one could construct a portfolio consisting of short positions of the asset S_t ; if $\alpha > r$, one could purchase the asset S_t , hold it, and sell it later for more than the purchase price. In this way, one would be guaranteed to, ‘on average’, make a profit [10].

³To classically price complex options for which nice analytic solutions do not exist, a large number M of potential asset price paths, $\{\mathbf{S}_i\}_{i=1}^M$ (where \mathbf{S}_i is a single realisation of the asset’s value evolution over the length of the contract), are first sampled from the underlying probability distribution \mathbb{P} (which can be either known, or implied by market data). An estimator $\hat{\mathbb{E}}_{\mathbb{P}}[f]$ of the true expected payoff $\mathbb{E}_{\mathbb{P}}[f]$ is then constructed by simply taking the average over the paths, i.e.

$$\hat{\mathbb{E}}_{\mathbb{P}}[f] = \frac{1}{M} \sum_{i=1}^M f(\mathbf{S}_i). \quad (6)$$

Finally, from the central limit theorem, as M grows, the difference between $\hat{\mathbb{E}}_{\mathbb{P}}[f]$ and $\mathbb{E}_{\mathbb{P}}[f]$ converges as $\mathcal{O}(M^{-1/2})$ [12]

S_T . Second, a series of controlled entangling gates is applied, with the probability register functioning as the control, to load the payoff function into the amplitude of a target ancilla qubit's $|1\rangle$ basis state. As such, this amplitude becomes a weighted combination of payoffs, with the weights being the probabilities of the payoffs themselves. This is akin to classically integrating the payoff over the probability density. And finally, QAE is used to estimate the amplitude of the ancilla's $|1\rangle$ basis state. The resulting quantity can easily be classically mapped to the expected payoff $E_P[f(S_T)]$ and the European option's fair value computed (as discounting $E_P[f(S_T)]$ is computationally cheap).

The chief advantage of pricing options using this algorithm on quantum computers rather than employing high performance classical resources is due to the quadratic speedup inherent in the quantum algorithm. Indeed, as outlined in footnote 3, the discrepancy between the true expected payoff and the estimator that classical Monte Carlo integration outputs, scales as $\mathcal{O}(M^{-1/2})$. The more samples paths M taken, the smaller the error. However, this convergence is not particularly fast: if one takes 100 times as many samples (i.e. $M \rightarrow 100M$), one gains only an order of magnitude of accuracy ($M^{-1/2} \rightarrow 10M^{-1/2}$). In the quantum algorithm to be outlined, however, this discrepancy scales as $\mathcal{O}(M^{-1})$ [16], where M here is the quantum analogue of the number of samples taken (in particular, M in the quantum case is the number of calls to a particular operator, called an oracle, which we shall discuss later). This is the much-touted speedup quantum computers provide.

Before the qudit extension of this algorithm is outlined in detail and simulated below, it is worth mentioning that the 'phase estimation' subroutine in QAE currently presents the biggest bottleneck in experimental implementations of quantum derivative pricing [17, 18]. However, it has been numerically demonstrated that one can still asymptotically achieve the quadratic speedup by forgoing the phase estimation subroutines and instead employing maximum-likelihood estimation [19]. This is also the case in a qudit-based algorithm, and as such, we shall discuss and implement this version of amplitude estimation without phase estimation.

2.1 Step 1: Load the distribution

Suppose that we have access to a quantum system comprised of n fully-connected qudits, each of which lies in a d -dimensional Hilbert space (such systems may arise from, for example, the linear modes of superconducting cavities coupled to transmon qubits [20]). The computational basis states of qudit j , where $j \in \{0, 1, \dots, n-1\}$, are $\{|0\rangle_j, |1\rangle_j, \dots, |d-1\rangle_j\}$. As such, the composite Hilbert space \mathcal{H} is spanned by the computational basis states of the form $|i\rangle_{(n)} = |i_0\rangle_0 \otimes |i_1\rangle_1 \otimes \dots \otimes |i_{n-1}\rangle_{n-1}$. Note that $i \in \{0, 1, \dots, d^n - 1\}$, expressed in base- d , is $i = i_0 + d^1 i_1 + \dots + d^{n-1} i_{n-1}$, and hence i holds $\log_2 d^n$ bits of information.

Our first task is to create an operator \hat{P} which loads the probability distribution of the asset S_T into the register of qudits. This probability distribution could be derived from either a financial model or implied from market data the modeller has at hand. For our single asset in the BSM model, the distribution is $p(S_T; T, r, \sigma, S_0)$ in Eq. 4 (we will suppress the parameters and denote this distribution by $p(S_T)$ from here onwards).

The asset price S_T can theoretically assume any non-negative real value. However, a finite register cannot faithfully represent this uncountable set (if the probability distribution for the asset were discrete, the situation would simplify slightly). So, we first truncate the asset price domain $[0, \infty)$ to $S = [S_{\min}, S_{\max}]$ by choosing S_{\min} and S_{\max} such that most of the probability mass lies within the truncated region⁴. The choice of this truncation is ultimately a design choice influenced by the desired accuracy, the distribution itself and available register size. However, for an underlying distribution P with mean μ and standard deviation σ , it is reasonable to set $S_{\min} = \max(0, \mu - 3\sigma)$ and $S_{\max} = \mu + 3\sigma$ to capture 3σ worth of data.

Next, we approximate P by sampling it at discrete points in S , since our qudit register can only hold d^n values. This is done by first dividing the region S into d^n partitions of equal length ω , where $\omega = (S_{\max} - S_{\min})/d^n$. Sampling S at the centre of each of these intervals gives a set $\{s_i\}$, where

$$s_i = S_{\min} + \left(i + \frac{1}{2}\right)\omega, \text{ with } i \in \{0, \dots, d^n - 1\}. \quad (7)$$

⁴This assumption necessarily introduces an error into the calculations as unlikely but 'extreme' events, which reside within the tails of the distribution, are discarded. The error introduced should be acceptable for short-tailed distributions, but it should be noted that potential extreme, unlikely events still affect the market in practice [21].

Eq. 7 provides the affine transformation between the integers i stored in the qudit register and the asset values $\{s_i\}$ they represent. This mapping will become important later.

Next, the distribution P is discretised by sampling it at points $\{s_i\}$. Finally, a set of probability values $\{p_i\}$, which will represent P , is formed by normalising the samples of P , i.e. $p_i = \frac{1}{\mathcal{N}} \times p(s_i)$ with $\mathcal{N} = \sum_{i=0}^{d^n-1} p(s_i)$. In a rough sense, p_i is the chance of the asset’s value assuming a value ‘close to’ s_i at the expiration date T .

With this setup, the first step of the qudit option pricing scheme is:

- 1 Construct a unitary operator \hat{P} which loads the probabilities $\{p_i\}$ into the coefficients of a register of qudits initialised in the ground state, i.e.

$$|0\rangle_{(n)} \xrightarrow{\hat{P}} \sum_i \sqrt{p_i} |i\rangle_n = \hat{P} |0\rangle_{(n)}. \quad (8)$$

Creating a general state preparation subroutine \hat{P} (Fig. 2) often requires very large numbers of fundamental gates (see Sec. 2.1.1). Whatever our choice for the unitary \hat{P} however, the state preparation step above only requires that the first column in the matrix representation of \hat{P} , denoted \mathbb{P} , be equal to $\mathbf{p} = [\sqrt{p_0}, \sqrt{p_1}, \dots, \sqrt{p_{d^n-1}}]^T$ (up to permutation, to match the chosen basis). While many potential choices of \mathbb{P} with this property exist, one simple unitary matrix with the requisite first column can be constructed using Householder transformations [22]: define $\mathbf{w} = \mathbf{p} - \mathbf{e}_0$ (with $\mathbf{e}_0 = [1, 0, 0, \dots, 0]^T$), and note that $\mathbf{w}^\dagger \mathbf{w} = 2(1 - \sqrt{p_0})$. Then, a \mathbb{P} with the correct first column is given by

$$\mathbb{P} = \mathbb{I} - \frac{\mathbf{w} \mathbf{w}^\dagger}{1 - \sqrt{p_0}}. \quad (9)$$

It can easily be verified that $\mathbb{P} \mathbf{e}_0 = \mathbf{p}$, which is the matrix representation equivalent of Eq. 8. We use this choice of \hat{P} in the forthcoming simulations.




Figure 2: Probability loading subroutine

2.1.1 Resource comments for \hat{P} (and general unitary gates)

Given the relative dearth of studies of quantum computers based on qudits, both theoretical and experimental, and what their fundamental gate set(s) would be (with which one must decompose the state one wishes to prepare, into), it is difficult to give accurate resource estimates for arbitrary state preparation as well as unitary compilation. Furthermore, one cannot make a meaningful ‘apples to apples’ gate resource comparison between qubit and qudit architectures without clearly analogous fundamental gate sets. However, some important comments can nonetheless be made.

Arbitrary state preparation and unitary compilation, using a conventional straightforward ‘decomposition’ scheme, on an n qubit gate-based system with a fundamental gate set comprised of arbitrary single-qubit gates and CNOT gates, generally requires $O(a^n)$ single qubit gates as well as $O(a^n)$ CNOT gates (where $a = 2$ in the case of state preparation [23, 24] and $a = 4$ for unitary compilation [25, 26]). This exponential scaling renders probability encoding schemes, via state preparation as above, infeasible for large n . This is because the current NISQ era of qubit-based quantum computing (which is characterised by imperfect gates, short qubit coherence times and no algorithmic error correction [27]) precludes algorithms with many operations and/or long execution times. These same concerns arise in a qudit-based computer comprised of the natural qudit analogues of the well-known X , Z and CZ gates (see [28, 13] for details). Indeed,

decomposing an arbitrary unitary gate $U \in \text{SU}(d^n)$, given n d -dimensional qudits and equipping the system with the aforementioned analogue qudit gates, requires $\mathcal{O}(nd^n)$ primitive operations [28], which is still a prohibitive scaling.

For the particular case, however, where a state encodes a *classically efficiently-integrable* distribution (such as the log-concave BSM model, Eq. 4) in a register of qubits (and qudits), state preparation can be efficiently accomplished [29]. But, in any event, the efficient loading of log-normal distributions requires pre-computing integrals of the probability distribution - often using classical Monte Carlo integration - whereas the crux of employing quantum amplitude estimation-based option pricing in the first place was to avoid the need for classical integration. Furthermore, the underlying distributions of more realistic financial models which capture important market dynamics which the BSM model does not, or distributions suggested purely by market data, may not have analytic representations at all. In these cases, the distributions may need to be encoded using other, contemporary techniques. With this in mind, it almost goes without saying that machine learning and neural networks are becoming increasingly important tools in classical finance to capture highly complex and inter-connected market dynamics efficiently (see, for example, [30, 31]), as well as model the market without assuming an underlying financial mathematical model [32]. The same holds true in the quantum context. Indeed, in [33], the authors outlined how a quantum version of a Generative Adversarial Network [34] can be used to create a circuit encoding the underlying distribution of training data samples by tuning the parameters of a parameterised model ansatz circuit. The parameters are tuned according to the feedback from a classical discriminator which is fed both the training data as well as the corresponding data generated from measurements of the parameterised quantum circuit itself. The upshot of such a protocol is the freedom to choose a polynomial-length parameterised ansatz circuit which can be efficiently deployed on the available quantum computing platform [35]. While incorporating neural network and machine learning techniques into quantum computing is still at the relative cutting edge of qubit-based systems, such techniques hold enormous promise and are almost certain to eventually feature in qudit-based systems too.

Finally, it is foolish to completely separate theoretical discussions of quantum circuit compilation from hardware considerations. The X , Z and CZ qudit gate analogues mentioned above naturally lend themselves to linear ion-trap [36, 13] and multi-port photonics [37, 38] quantum computing systems. As such, future commercial quantum computers based on such hardware would likely have access to these fundamental gate operations. On the other hand, circuit QED systems [39] and/or microwave cavities [40] are a novel technology which also hold much promise, as they have the potential for very long coherence times [41], relatively efficient universal control [42] and error correction protocols [43]. A popular potential gate set for cavity systems is one comprised of so-called SNAP gates along with cavity displacements [44, 42], which are intuitively different operations from qubit gates. SNAP and displacement gates are universal in that they can implement any arbitrary operation on the cavity qudit/transmon system, and importantly, if one wishes to address the first N Fock states of the cavity mode, the number of SNAP and displacement gate pairs scales as $\mathcal{O}(N)$ [42]. This potential linear scaling of the number of gates required to affect a chosen N -dimensional chosen unitary could offer a potential boon for microwave cavity based systems. As always though, much more work remains to be done to fully understand such systems.

2.2 Step 2: Compute the payoff

Suppose that we have constructed the operator \hat{P} , which loads the random variable (namely the asset price S_T at time T) into the register of qudits. We next move on to the subroutine which computes the payoff for the European call option. This payoff, Eq. 1, is $f(S_T) = \max(0, S_T - K)$ for a predetermined, classical strike price K (with $K \in [S_{\min}, S_{\max}]$), and will be loaded into the amplitude of a ‘payoff’ ancilla qubit⁵. The subroutine for this part of the option pricing protocol comprises two parts: a part which first compares the integer representation of S_T (i.e. i) with the integer representation of K (which we call k), and then a part which loads either 0 or $S_T - K$ into the ancilla’s amplitude, depending on the result of this comparison⁶. In both the qubit [35] as well as the current qudit option pricing schemes, these comparison-payoff encoding

⁵One could use a qudit instead of a qubit, but the most natural qudit analogue of the ordinary qubit option pricing protocol does not need the higher-dimensional nature of a potential payoff qudit.

⁶For a simpler payoff, such as that of a futures derivative (for which $f(S_T) = S_T - K$), this need for a comparison would fall away and result in a simpler payoff circuit subroutine.

steps happen ‘in parallel’ for all asset prices $|i\rangle_{(n)}$ in the superposition state of Eq. 8. Finally, note that k can easily be found using the mapping in Eq. 7: with $K = S_{\min} + (k^* + \frac{1}{2})\omega$ for some k^* , k is found by rounding k^* to the nearest integer. k , an integer in $\{0, 1, \dots, d^n - 1\}$, has a base- d representation of $k = k_0 + d^1 k_1 + d^2 k_2 + \dots + d^{n-1} k_{n-1}$.

The first ‘part’ needed in the subroutine is an operator which compares the classical variable k with i from the qudit basis state $|i\rangle_{(n)}$ and then changes the state of a ‘comparator’ ancilla qubit depending on which integer is greater. Specifically:

2.1 Create an operator \hat{C}_k which flips the state of an ancilla qubit $|0\rangle_c$ depending on the computational basis state of the qudit register, according to

$$|i\rangle_{(n)} |0\rangle_c \xrightarrow{\hat{C}_k} \begin{cases} |i\rangle_{(n)} |0\rangle_c & \text{if } i < k, \\ |i\rangle_{(n)} |1\rangle_c & \text{if } i \geq k. \end{cases} \quad (10)$$

Once such a circuit for \hat{C}_k has been constructed, it can be used to split the superposition state $\hat{P}|0\rangle_{(n)}$ from the first step, into two parts: one part corresponds to the cases where the asset price is less than the strike price K , and the other when it’s greater than K . In other words,

$$\hat{P}|0\rangle_{(n)} |0\rangle_c = \sum_i \sqrt{p_i} |i\rangle_n |0\rangle_c \xrightarrow{\hat{C}_k} \sum_{i < k} \sqrt{p_i} |i\rangle_{(n)} |0\rangle_c + \sum_{i \geq k} \sqrt{p_i} |i\rangle_{(n)} |1\rangle_c. \quad (11)$$

We shall see that the circuit for \hat{C}_k corresponds to a qudit version of the qubit quantum comparator (see, for example, [45]) and consists of a number of steps. First, the so-called $(d-1)$ ’s complement of k is found⁷. Next, the integer 1 is added to this complement, giving a string of digits we call k^c . Then, the sum of k^c and the quantum variable i , which is encoded in the state $|i\rangle_{(n)}$, is computed by creating the quantum circuit equivalent of a full-adder for each individual qudit comprising the state $|i\rangle_{(n)}$ [47, 48, 45]. A single full adder subroutine computes the sum of the j ’th qudit state $|i_j\rangle_j$, the corresponding j ’th digit in k^c , as well as any input carry digit from the previous step, and outputs the sum of these digits along with any output carry digit (which is fed into the full adder of the next step). As an example in Appendix 6.1 demonstrates, adding k^c to i in this manner is equivalent to calculating the difference $i - k$. Finally, the overall carry digit in the classical sum of i and k^c is identically encoded in the state of the comparator qubit of step 2.1 above. This qubit state hence indicates whether k is larger than or equal to i (in which case the comparator’s state is set to $|1\rangle_c$) or not (in which case it is $|0\rangle_c$). These steps are potentially confusing, especially for those unfamiliar with classical computing arithmetic. However, the logic behind these steps as well as potential gate representations for the subroutine \hat{C}_k , are outlined in great detail in Appendix 6.1.




Figure 3: The state of play so far: the probability loading and comparator subroutines. Note that any intermediate carry ancilla qubits aren’t shown.

Given \hat{C}_k , the next step is to encode the linear part of the payoff function in the amplitude of a payoff ancilla qubit (our analogous qudit protocol only requires a payoff qubit, and not a qudit). This is done

⁷The method of complements, particularly finding nine’s complement of an integer, is a well-known procedure in classical computing which is used, among other things, to find the difference between integers using a machine which can only perform addition (for more on arithmetic on classical computers, see [46]).

using controlled- Y rotations⁸ where, in our case, the asset price qudit register will act as the control. This encoding isn't exact: it is not possible to perfectly encode a linear function in the amplitude of the ancilla using qubit rotations on a machine with finite resources since these rotations only faithfully encode trigonometric functions. However, as outlined in [49], we will shift and scale the rotations such that an accurate approximation of the true payoff function is encoded.

Since the Y rotations will approximately encode the payoff function samples $f(i) \equiv f(s_i) = \max(0, s_i - K)$ (for the sampled asset values $\{s_i\}$ in Eq. 7), and since quantum amplitude estimation (either with or, in our case, without phase estimation) requires that the rotation angles be restricted such that the mapping from the states $|i\rangle_{(n)}$ to the angles be one-to-one [16], we must shift and scale the payoffs $f(i)$ accordingly. First, define

$$\tilde{f}(i) = 2 \frac{f(i) - \min_i(f(i))}{\max_i(f(i)) - \min_i(f(i))} - 1 = 2 \frac{\max(0, s_i - K)}{s_{d^n-1} - K} - 1 \in [-1, 1]. \quad (12)$$

Next, the thing that is measured in practice for the traditional qubit, as well as the current qudit, option pricing scheme (when amplitude estimation without phase estimation [19] is employed), is the state of the payoff ancilla qubit. Per Born's rule [50], the probability of measuring either of the two possible computational basis states is contained in the probability amplitudes of this payoff qubit. Since we're encoding the payoff f by performing rotations on its state, the amplitude of measuring the $|1\rangle$ state will be $\sin^2(\theta)$ (c.f. footnote 8), or, more accurately - since our qudit register is in a superposition of states - a convex combination of the sine-squared function, $\sum_{i=0}^{d^n-1} p_i \sin^2(\theta_i)$, for some θ_i 's. This function, $\sin^2(\theta)$, is one-to-one and onto on the domain $[0, \frac{\pi}{2}]$ and codomain $[0, 1]$. Therefore, we shift and scale $\tilde{f}(i)$ further, to $c\tilde{f}(i) + s$, where $s \equiv \frac{\pi}{4}$ is the domain midpoint and $0 < c \leq \frac{\pi}{4}$ is a small constant such that $c\tilde{f}(i) + s \in [0, 1]$ for all $\tilde{f}(i)$ (more detailed reasons for these choices, as well as an optimal scheme for choosing c , can be found in [49]). As such, $c\tilde{f}(i) + s$ equals either $\frac{\pi}{4} - c$ if $i < k$, or $\frac{2c(s_i - K)}{s_{d^n-1} - K} + \frac{\pi}{4} - c$ if $i \geq k$.

We are now in a position to outline the second part of the European call payoff loading subroutine:

2.2 Construct an operator \hat{L}_f , using controlled- Y rotations with the qudit register and comparator qubit as the controls, to encode f in the amplitude of the payoff qubit as

$$|i\rangle_{(n)} |x\rangle_c |0\rangle_p \xrightarrow{\hat{L}_f} \begin{cases} |i\rangle_{(n)} |0\rangle_c \left(\cos\left(\frac{\pi}{4} - c\right) |0\rangle_p + \sin\left(\frac{\pi}{4} - c\right) |1\rangle_p \right) & \text{if } x = 0, \\ |i\rangle_{(n)} |1\rangle_c \left(\cos\left(\frac{2c(s_i - K)}{s_{d^n-1} - K} + \frac{\pi}{4} - c\right) |0\rangle_p + \sin\left(\frac{2c(s_i - K)}{s_{d^n-1} - K} + \frac{\pi}{4} - c\right) |1\rangle_p \right) & \text{if } x = 1. \end{cases} \quad (13)$$

The circuit implementation of \hat{L}_f is explained in detail in Appendix 6.2. Finally, recall that we cannot perfectly encode f in the amplitude of the payoff ancilla. Indeed, \hat{L}_f does not encode f in the amplitude, but rather $\sin^2(c\tilde{f}(i) + s)$. However, for small enough θ , note that $\sin^2(\theta + \frac{\pi}{4}) = \theta + \frac{1}{2} + \mathcal{O}(\theta^3)$: $\theta + \frac{1}{2}$ provides an adequate approximation⁹. As such, combining Eqs. 11 and 13 and defining $\hat{A} \equiv \hat{L}_f \hat{C}_k \hat{P}$, we get

$$\begin{aligned} |0\rangle_{(n)} |0\rangle_c |0\rangle_p &\xrightarrow{\hat{A}} \sum_{i < k} \sqrt{p_i} |i\rangle_{(n)} |0\rangle_c \left(\cos\left(\frac{\pi}{4} - c\right) |0\rangle_p + \sin\left(\frac{\pi}{4} - c\right) |1\rangle_p \right) \\ &+ \sum_{i \geq k} \sqrt{p_i} |i\rangle_{(n)} |1\rangle_c \left(\cos\left(\frac{2c(s_i - K)}{s_{d^n-1} - K} + \frac{\pi}{4} - c\right) |0\rangle_p + \sin\left(\frac{2c(s_i - K)}{s_{d^n-1} - K} + \frac{\pi}{4} - c\right) |1\rangle_p \right) \\ &= |\Psi_0\rangle |0\rangle_p + |\Psi_1\rangle |1\rangle_p \equiv |\Psi\rangle, \end{aligned} \quad (14)$$

(see Fig. 4). Finally, after having prepared the composite state $|\Psi\rangle$, the probability of measuring the payoff qubit in the $|1\rangle_p$ state is

⁸A simple, 'uncontrolled' rotation of a qubit $|0\rangle$ about the Y axis of its Bloch sphere gives $|0\rangle \rightarrow \cos(\theta)|0\rangle + \sin(\theta)|1\rangle$, where θ is a fixed angle. However, for a *controlled* rotation, with the qudit register in the current context acting as the control, θ is not fixed but is generally a function of the register's state. See [18] for further details.

⁹We could better encode $c\tilde{f}(i) + s$ in the amplitude of the payoff qubit by modifying \hat{L}_f to include higher-order terms. However, this would involve a deeper circuit with far more gates.

$$\begin{aligned}
P_1 &= \langle \Psi_1 | \langle 1 | \Psi \rangle = \langle \Psi_1 | \langle 1 | \Psi_1 \rangle | 1 \rangle_p = \sum_{i < k} p_i \sin^2 \left(\frac{\pi}{4} - c \right) + \sum_{i \geq k} p_i \sin^2 \left(\frac{2c(s_i - K)}{s_{d^n-1} - K} + \frac{\pi}{4} - c \right) \\
&\approx \sum_{i < k} p_i \left(\frac{1}{2} - c \right) + \sum_{i \geq k} p_i \left(\frac{2c(s_i - K)}{s_{d^n-1} - K} + \frac{1}{2} - c \right) = \frac{1}{2} - c + \frac{2c}{s_{d^n-1} - K} \text{E}_P[f(S_T)], \tag{15}
\end{aligned}$$

up to third-order terms. Once we have outlined an efficient way of measuring P_1 , Eq. 15 gives us a way of recovering the expected payoff $\text{E}_P[f(S_T)]$.




Figure 4: Probability loading \hat{P} , comparator \hat{C}_k and payoff encoding \hat{L}_f subroutines.

However, how do we measure P_1 ? Naively repeatedly preparing $|\Psi\rangle$ and then measuring the state of the payoff qubit in an effort to create an estimator of P_1 provides no quantum advantage: indeed, while doing this would eventually give the correct answer, many measurement samples would be needed as this approach is essentially classical Monte Carlo integration which, as mentioned previously, converges as $\mathcal{O}(M^{-1/2})$ (where M is the number of measurement samples taken).

The real advantage offered by quantum computers in pricing complex derivatives stems ultimately from the amplitude amplification inherent to QAE. In a quantum setting, the notion that is equivalent to the number of classical measurement samples taken, is the number of times the operator, or oracle, \hat{A} , is called (we also call this number M). As we shall see in the next section, the error in P_1 converges as $\mathcal{O}(M^{-1})$, a quadratic speed-up over the equivalent classical algorithm.

2.3 Step 3: Estimate the probability amplitude

While a detailed explanation of QAE is given in [16], we here outline only the salient points for brevity.

From Eq. 14, note that \hat{A} , acting on $|0\rangle \equiv |0\rangle_{(n)} |0\rangle_c |0\rangle_p$, creates a state which can be partitioned into a ‘good’ part, $|\Psi_1\rangle |1\rangle$ (where we have dropped the p subscript on the payoff ancilla), whose amplitude $P_1 = \langle \Psi_1 | \langle 1 | \Psi_1 \rangle | 1 \rangle$ we want, and a ‘bad’ mutually orthogonal part, $|\Psi_0\rangle |0\rangle$. The first step in QAE is to use \hat{A} (see Fig. 4) to construct the so-called Grover rotation operator, which is usually denoted by \hat{Q} . This operator is a ‘rotation’ in the sense that repeated applications of \hat{Q} to $|\Psi\rangle$ simply rotate the state in the subspace spanned by the vectors $|\Psi_0\rangle |0\rangle$ and $|\Psi_1\rangle |1\rangle$. In other words, \hat{Q} simply changes the relative amplitudes of $|\Psi_0\rangle |0\rangle$ and $|\Psi_1\rangle |1\rangle$ depending on how many times it is applied.

First, note that $\hat{S}_0 := \mathbb{I}_{d^{n+2}} - 2|0\rangle\langle 0|$ is an operator which flips the sign of only the ground state and does nothing otherwise¹⁰. Given this, define

$$\hat{S}_\Psi := \hat{A} \hat{S}_0 \hat{A}^\dagger = \mathbb{I}_{d^{n+2}} - 2|\Psi\rangle\langle\Psi|, \text{ and } \hat{S}_{\Psi_1} := \mathbb{I}_{d^{n+2}} - 2(\mathbb{I}_{d^{n+1}} \otimes |1\rangle\langle 1|). \tag{16}$$

Here, \hat{S}_Ψ represents a reflection about the state $|\Psi\rangle$, whereas \hat{S}_{Ψ_1} has the action of flipping the sign of states in which the payoff ancilla is in the $|1\rangle$ basis state, and does nothing otherwise¹¹. These operators are used to define \hat{Q}

¹⁰In reality, any ancilla qubits/qudits, from the comparator subroutines outlined in the Appendix, are also included in the system’s composite Hilbert space. These ancillas are just not shown since they are uncomputed in the end.

¹¹As already noted, creating these operators, and hence \hat{Q} , in practice will likely require either decomposing them into whatever fundamental gate set will be available on any potential qudit hardware, or will necessitate feeding these desired operators as input into quantum optimal control algorithms [51, 52]. While we have given a scheme to decompose \hat{A} into fundamental matrices (and, as it turns out, \hat{S}_{Ψ_1} can simply be written as a single qubit gate), it isn’t clear how to decompose \hat{S}_0 . The purpose of this manuscript is not to discuss either general gate compilation nor optimal control in detail but rather simply serves to both outline the steps needed to calculate the theoretical matrices that represent these operators in a qudit context as well as give a tutorial of quantum option pricing.

$$\hat{Q} := -\hat{S}_\Psi \hat{S}_{\Psi_1}. \quad (17)$$

Since \hat{Q} is unitary, its two eigenvalues are pure phases (and are functions of P_1). See [16] for further details of this and the following result but, suffice it to say, this fact can be used to easily compute the action of \hat{Q}^j on $|\Psi\rangle$, for any non-negative integer j

$$\begin{aligned} \hat{Q}^j |\Psi\rangle &= \hat{Q}^j \hat{A} |0\rangle \\ &= \frac{1}{\sqrt{P_1}} \sin(\theta_{P_1}(2j+1)) |\Psi_1\rangle |1\rangle + \frac{1}{\sqrt{1-P_1}} \cos(\theta_{P_1}(2j+1)) |\Psi_0\rangle |0\rangle, \end{aligned} \quad (18)$$

where $\sin^2(\theta_{P_1}) = P_1$ (with $0 \leq \theta_{P_1} \leq \frac{\pi}{2}$). Finding θ_{P_1} hence immediately gives P_1 . Also, note that $2j+1$ is the number of queries to the oracle \hat{A} (once from preparing state $|\Psi\rangle$, and two from each \hat{Q}); recall that the number of calls to this oracle is what is used to gauge the efficiency or convergence of an algorithm. And finally, Eq. 18 roughly gives some intuition behind the quadratic scaling arising from the amplification process: $\sin^2(\theta_{P_1}(2j+1))/P_1 \sim (2j+1)^2$ is approximately the probability of measuring the ‘good’ state for small θ_{P_1} .

Grover’s original search paper [53] used \hat{Q} , in a far simpler context with only one qubit, to enhance the amplitude of the ‘good’ state. This meant that the good state was more likely to be measured in an experiment. The full quantum amplitude amplification and estimation algorithm [16] goes further by employing quantum phase estimation. This process, in general as well as in the current context, involves applying many controlled versions of Grover’s operator \hat{Q} to create various powers of the eigenvalues of \hat{Q} (which, recall, encode P_1). These phase eigenvalues combine together, via phase kickback, in the complex amplitudes of a second qubit register. This register, which acts as the aforementioned control, would simply be a register of *qudits* in a full qudit implementation of QAE [54]. Finally, applying an inverse quantum Fourier transform [55, 18] and measuring the state of the control register gives an integer whose value maps to a very good estimator of θ_{P_1} with a high probability. See [18] for a full explanation of these concepts.

However, phase estimation and the Fourier transform subroutines are themselves the prohibitively computationally expensive bottlenecks in full amplitude estimation algorithms carried out on current, noisy qubit or qudit NISQ-era hardware. As such, various approaches which maintain the quadratic speed-up while forgoing both the phase estimation and Fourier transform subroutines have been proposed [19, 56, 57]). Arguably the simplest method among these is that outlined in [19], where maximum likelihood estimation is carried out on measurement data collected after various numbers of ordinary, un-controlled \hat{Q} operations are applied to the quantum system. This is the technique we’ll detail and simulate below.

We can now state the final step in the option pricing scheme

- 3 Using \hat{A} , create the Grover operator \hat{Q} as per Eq. 17. Then, apply it j times to the prepared state $|\Psi\rangle = \hat{A} |0\rangle$ and measure the state of the payoff qubit (see Fig. 5). Repeat this single shot measurement a number of times and for various j ’s, according to the scheme to be outlined shortly. Finally, use the measurement data to construct a maximum likelihood (ML) function from which the ML estimate of θ_{P_1} can be gleaned.




Figure 5: Maximum-likelihood amplitude estimation algorithm using Grover’s operator \hat{Q} .

Maximum likelihood estimation (MLE) is an important technique in modern statistics. Here we give a brief summary of MLE and tailor the ideas to the current option pricing context (see [58] for an illuminating

tutorial on MLE). In general, suppose that one has a statistical model for a random variable and that the model's distribution itself depends on one or more unknown parameters. Furthermore, assume that one can draw samples from the distribution as needed. In this case, MLE provides a technique to deduce the unknown parameters for which, under the chosen mathematical model, the observed sample data is 'most likely' to have been observed. This setup can be easily applied to the current context: Eq. 18 represents the probability model for measurements of the random variable (namely the payoff qubit's state) and the unknown parameter we are attempting to deduce is θ_{P_1} .

Following the optimal scheme outlined in [19], define $m_0 = 0$ and $m_\ell = 2^{\ell-1}$, for $\ell = 1, \dots, T$ for some cutoff T ¹². The m_ℓ 's will correspond to the number of times \hat{Q} is applied to the state $\hat{A}|0\rangle$, i.e. $\hat{Q}^{m_\ell}\hat{A}|0\rangle$, in different experimental runs. Furthermore, take N to be the fixed number of single-shot measurements that will be taken of the payoff qubit's state for each run m_ℓ .

Any single measurement of the payoff qubit's state is a Bernoulli trial with $\sin^2(\theta_{P_1}(2m_\ell + 1))$ being the probability of measuring the 'good' state $|1\rangle$ and $\cos^2(\theta_{P_1}(2m_\ell + 1))$ the probability of measuring the 'bad' state $|0\rangle$. After performing this single measurement, set $s_\ell = 1$ if we measure the good state and $s_\ell = 0$ if we measure the bad one. In this case, the likelihood function is

$$\mathcal{L}_\ell(\theta_{P_1}|s_\ell) = [\sin^2(\theta_{P_1}(2m_\ell + 1))]^{s_\ell} [\cos^2(\theta_{P_1}(2m_\ell + 1))]^{1-s_\ell}. \quad (19)$$

If we repeat this measurement N times for a fixed m_ℓ (so that the events are independent and identically distributed) and let $s_\ell \in \{0, \dots, N\}$ denote the number of times the 'good' state is measured, this experiment is described by a binomial distribution with the likelihood function

$$\mathcal{L}_\ell(\theta_{P_1}|s_\ell) = \binom{N}{s_\ell} [\sin^2(\theta_{P_1}(2m_\ell + 1))]^{s_\ell} [\cos^2(\theta_{P_1}(2m_\ell + 1))]^{N-s_\ell}. \quad (20)$$

Next, even though each s_ℓ is not identically distributed for different ℓ 's, the samples are generated independently from the same joint distribution. As such, after repeating the single shot measurement N times for each m_ℓ and recording each s_ℓ , we can arrive at a single likelihood function for the unknown parameter θ_{P_1} by simply multiplying the individual \mathcal{L}_ℓ 's together

$$\mathcal{L}(\theta_{P_1}|\{s_\ell\}) = \prod_{\ell=0}^T \mathcal{L}_\ell(\theta_{P_1}|s_\ell). \quad (21)$$

In a rough sense, \mathcal{L} represents the probability of observing the set of random variables $\{s_\ell\}$, as a function of θ_{P_1} . The last step of MLE entails 'maximising' this probability by finding the value of the parameter which maximises \mathcal{L} . This estimate of θ_{P_1} , which we denote $\hat{\theta}_{P_1}$, is hence

$$\hat{\theta}_{P_1} = \arg_{0 \leq \theta_{P_1} \leq \frac{\pi}{2}} \max \mathcal{L}(\theta_{P_1}|\{s_\ell\}). \quad (22)$$

$\hat{\theta}_{P_1}$ is the value for which the data samples $\{s_\ell\}$ are 'most likely' to have been observed and is generally found in practice using numerical solvers.

We finally have all of the ingredients that are needed to price a European call option on a quantum computer with a qudit register, since the estimate $\hat{\theta}_{P_1}$ is mapped to an estimate for P_1 , which in turn can be easily classically mapped to the expected payoff via Eq. 15. Note that the total number of oracle queries (i.e. the number of times \hat{A} is applied), namely M , in this experiment is $M = \sum_{\ell=0}^T N(2m_\ell + 1)$. Although proving so is beyond the scope of the current text, it can be shown that for the amplitude estimation without phase estimation scheme in both a qubit and qudit context, the estimation error $\hat{\epsilon}_{P_1} = \sqrt{\mathbb{E}[(\theta_{P_1} - \hat{\theta}_{P_1})^2]}$ asymptotically scales as $\mathcal{O}(M^{-1})$ (see [19, 56] for details). As such, the full quantum algorithm here presents a theoretical quadratic speedup compared with typical derivative pricing calculations employing classical Monte Carlo techniques. However, as mentioned before, path-independent options, such as the European call option presented here, are simple enough to unlikely benefit much from quantum derivative pricing. On the other hand, path-dependent financial options, such as the barrier options briefly mentioned at the end of section 1 as well as the real-world financial options discussed in [11], are likely to benefit the most from


¹²Note that one could choose a different equation for the m_ℓ 's. However, this optimal exponential sequence asymptotically maintains the quadratic speed-up [19].

advances in quantum computing [11]. While we close off this manuscript by simulating the entire stack for a relatively simple path-independent option, its purpose wasn't to delve deeply into more complex financial derivatives but rather to both discuss the basics of derivative pricing for those unfamiliar with the subject as well as to outline the quantum computing algorithm behind a simple use case, in enough mathematical detail, so as to easily allow the interested reader to disassemble and then 'stitch together' the constituent subroutines of the algorithm when applying quantum computing to more complex derivatives.


3 European call simulations

We finally simulate the entire European call option pricing scheme using a custom PYTHON package, from sampling the asset's theoretical distribution and constructing the appropriate circuit subroutines, to performing a maximum likelihood estimation on the simulated measurement results. The scheme presented in section 2 above is similar to what a future implementation of financial derivative pricing on a qudit-based quantum computer could potentially entail.

Fig. 6a gives the distribution of the random price of a single asset arising from the Black-Scholes-Merton model, with the chosen parameters in the caption. Since this distribution has relatively flat tails, truncating and discretising the distribution doesn't introduce significant errors. This may not be the case for more volatile or some realistic, market-implied distributions. However, as quantum computers improve over time, such cases could be studied by increasing either the logical qubit/qudit count or increasing the effective Hilbert space dimension of existing qudits. This would allow one to produce a more fine-grained discrete approximation of the true distribution. The latter, namely exploiting the theoretically larger state space of qudit-based hardware, isn't possible when working strictly with qubits and hence presents one possible boon of qudits, all else being equal. However, a trade-off between the experimental time and effort required to increase the dimension of a qudit and the logarithmic scaling of the information that is able to be stored in a qudit, would need to be considered.



(a) BSM model asset price distribution and sampled points, assuming an 8-dimensional asset price register.




(b) Corresponding expected payoff of European call option vs. total Hilbert space dimension.

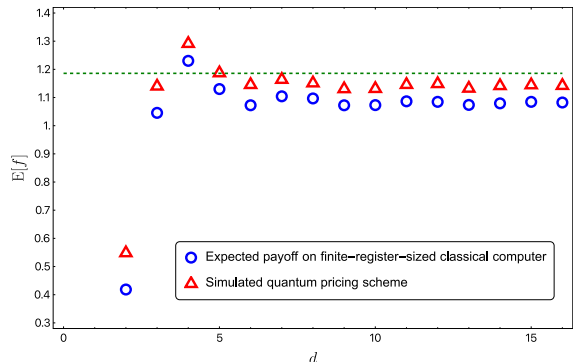
Figure 6: Simulation of the expected payoff of a European call option. (a) The encoded distribution, namely Eq .4, with $S_0 = 2.0, t = 365/365, \alpha = 0.07$ and $\sigma = 0.3$. (b) The analytic expected payoff (dashed line - assuming a strike price K of 1.7), the expected payoff given a classical register of size d (blue circles), and the corresponding quantum simulation for a single asset qudit register ($n = 1$) with a Hilbert space dimension of d (red triangles). $N = 100$ shots are taken for each m_ℓ , and $T = 7$ (giving $M = 26200$).

Fig. 6b gives the corresponding expected payoff as a function of the size of the asset price's qudit register, or equivalently, the number of samples of the underlying probability distribution taken. The discrepancy between both the classical and quantum simulation data points with the analytic payoff is an inevitable result of truncating the domain of the distribution, a process which discards the large but relatively unlikely asset values. For a fixed number of qubits/qudits, increasing the size of the truncated region would capture more extreme prices but with the drawback of producing a more coarse-grained discretisation. However, more

important is the fact that the qudit algorithm’s expected payoff data points (the red triangles) are close to the payoff calculated on a finite-resource classical computer (blue circles): the former even approaches the latter as the size of the Hilbert space d increases. The classical and quantum payoffs are close to one another after a relatively small increase in Hilbert space size, suggesting that a qudit register of modest size may be sufficient to match the payoff possible on a finite-resource classical device, within error. The remaining small discrepancy between the quantum and classical simulation data points is due to a combination of the constant error arising from encoding the linear payoff function using trigonometric rotations [49] as well as the $\mathcal{O}(M^{-1})$ error arising from QAE itself. This discrepancy could be decreased further by including higher-order terms in the payoff loading subroutine \hat{L}_f (which would result in a deeper circuit), by perhaps increasing the number of shots N for each unique circuit ℓ , as well as increasing the number of unique circuits ℓ that are measured too. It should be noted, however, that these aforementioned discrepancies arise from the theoretical derivative pricing algorithm itself, to say nothing about the physics- and device-level errors of current NISQ-era hardware. For example, a future quantum compiler (a good explanation of the technical details of which, in the case of IBM’s widely-used qubit compiler, can be found in [59]) that would compile the required unitary operations (\hat{A} and \hat{Q} in our case) into hardware-native operations, will be subject to computational and coherent errors. Furthermore, both physical qubit and qudit systems have inherent sources of noise on a fundamental physics level, the study of which is an active area of research [60].



(a) BSM model, with a higher volatility, and a 10-dimensional asset price register’s sample points.



(b) Corresponding expected payoff of option vs. total Hilbert space dimension.

Figure 7: Simulation of a European call option with a more volatile underlying asset. (a) The same analytic distribution and parameters as Fig. 6, except now with $S_0 = 3.0$ and $\sigma = 0.5$. (b) The corresponding simulated payoff data (with a strike price now of 2.2), with the dashed line indicating the analytic value of the expected payoff.

Fig. 7 gives simulated results, also for a European call option, but one based on a more volatile underlying asset. The larger volatility increases the statistical spread of the distribution which in turn increases the probability mass of extreme asset prices. The simulation in Fig. 7b, as was the case in Fig. 6b, truncates the asset price to within three standard deviations of the mean. The larger spread of this distribution means that extreme asset prices are not as well accounted for in the probability loading subroutine \hat{P} . As a result, there is a larger discrepancy in Fig. 7b between the analytic payoff and both the classical finite-resource and quantum payoffs, as expected. However, as before, the classical finite-resource and quantum payoff simulations again match, up to the expected theoretical error.

4 Conclusion

We have presented a tutorial on the problem of European option pricing and the algorithm steps one might go through when realistically pricing such an option on a future quantum computer with qudit information carriers. While the primer on derivatives in section 1 is admittedly brief and doesn’t do the enormous field of financial mathematics much justice, we have endeavoured to both provide the necessary basic details for those unfamiliar with the theory behind financial options as well as outline the BSM model, which is

the starting point both in the historical development of financial mathematics as well as a natural point from which the interested reader may begin exploring the topic. It must again be emphasised that the simple European options discussed in section 1 can be efficiently priced classically; more complex, realistic, path-dependent financial options are those which stand to benefit the most from quantum computers (see [11] for an excellent explanation of real-world derivatives as well as resource estimates for their pricing on a qubit-based quantum computer). However, given that quantum computing hardware is still in its relative infancy, comprised of relatively small numbers of noisy qubits or qudits, studying the simplest financial use cases will remain the benchmark for some time.

Armed with this basic understanding, section 2 described the probability loading, payoff computation and amplitude estimation routines in quantum derivative pricing. While loading a log-normal distribution - of which the European call's distribution is an example - into a register of qudits can be done efficiently, detailed comments on the ease with which more general distributions could be encoded can only be made in light of knowledge of the native operations a potential qudit-based quantum computer will have access to. For example, for microwave cavities coupled to superconducting qubits ([40, 39]), gate sets comprised of SNAP plus displacement operations [44, 42], or echo-conditional displacement gates plus single qubit rotations [61], provide universal control over their system's Hilbert space. However, these gates differ from those of, say, traditional 2D superconducting qubit chip systems. Regardless, given a universal gate set, it is theoretically possible to encode any arbitrary probability distribution as long as the system is coherent enough to accommodate a deep enough circuit and/or has enough information carriers. All else being equal however, a register of higher-dimensional logical qudits can encode an analytic or market-implied distribution better than a register of logical qubits consisting of an equal number of information carriers. This is important since, as discussed, better distribution encoding is essential in order to better capture outlying events, which tend to have an outsized effect on market dynamics.

In section 2.2 and Appendix 6.1, we discussed a couple of natural qudit analogues of the ordinary qubit comparator: one choice uses fewer gates but requires more ancillas, whereas the other uses only one ancilla but requires a far deeper circuit. Experimental hardware considerations, such as device coherence times and qudit/qubit connectivity, would best dictate which to use in practice. Furthermore, while the ancilla carry qubits in the comparator subroutine as well as the comparator and payoff qubits themselves were indeed qubits, they could very well have been qudits. Ancilla qudits, however, would not provide any advantage over qubits in the option pricing scheme discussed here since these ancillas are merely required to carry one bit of information. However, given that many proposed cavity QED schemes use traditional superconducting qubits as non-linear elements, it is not unreasonable to suppose that a future qudit-based quantum computer may in fact be hybrid in nature, employing both qubits and qudits.

Using this comparator, sections 2.2, and 2.3 detail how the European call payoff is both encoded in the amplitude of a qubit and the expected payoff subsequently recovered using amplitude estimation without phase estimation. The latter, which uses statistical maximum likelihood estimation to forgo phase estimation, will be required in the current NISQ era of quantum computing as a full phase estimation subroutine [17, 18] increases both the depth and width of the qudit circuit beyond near-term capabilities. The scheme that we've discussed is constructed in such a manner so as to reduce practical resource requirements while maintaining the speedup quantum option pricing offers over the usual method of classically pricing an option using Monte Carlo integration.

The subroutines outlined in section 2 also provide many of the ingredients required to construct the qudit algorithms needed to price other types of options. With the addition of further qudit registers to hold multiple assets' prices as well as subroutines which sum and average the integers encoded in these registers, it is possible to use a qudit-based quantum computer to price, for example, the Asian options and barrier options discussed in section 1. Both of these are path-dependent options and contain structures similar to those appearing in other modern, exotic options [62]. Section 3 finally provides a couple of simulations of the entire pricing scheme for the European call option, with different parameters. While the focus of this manuscript has been far more on the journey towards understanding the detailed intersection between quantum computing and derivative pricing and not this simulation 'destination', it is important to note that a logical qudit register of only modest size is capable of pricing a European call option with comparable accuracy to a finite-resource classical computer of similar register size. In the current NISQ era, the logarithmic scaling potential of logical qudits still provide a worthwhile advantage over qubits. Furthermore, qudits with correspondingly modest Hilbert space dimensions of the order of 1 – 10 are currently being actively studied

in existing systems .

Finally, it is worth mentioning that qudits and higher-dimensional Hilbert spaces in general offer other potential advantages. For example, it is possible to encode information in robust continuous-variable states, such as cat states [63]. This is an example of bosonic encoding [64] whereby the full Hilbert space of electromagnetic signals (in microwave or SRF cavities [20], for example) are exploited in quantum information encoding and processing, in contrast to constraining oneself to the first two states for each qubit in a superconducting 2D chip. Bosonic encoding schemes allow for novel quantum error mitigation schemes and correction protocols [65, 66], an important step on the road to creating a fault-tolerant quantum computer beyond the NISQ era. And beyond quantum computing, into the broader quantum technology sphere, quantum communication channels employing qudits possess larger information bandwidths as well as potentially greater robustness against quantum state cloning attacks [67].

5 Acknowledgements

The author thanks Andy C. Y. Li for his helpful comments. This material is based upon work supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science, Research Centers, Superconducting Quantum Materials and Systems Center (SQMS) under contract number DE-AC02-07CH11359.

References

- [1] G. Ammendola, *Devil Take the Hindmost: A History of Financial Speculation*. Taylor & Francis, 2000.
- [2] F. Black and M. Scholes, “The pricing of options and corporate liabilities,” *Journal of Political Economy*, vol. 81, no. 3, pp. 637–654, 1973.
- [3] J. Baz and G. Chacko, *Financial derivatives: pricing, applications, and mathematics*. Cambridge University Press, 2004.
- [4] J. C. Hull, *Options futures and other derivatives*. Pearson Education India, 2003.
- [5] A. D. Fokker, “Die mittlere energie rotierender elektrischer Dipole im Strahlungsfeld,” *Annalen der Physik*, vol. 348, no. 5, pp. 810–820, 1914.
- [6] M. Planck, “Über einen Satz der statistischen Dynamik und seine Erweiterung in der Quantentheorie,” *Sitzungsberichte der Preussischen Akademie der Wissenschaften zu Berlin*, vol. 24, pp. 324—341, 1917.
- [7] H. Risken, *Fokker-Planck Equation*, pp. 63–95. Berlin, Heidelberg: Springer Berlin Heidelberg, 1996.
- [8] A. Damodaran, “Investor irrationality.” https://pages.stern.nyu.edu/~adamodar/New_Home_Page/invfables/investorirrationality.htm. Accessed: 2022-12-25.
- [9] G. G. Roussas, *An introduction to measure-theoretic probability*. Academic Press, 2014.
- [10] P. Rebentrost, B. Gupt, and T. R. Bromley, “Quantum computational finance: Monte carlo pricing of financial derivatives,” *Phys. Rev. A*, vol. 98, p. 022321, Aug 2018.
- [11] S. Chakrabarti, R. Krishnakumar, G. Mazzola, N. Stamatopoulos, S. Woerner, and W. J. Zeng, “A threshold for quantum advantage in derivative pricing,” *Quantum*, vol. 5, p. 463, 2021.
- [12] R. E. Caffisch, “Monte carlo and quasi-monte carlo methods,” *Acta numerica*, vol. 7, pp. 1–49, 1998.
- [13] Y. Wang, Z. Hu, B. C. Sanders, and S. Kais, “Qudits and high-dimensional quantum computing,” *Frontiers in Physics*, vol. 8, p. 589504, 2020.
- [14] Y. Chi, J. Huang, Z. Zhang, J. Mao, Z. Zhou, X. Chen, C. Zhai, J. Bao, T. Dai, H. Yuan, *et al.*, “A programmable qudit-based quantum processor,” *Nature communications*, vol. 13, no. 1, p. 1166, 2022.

- [15] A. Montanaro, “Quantum speedup of monte carlo methods,” *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 471, no. 2181, p. 20150301, 2015.
- [16] G. Brassard, P. Hoyer, M. Mosca, and A. Tapp, “Quantum amplitude amplification and estimation,” *Contemporary Mathematics*, vol. 305, pp. 53–74, 2002.
- [17] A. Y. Kitaev, “Quantum measurements and the abelian stabilizer problem,” *Electron. Colloquium Comput. Complex.*, vol. TR96, 1995.
- [18] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.
- [19] Y. Suzuki, S. Uno, R. Raymond, T. Tanaka, T. Onodera, and N. Yamamoto, “Amplitude estimation without phase estimation,” *Quantum Information Processing*, vol. 19, pp. 1–17, 2020.
- [20] M. S. Alam, S. Belomestnykh, N. Bornman, G. Cancelo, Y.-C. Chao, M. Checchin, V. S. Dinh, A. Grassellino, E. J. Gustafson, R. Harnik, *et al.*, “Quantum computing hardware for HEP algorithms and sensing,” *arXiv preprint arXiv:2204.08605*, 2022.
- [21] N. N. Taleb, *The Black Swan: The Impact of the Highly Improbable*, vol. 2. Random House, 2007.
- [22] M. Taboga, “Householder matrix,” 2021. <https://www.statlect.com/matrix-algebra/Householder-matrix>. Accessed: 2022-12-08.
- [23] M. Plesch and Č. Brukner, “Quantum-state preparation with universal gate decompositions,” *Phys. Rev. A*, vol. 83, no. 3, p. 032302, 2011.
- [24] M. Möttönen, J. J. Vartiainen, V. Bergholm, and M. M. Salomaa, “Transformation of quantum states using uniformly controlled rotations,” *Quantum Info. Comput.*, vol. 5, p. 467–473, sep 2005.
- [25] J. J. Vartiainen, M. Möttönen, and M. M. Salomaa, “Efficient decomposition of quantum gates,” *Physical Review Letters*, vol. 92, no. 17, p. 177902, 2004.
- [26] M. Möttönen, J. J. Vartiainen, V. Bergholm, and M. M. Salomaa, “Quantum circuits for general multiqubit gates,” *Phys. Rev. Lett.*, vol. 93, no. 13, p. 130502, 2004.
- [27] J. Preskill, “Quantum computing in the NISQ era and beyond,” *Quantum*, vol. 2, p. 79, 2018.
- [28] M. Luo and X. Wang, “Universal quantum computation with qudits,” *Science China Physics, Mechanics & Astronomy*, vol. 57, pp. 1712–1717, 2014.
- [29] L. Grover and T. Rudolph, “Creating superpositions that correspond to efficiently integrable probability distributions,” *arXiv preprint quant-ph/0208112*, 2002.
- [30] P. B. Matthew F. Dixon, Igor Halperin, *Machine Learning in Finance: From Theory to Practice*. Springer Cham, 2020.
- [31] S. Liu, C. W. Oosterlee, and S. M. Bohte, “Pricing options and computing implied volatilities using neural networks,” *Risks*, vol. 7, no. 1, p. 16, 2019.
- [32] B. Horvath, A. Muguruza, and M. Tomas, “Deep learning volatility: a deep neural network perspective on pricing and calibration in (rough) volatility models,” *Quantitative Finance*, vol. 21, no. 1, pp. 11–27, 2021.
- [33] C. Zoufal, A. Lucchi, and S. Woerner, “Quantum generative adversarial networks for learning and loading random distributions,” *npj Quantum Information*, vol. 5, no. 1, p. 103, 2019.
- [34] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” in *Advances in Neural Information Processing Systems* (Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, eds.), vol. 27, Curran Associates, Inc., 2014.

- [35] N. Stamatopoulos, D. J. Egger, Y. Sun, C. Zoufal, R. Iten, N. Shen, and S. Woerner, “Option Pricing using Quantum Computers,” *Quantum*, vol. 4, p. 291, July 2020.
- [36] D. Kielpinski, C. Monroe, and D. J. Wineland, “Architecture for a large-scale ion-trap quantum computer,” *Nature*, vol. 417, no. 6890, pp. 709–711, 2002.
- [37] E. Knill, R. Laflamme, and G. J. Milburn, “A scheme for efficient quantum computation with linear optics,” *Nature*, vol. 409, no. 6816, pp. 46–52, 2001.
- [38] P. Kok, W. J. Munro, K. Nemoto, T. C. Ralph, J. P. Dowling, and G. J. Milburn, “Linear optical quantum computing with photonic qubits,” *Reviews of Modern Physics*, vol. 79, no. 1, p. 135, 2007.
- [39] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff, “Circuit quantum electrodynamics,” *Reviews of Modern Physics*, vol. 93, p. 025005, May 2021.
- [40] S. Haroche, M. Brune, and J. M. Raimond, “From cavity to circuit quantum electrodynamics,” *Nature Physics*, vol. 16, no. 3, pp. 243–246, 2020.
- [41] A. Romanenko, R. Pilipenko, S. Zorzetti, D. Frolov, M. Awida, S. Belomestnykh, S. Posen, and A. Grassellino, “Three-dimensional superconducting resonators at $T \leq 20$ mk with photon lifetimes up to $\tau = 2$ s,” *Phys. Rev. Appl.*, vol. 13, p. 034032, Mar 2020.
- [42] T. Fosel, S. Krastanov, F. Marquardt, and L. Jiang, “Efficient cavity control with SNAP gates,” *Bulletin of the American Physical Society*, 2020.
- [43] Google Quantum AI, “Suppressing quantum errors by scaling a surface code logical qubit,” *Nature*, vol. 614, no. 7949, pp. 676–681, 2023.
- [44] S. Krastanov, V. V. Albert, C. Shen, C.-L. Zou, R. W. Heeres, B. Vlastakis, R. J. Schoelkopf, and L. Jiang, “Universal control of an oscillator with dispersive coupling to a qubit,” *Phys. Rev. A*, vol. 92, p. 040303, Oct 2015.
- [45] S. A. Cuccaro, T. G. Draper, S. A. Kutin, and D. P. Moulton, “A new quantum ripple-carry addition circuit,” *arXiv preprint quant-ph/0410184*, 2004.
- [46] I. Koren, *Computer arithmetic algorithms*. AK Peters/CRC Press, 2018.
- [47] F. Orts, G. Ortega, E. Combarro, and E. Garzón, “A review on reversible quantum adders,” *Journal of Network and Computer Applications*, vol. 170, p. 102810, 2020.
- [48] T. G. Draper, “Addition on a quantum computer,” *arXiv preprint quant-ph/0008033*, 2000.
- [49] S. Woerner and D. J. Egger, “Quantum risk analysis,” *npj Quantum Information*, vol. 5, no. 1, p. 15, 2019.
- [50] M. Born, “Zur quantenmechanik der Stoßvorgänge,” *Zeitschrift für Physik*, vol. 37, pp. 863–867, 1926.
- [51] J. Werschnik and E. Gross, “Quantum Optimal Control Theory,” *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 40, no. 18, p. R175, 2007.
- [52] C. P. Koch, U. Boscain, T. Calarco, G. Dirr, S. Filipp, S. J. Glaser, R. Kosloff, S. Montangero, T. Schulte-Herbrüggen, D. Sugny, *et al.*, “Quantum optimal control in quantum technologies. strategic report on current status, visions and goals for research in Europe,” *EPJ Quantum Technology*, vol. 9, no. 1, p. 19, 2022.
- [53] L. K. Grover, “A fast quantum mechanical algorithm for database search,” in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC ’96, (New York, NY, USA), p. 212–219, Association for Computing Machinery, 1996.
- [54] Y. Cao, S.-G. Peng, C. Zheng, and G.-L. Long, “Quantum Fourier Transform and Phase Estimation in Qudit System,” *Communications in Theoretical Physics*, vol. 55, no. 5, p. 790, 2011.

- [55] D. Coppersmith, “An approximate Fourier transform useful in quantum factoring,” in *IBM: Research Report RC19642*, 1994.
- [56] S. Aaronson and P. Rall, “Quantum Approximate Counting, Simplified,” in *Symposium on Simplicity in Algorithms*, pp. 24–32, SIAM, 2020.
- [57] D. Grinko, J. Gacon, C. Zoufal, and S. Woerner, “Iterative quantum amplitude estimation,” *npj Quantum Information*, vol. 7, no. 1, p. 52, 2021.
- [58] I. J. Myung, “Tutorial on maximum likelihood estimation,” *Journal of Mathematical Psychology*, vol. 47, no. 1, pp. 90–100, 2003.
- [59] A. W. Cross, L. S. Bishop, S. Sheldon, P. D. Nation, and J. M. Gambetta, “Validating quantum computers using randomized model circuits,” *Phys. Rev. A*, vol. 100, p. 032328, Sep 2019.
- [60] C. Müller, J. H. Cole, and J. Lisenfeld, “Towards understanding two-level-systems in amorphous solids: insights from quantum circuits,” *Reports on Progress in Physics*, vol. 82, p. 124501, Oct 2019.
- [61] A. Eickbusch, V. Sivak, A. Z. Ding, S. S. Elder, S. R. Jha, J. Venkatraman, B. Royer, S. Girvin, R. J. Schoelkopf, and M. H. Devoret, “Fast universal control of an oscillator with weak dispersive coupling to a qubit,” *Nature Physics*, vol. 18, no. 12, pp. 1464–1469, 2022.
- [62] C. Dodds, “What Are Exotic Options? 11 Types of Exotic Options,” 2021. <https://www.sofi.com/learn/content/exotic-options/>. Accessed: 2023-10-23.
- [63] C. Gerry and P. Knight, “Quantum superpositions and Schrödinger cat states in quantum optics,” *American Journal of Physics*, vol. 65, no. 10, pp. 964–974, 1997.
- [64] W.-L. Ma, S. Puri, R. J. Schoelkopf, M. H. Devoret, S. M. Girvin, and L. Jiang, “Quantum control of bosonic modes with superconducting circuits,” *Science Bulletin*, vol. 66, no. 17, pp. 1789–1805, 2021.
- [65] L. Li, C.-L. Zou, V. V. Albert, S. Muralidharan, S. Girvin, and L. Jiang, “Cat Codes with Optimal Decoherence Suppression for a Lossy Bosonic Channel,” *Physical Review Letters*, vol. 119, no. 3, p. 030502, 2017.
- [66] W. Cai, Y. Ma, W. Wang, C.-L. Zou, and L. Sun, “Bosonic quantum error correction codes in superconducting quantum circuits,” *Fundamental Research*, vol. 1, no. 1, pp. 50–67, 2021.
- [67] D. Cozzolino, B. Da Lio, D. Bacco, and L. K. Oxenløwe, “High-dimensional Quantum Communication: Benefits, Progress, and Future Challenges,” *Advanced Quantum Technologies*, vol. 2, no. 12, p. 1900038, 2019.
- [68] C. C. Pinter, *A Book of Abstract Algebra*. Dover Publications, 2010.

6 Appendix

6.1 Comparator subroutine

Given two classical integers k and i , expressed as strings of n digits in base- d (e.g. $k = k_{n-1}k_{n-2} \cdots k_1k_0$), the method of complements - used to compute the difference $i - k$ using addition - is as follows:

- For each digit $k[j] \equiv k_j$, compute the digit’s $d - 1$ complement and write the result as a string, namely: $(d - 1 - k_{n-1})(d - 1 - k_{n-2}) \cdots (d - 1 - k_0)$
- Add 1 to this, calling the result k^c
- Add i and k^c , digit by digit. The extra leftmost digit of this sum, at index n , is either 0 or 1: if it’s 0, then i is less than k ; if it is 1, then i is bigger than or equal to k

For example, in base 5, with four digits ($n = 4$), suppose $i = 382_{10} = 3012_5$ and $k = 329_{10} = 2304_5$. In this case, $k^c = 2141_5$ and hence $i + k^c = 10203_5$. Since the extra digit, at index position 4, is 1, this indicates that $i \geq k$ (and the 4 least significant digits, namely 0203_5 , gives the difference $i - k = 53_{10}$ itself).

This is simple enough in a purely classical case, but we're comparing the qudit quantum state $|i\rangle_{(n)}$ with the classical variable k . We hence use a qudit version of the quantum ripple-carry full adder circuit for qubits [45], tailored to k , to perform this sum. In what follows, we assume that our quantum system, which is comprised of qudits along with some qubits, is equipped with any arbitrary single qudit/qubit gate along with two-qudit/qubit controlled gates. Any unitary matrix operation on the overall composite Hilbert space of the system can then be decomposed into these fundamental operations [13]. In this manuscript, the elementary gates used to construct circuits are chosen simply to draw clean analogies with existing gates in the qubit literature as well as to be able to calculate numerical matrix representations of the required unitary operations. As mentioned in section 2.1.1, in practice, the matrix operation for the current option pricing algorithm would need to be decomposed into a fundamental gate set offered by the particular 'qudit' hardware platform that is available; this is, as yet, unknown.

There are at least two potential approaches to construct circuits for \hat{C}_k , depending on how many ancilla carry qubits we wish to include. One approach assumes that we have $\mathcal{O}(n)$ carry qubits available and results in a circuit of depth $\mathcal{O}(n)$ (counted with respect to single-control or double-control qudit/qubit gates). This approach is the qudit analogue of the comparator put to use in [35]. See Fig. 8 and its caption for the circuit and detailed explanation thereof (note that any ancilla carry qubits are not shown in the states in the main text of this manuscript).

A second approach to constructing \hat{C}_k instead only requires $\mathcal{O}(1)$ ancilla carry qubits. Such a circuit subroutine may be preferable for quantum hardware with limited resources yet long coherence times. Figs. 9a and 9b outline the $n = 2$ and $n = 3$ cases, respectively; the subroutine for general n can be found recursively. However, having fewer ancilla qubits comes at the expense of requiring a much deeper circuit: at least $\mathcal{O}(n^2)$, counted with respect to the controlled- X gates with an arbitrary number of control qudits¹³. Recall that it is indeed possible to employ qudits, rather than qubits, as the 'carry' ancillas.

6.2 Payoff encoding subroutine

To construct \hat{L}_f , first recall that the base- d digit representation of i , is

$$i = i_0 + d^1 i_1 + \dots + d^{n-2} i_{n-2} + d^{n-1} i_{n-1}, \quad (23)$$

so, using Eqs. 7 and 12, $c\tilde{f}(i) + s$ can be written as

$$c\tilde{f}(i) + s = \frac{2c}{s_{d^{n-1}} - K} \times \max \left(0, \omega i_0 + \omega d^1 i_1 + \dots + \omega d^{n-1} i_{n-1} + S_{\min} + \frac{1}{2} \omega - K \right) - c + \frac{\pi}{4}.$$

The choice between the max function arguments is controlled by the comparator qubit. This relatively straightforward circuit is set out in Fig. 10. Since the mapping from $(\mathbb{R}, +)$ (the group of all real numbers, equipped with ordinary addition) to (\mathbb{R}_Y, \cdot) (the group of all R_Y rotations, equipped with matrix multiplication) is a group homomorphism [68], the contributions from each of the individual Y rotations combine such that $R_Y(c\tilde{f}(i) + s)$ is ultimately the rotation applied to the payoff qubit.

¹³Since it's unclear how to efficiently decompose multi-controlled qudit operations - as in Fig. 9 - into single- or double-control gates - as in Fig. 8 - a fair comparison between these two approaches' scaling isn't practicable yet.


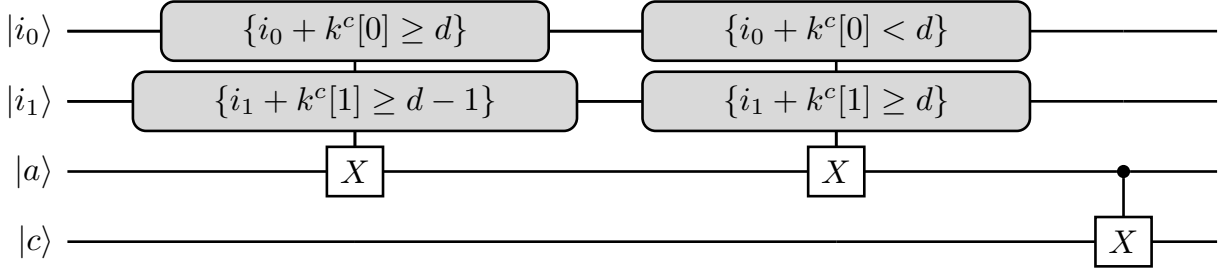



Figure 8: Circuit for qudit quantum comparator with $\mathcal{O}(n)$ available ancilla carry qubits. The information carriers denoted with i 's are the d -dimensional qudits, the a 's are the 2-dimensional ancilla qubits (which we initialise to $|0\rangle$ at the start of the experiment), and qubit $|c\rangle$ is the comparator qubit. Each greyed 'control bubble' denotes a set of i_j values (for $j \in \{0, \dots, n-1\}$) for which the target operation - here only X gates - is applied to the corresponding ancilla qubit $|a_j\rangle$ (since k , which is uniquely defined by the strike price K , is fixed and known, this circuit is unique and customised for each k). For example, the leftmost gate inside the very first dashed box flips the state of $|a_1\rangle$ to $|1\rangle$ only for the basis state $|i_1\rangle$ for which the digit sum of $|i_1\rangle$ and $k^c[1]$ equals $d-1$. In particular, this gate accounts for the case in which the classical addition of i_1 and $k^c[1]$ would(not) produce an output carry bit, given a $1(0)$ input carry bit, since the gate itself is also conditioned on the state of the input carry qubit, namely $|a_0\rangle$. The rightmost gate inside the same dashed box accounts for the cases in which $i_1 + k^c[1]$ is large enough to produce an output carry bit of 1, regardless of the value of the input carry bit (if any control set in this circuit is empty, the entire gate falls away). The presence of an output carry bit is hence encoded in $|a_1\rangle$, which is in turn fed, as an input carry bit, into the gates operating non-trivially on $|i_2\rangle$ and $|a_2\rangle$. This process is repeated in turn for each qudit until the final ancilla, $|a_{n-1}\rangle$, is set to $|0\rangle$ if $i - k$ (equivalently, $i + k^c$) would not produce a carry bit (in which case $i < k$), or $|1\rangle$ if it would. Finally, in order to be able to 're-use' the a qubits in later iterations of C_k , the state of $|a_{n-1}\rangle$ is copied to the comparator qubit $|c\rangle$ (also initialised to $|0\rangle$) using one final controlled- X gate, and all of the gates (except the last CX) are uncomputed by repeating all of the gates, in reverse order (not shown above). Qubit $|c\rangle$ is the comparator ancilla qubit of step 2.1 in section 2.2.



(a) Two-qudit quantum comparator circuit



(b) Three-qudit quantum comparator circuit

Figure 9: Circuits for \hat{C}_k comparator, given $\mathcal{O}(1)$ available ancilla carry qubits, for an (a) $n = 2$ and (b) $n = 3$ qudit register. The diagrams follow the same notation as in Fig. 8. Note that while these choices for \hat{C}_k only require one ancilla carry qubit $|a\rangle$ (which, as always, is initialised to $|0\rangle$), the resulting circuit is much deeper, with many gates containing multiple qudits acting as controls. Each of these gates encodes the logic for a particular scenario in which sums of the individual digits in i and k^c result in an overall output carry bit. For example, for second gate in (b) above, the classical sum $i_0 + k^c[0]$ is large enough to result in an output carry bit, but $i_1 + k^c[1]$, plus the input carry bit from index 0, is not big large enough to give one. However, $i_2 + k^c[2]$ is, even without the input carry bit (if any ‘control set’ in any gate is empty, the gate falls away). As before, the bit value of the most significant qubit is encoded in the comparator qubit $|c\rangle$ with the final controlled- X gate, and all of the other gates are uncomputed after this controlled- X gate (not shown above).




Figure 10: \hat{L}_f subroutine. Circuit for \hat{L}_f , with $|p\rangle$ the payoff qubit (initialised to $|0\rangle$) from step 2.2 of section 2.2. The circuit follows a similar notation as in Figs. 8 and 9 (with any comparator subroutine ancillas $|a\rangle$ not shown here). However, here the blank grey ‘control bubbles’ for each gate indicate which qudit is acting as the control, with the relationship between the control qudit’s state and the R_Y rotation angle given by the argument of the target rotation operation itself. For example, for qudit 0, if $|i_0\rangle = |0\rangle$, then $R_Y(0) = I$ is applied to $|p\rangle$; if $|i_0\rangle = |1\rangle$, then $R_Y\left(\frac{2c\omega}{s_{d^{n-1}-K}}\right)$; if $|i_0\rangle = |2\rangle$, then $R_Y\left(\frac{4c\omega}{s_{d^{n-1}-K}}\right)$; etc., and so on for all qudits. However, these gates (except the very first one) are only applied if the state of the comparator qubit is $|1\rangle$ (this corresponds to the ‘max’ function in the payoff f). Finally, note the different power of d coefficients for each unique controlled rotation, which corresponds with the different digits in i , Eq. 23.