Quantum Monte Carlo simulations for financial risk analytics: scenario generation for equity, rate, and credit risk factors

Titos Matsakos and Stuart Nield

Financial Risk Analytics, Credit & Risk Solutions, Market Intelligence, S&P Global, 25 Ropemaker St, London, EC2Y 9LY, UK

arXiv:2303.09682v2 [quant-ph] 24 Mar 2024

Monte Carlo (MC) simulations are widely used in financial risk management, from estimating value-at-risk (VaR) to pricing over-the-counter derivatives. However, they come at a significant computational cost due to the number of scenarios required for convergence. If a probability distribution is available, Quantum Amplitude Estimation (QAE) algorithms can provide a quadratic speed-up in measuring its properties as compared to their classical counterparts. Recent studies have explored the calculation of common risk measures and the optimisation of QAE algorithms by initialising the input quantum states with pre-computed probability distributions. If such distributions are not available in closed form, however, they need to be generated numerically, and the associated computational cost may limit the quantum advantage. In this paper, we bypass this challenge by incorporating scenario generation — i.e. simulation of the risk factor evolution over time to generate probability distributions — into the quantum computation; we refer to this process as Quantum MC (QMC) simulations. Specifically, we assemble quantum circuits that implement stochastic models for equity (geometric Brownian motion), interest rate (mean-reversion models), and credit (structural, reduced-form, and rating migration credit models) risk factors. We then integrate these models with QAE to provide end-to-end examples for both market and credit risk use cases.

Titos Matsakos: titos.matsakos@spglobal.com

1 Introduction

Several papers have been published recently that demonstrate how quantum algorithms can significantly speed up the solution of typical numerical problems in quantitative finance (for reviews, see [1–7]). In the area of financial risk management, applications include calculating value-at-risk and expected shortfall [8–10], pricing of options [11–15] and collateralised debt obligations [16], credit valuation adjustments [17, 18], and calculating Greeks [19]. While the current state of quantum computing hardware is years away from being used in production systems, theoretical advances on the development and implementation of quantum algorithms are necessary to prepare financial institutions for the transition. The so-called Noisy Intermediate-Scale Quantum (NISQ) technology is considered to outperform current classical computers in the limit of 50-100 qubits ("Intermediate-Scale") and shallow circuit depths ("Noisy", i.e. due to gate errors stacking up) [20]. In this paper we consider fault-tolerant quantum devices; namely, we focus on the representation of stochastic models by quantum gates ignoring the effect of quantum errors.

1.1 Monte Carlo simulations

Financial risk management applications estimate the likelihood and size of potential losses due to hypothetical future events, such as the adverse move of equity prices, interest rates, credit spreads, or foreign exchange rates (market risk), the default of a debt instrument (credit risk), and the default of a counterparty (counterparty credit risk). The risk is generally represented by a statistical measure that depends on the properties of the underlying risk factors. A typical calculation consists of the following three stages: i) gen-

erating a probability distribution $P(\mathbf{X}_t)$ of the risk factor vector $\mathbf{X}_t$ at a future time t, ii) defining a risk measure $F(\mathbf{X}_t)$ as a function of $\mathbf{X}_t$, and iii) employing a statistical method to estimate its value.[1] For example, consider the value V_t of a portfolio consisting of q shares of a given stock, with S_0 its price at $t = 0$. Here, $\mathbf{X}_t$ is one-dimensional with its only element being S_t. The price and volatility of the stock can be used as inputs to model the evolution of S_t and generate a distribution $P(S_t)$. Suppose we define our risk measure as the expected value of the portfolio at time t, $E(V_t) = qE(S_t)$; then, we can estimate $E(V_t)$ with either analytical or numerical methods. Apart from special cases where closed-form expressions are available, the probabilistic nature of the risk factor evolution often requires repeated random sampling followed by statistical estimation — this type of numerical approach is called Monte Carlo (MC) simulation.

Financial risk management calculations that rely on MC simulations often require 10,000 to 1,000,000 experiments to achieve the desired precision. This relatively-slow convergence can be understood by considering the following example: suppose we want to estimate the probability p of a random variable X that takes the value $X = 1$ with p and the value $X = 0$ with $1-p$. In an MC simulation the experiment is repeated N times, obtaining N_1 observations of the outcome $X = 1$ and N_0 of the outcome $X = 0$, with $N_1+N_0 = N$. This is a binomial distribution with the expected value of the number of occurrences of $X = 1$ being $\bar{N}_1 = pN$ and the standard deviation being $\delta N_1 = [p(1-p)N]^{1/2}$. Since N_1 is known but not p, we can invert the expressions to approximate p based on the observed outcomes:

$$p \simeq \bar{p} = \frac{N_1}{N}\,, \tag{1}$$

$$\delta p \simeq \frac{\delta N_1}{N} = \left[\frac{\bar{p}(1-\bar{p})}{N}\right]^{1/2} \propto \left(\frac{1}{N}\right)^{1/2}. \tag{2}$$

From the second equation we can infer that the precision of the estimate is inversely proportional to the square root of N.

[1] The same steps apply for the valuation of many types of exchange-traded or over-the-counter derivatives, in which case $F(\mathbf{X}_t)$ is the pricing function.

1.2 Quantum Amplitude Estimation algorithms

Equation (2) implies that, in order to improve the precision by a decimal digit, classical algorithms require 100 times more experiments; therefore, typical values of N are between 10,000–1,000,000. In quantum computing (see App. A for an introduction and definitions), it has been shown that Quantum Amplitude Estimation (QAE) algorithms can achieve a quadratic speed-up as compared to classical algorithms [21]. For the example of Sect. 1.1, we can estimate the probability p with a quantum computer by encoding the probability distribution of the random variable in the quantum state of a qubit:

$$\begin{aligned}|\psi\rangle &= \sqrt{1-p}\,|0\rangle + \sqrt{p}\,|1\rangle \\ &= \cos(\theta/2)\,|0\rangle + \sin(\theta/2)\,|1\rangle\,,\end{aligned} \tag{3}$$

such that the state $|1\rangle$ is measured with probability p. However, repeating this quantum experiment multiple times (each one called a shot) and measuring the outcomes does not bypass the classical constraint: the precision would again be proportional to the square root of the number of shots as in Eq. (2). Instead, the quadratic gain of QAE algorithms is achieved by leveraging quantum interference. Specifically, consider an input qubit initialised in the superposition of Eq. (3) and a quantum register of n output qubits. The QAE algorithm is based on phase kickback and quantum amplitude estimation [21]. Essentially, the probability p — i.e. the angle θ (Eq. 3) — is imprinted as a phase $\phi = \pm k\theta$ onto the output qubits, where k takes the values $2^0, 2^2, ..., 2^{n-1}$ for the output qubits labelled $0, 1, ..., n-1$, respectively. This operation is then followed by an inverse quantum Fourier transform which, through interference, converts these phases to a binary number; this enables to directly read off the angle θ and hence the value of p. With this approach, the estimated value of p and its error are (see Sect. 1.4.4):

$$p \simeq \bar{p} = \sin^2\left(\frac{\theta}{2}\right)\,, \tag{4}$$

$$\delta p \simeq \sin\theta\frac{\pi}{2N} \propto \frac{1}{N}\,, \tag{5}$$

where $N = 2^n$ is the number of possible outcomes, and, notably, δp scales as the inverse of N.

1.3 Paper motivation and structure

The majority of QAE papers on financial risk have studied quantum circuits with pre-computed probability distributions of the random variables, $P(\mathbf{X}_t)$. However, unless we can exactly encode them in a quantum state, algorithms that involve Monte Carlo integration for state preparation [22] may not be sufficiently efficient to preserve the quadratic speed-up [23]. When such distributions are readily available, there are approaches to efficiently load them on a quantum register [e.g. 24–26], as well as to simplify state preparation and optimise circuits which can help reduce circuit complexity [13, 27]. However, if $P(\mathbf{X}_t)$ is not given in closed form, an MC simulation is needed to generate the distribution numerically with a classical computer. In these cases, even though QAE algorithms can measure its properties with a quadratic speed-up, the computational cost to generate $P(\mathbf{X}_t)$ — associated with the slow convergence of classical algorithms — is paid upfront.

This challenge can be bypassed by incorporating scenario generation into the quantum circuit. In a single shot, the exact distribution of a random variable can be represented as a superposition of quantum states. By avoiding loading pre-computed distributions, this approach also reduces the dependence on classical computers, including the data transfer between classical and quantum systems. In this context, this paper focuses on the quantum implementation of stochastic risk models to generate risk factor scenarios — we refer to this process as QMC simulations.[2] The resulting distributions are then combined with QAE gates to provide end-to-end example circuits. We study the main risk factor groups — equities, rates, and credit — but not exchange rates as they rely on similar models and/or can be derived from interest rates. Depending on the complexity of the stochastic differential equations that describe the evolution of risk factors, theoretical studies have shown that MC simulations can be further optimised by leveraging quantum-accelerated multilevel MC methods [28–30]. Here, we limit our scope to assembling QMC circuits for the foundational stochastic models — i.e. geometric Brownian motion, mean-reversion models, Poisson processes, and multinomial trees — as most risk models are based on these or a variation thereof.[3]

The paper is structured as follows. Section 1.4 presents the main steps of QAE algorithms and their implementation in quantum circuits. In Sect. 2, we assemble quantum circuits to perform QMC simulations, namely to model the evolution of equity (Sect. 2.1), interest rate (Sect. 2.2), and credit (Sect. 2.3) risk factors. Section 3 estimates the number of qubits and the resulting circuit depth needed in typical financial risk use cases. We conclude in Sect. 4.

1.4 Overview of QAE circuits

1.4.1 The circuit

To estimate a statistical measure $F(X_t) \in [F_{\min}, F_{\max}]$ of a risk factor X_t, a QAE circuit would typically consist of:

1. m "risk factor" (rf) qubits to model the distribution of X_t (the initial state of which is denoted with $|0\rangle_{\text{rf}}^{\otimes m}$),

2. one "risk measure" (rm) qubit to encode the normalised value of the risk measure[4] $f(X_t) \in [0,1]$ in the angle $\theta \in [0,\pi]$ (the initial state of which is denoted with $|0\rangle_{\text{rm}}$),

3. n output (out) qubits to imprint multiples of θ onto their phases (the initial state of which is denoted with $|0\rangle_{\text{out}}^{\otimes n}$).[5]

Without loss of generality, in the following example we choose $f(X_t) = p$, i.e. the risk measure to be a probability of an outcome that depends on X_t, i.e. $f(X_t) = p$.

The general structure of a QMC/QAE quantum circuit is:

[2] The term "QMC simulations" is not related to "Quantum Monte Carlo methods" used to study complex quantum systems in Physics.

[3] For an overview of risk models, as well as their limitations and extensions, see [31].

[4] E.g. $f(X_t) = (F(X_t) - F_{\min})/(F_{\max} - F_{\min})$.

[5] We use the notation $|0\rangle \otimes |0\rangle = |0\rangle^{\otimes 2} = |0\rangle\,|0\rangle = |00\rangle$.

State	Description
$\|b_{m-1}...b_1b_0\rangle$	m qubits as a binary, e.g. $\|101\rangle$
$\|j\rangle$, $\|x\rangle$, $\|z\rangle$	qubits as an integer, e.g. $\|5\rangle$
$\|\psi\rangle$	superposition of states
$\|\psi\rangle_{\text{in}}$	input qubit(s)
$\|\psi\rangle_{\text{out}}$	output qubit(s)
$\|\psi\rangle_{\text{rf}}$	"risk factor" qubit(s) (input)
$\|\psi\rangle_{\text{rm}}$	"risk measure" qubit (input)
$\|\psi\rangle_{\text{st}}$	"state" qubit(s)
$\|\psi\rangle_{\text{c}}$	"count" qubit(s)
$\|\psi\rangle_{\text{anc}}$	"ancilla" qubit(s)

Table 1: Notation of states.

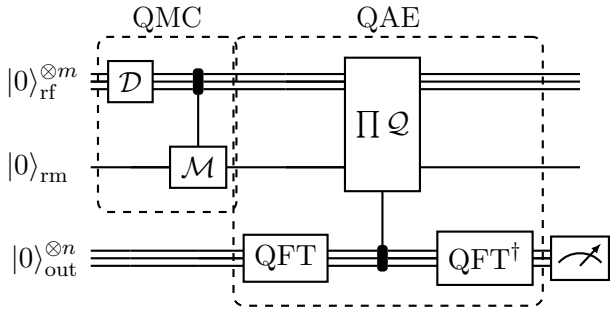

where $\mathcal{D}$ is the gate that generates the input distribution using m "risk factor" qubits (the focus of this paper), $\mathcal{M}$ the controlled gate that encodes the risk measure into the angle θ of the risk measure qubit, $\prod \mathcal{Q}$ the repeated application of the controlled gate $\mathcal{Q}$ to imprint θ onto the phases of the n output qubits, and QFT/QFT† the quantum Fourier transformation and its inverse to measure the phase of the output qubits with interference. These circuit components are described below in more detail, see Table 1 for state notation.

1.4.2 $\mathcal{D}$: preparing the distribution $P(X_t)$

The m qubits, $|0\rangle_{\text{rf}}^{\otimes m}$, can model a discrete probability distribution of 2^m possible outcomes, each one with a probability $|a_j|^2$, with $j \in \{0, 1, ..., 2^m - 1\}$. We use the gate $\mathcal{D}$ to load the distribution of X_t onto the "risk factor" qubits, which can be depicted as:

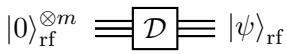

and expressed mathematically as:

$$\begin{aligned} |\psi\rangle_{\text{rf}} &= \mathcal{D}\,|0\rangle_{\text{rf}}^{\otimes m} \\ &= \sum_{j=0}^{2^m-1} a_j\,|b_{jm-1}...b_{j1}b_{j0}\rangle_{\text{rf}} = \sum_{j=0}^{2^m-1} a_j\,|j\rangle_{\text{rf}}\ , \end{aligned} \tag{6}$$

where j is an integer in the decimal number system representing the binary number $b_{jm-1}...b_{j1}b_{j0}$, and $b_{jl} \in \{0, 1\}$ is the l-th digit of the j-th state. If the distribution is precomputed, it can be loaded by leveraging Quantum Generative Adversarial Networks [24] or Fourier expansion [25], among other approaches [26, 27]. Here, the distributions will be generated with quantum gates that implement stochastic models for the risk factor evolution.

1.4.3 $\mathcal{M}$: calculating the risk measure $f(X_t) = p$

With the distribution $P(X_t)$ encoded in the state of the "risk factor" qubits, $|\psi\rangle_{\text{rf}}$, the next step is to encode a risk measure $f(X_t)$ — for example a probability such as $p = P(X_t < K)$ where $K \in [X_{\min}, X_{\max}]$ — in the state of the "risk measure" qubit, $|0\rangle_{\text{rm}}$. The value of p can be captured in the angle θ of a qubit with the help of a controlled gate $\mathcal{M}$ that reads the distribution $|\psi\rangle_{\text{rf}}$ and encodes $p \in [0, 1] \to \theta \in [0, \pi]$ into $|0\rangle_{\text{rm}}$. This can be represented as:

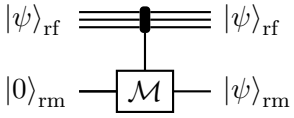

and described by the expression:

$$\begin{aligned} |\psi\rangle_{\text{in}} &= \mathcal{M}\,|\psi\rangle_{\text{rf}}\,|0\rangle_{\text{rm}} \\ &= \sqrt{1-p}\,|\psi_0\rangle_{\text{rf}}\,|0\rangle_{\text{rm}} + \sqrt{p}\,|\psi_1\rangle_{\text{rf}}\,|1\rangle_{\text{rm}} \\ &= \sqrt{1-p}\,|\psi_0\rangle_{\text{in}} + \sqrt{p}\,|\psi_1\rangle_{\text{in}}\ , \end{aligned} \tag{7}$$

where $\sqrt{1-p} = \cos(\theta/2)$, $\sqrt{p} = \sin(\theta/2)$, and we have simplified notation by writing $|\psi_0\rangle_{\text{in}} = |\psi_0\rangle_{\text{rf}}\,|0\rangle_{\text{rm}}$ and $|\psi_1\rangle_{\text{in}} = |\psi_1\rangle_{\text{rf}}\,|1\rangle_{\text{rm}}$.

1.4.4 $\prod \mathcal{Q}$ and QFT: estimation of p

Amplitude amplification The next part of the circuit is the controlled $\mathcal{Q}$ gate, an operator which is based on the Grover search algorithm [32]. $\mathcal{Q}$ consists of two reflections, $|\psi_0\rangle_{\text{in}} \to -|\psi_0\rangle_{\text{in}}$ and $|\psi\rangle_{\text{in}} \to -|\psi\rangle_{\text{in}}$:

$$\begin{aligned} \mathcal{Q} &= \mathcal{Q}_\psi\,\mathcal{Q}_{\psi_0} \\ &= (\mathbb{1} - 2\,|\psi\rangle_{\text{in}}\,\langle\psi|_{\text{in}})(\mathbb{1} - 2\,|\psi_0\rangle_{\text{in}}\,\langle\psi_0|_{\text{in}})\,, \end{aligned} \tag{8}$$

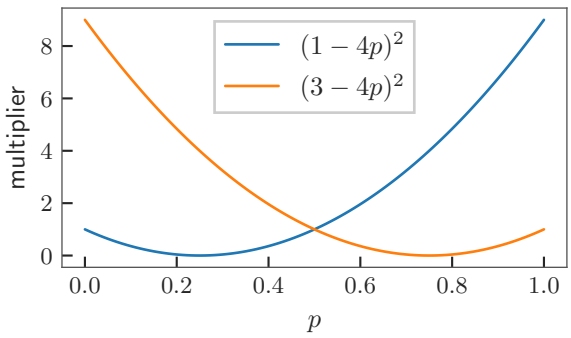

Figure 1: The amplification of the probabilities of the states $|\psi_0\rangle_{\text{in}}$ (blue) and $|\psi_1\rangle_{\text{in}}$ (orange) when $\mathcal{Q}$ acts on ψ (Eq. 9).

which amplify/deamplify the amplitudes of $|\psi_0\rangle_{\text{in}}$ and $|\psi_1\rangle_{\text{in}}$ depending on the value of p (see App. B.1):

$$\begin{aligned}&\mathcal{Q}\,|\psi\rangle_{\text{in}} = \\ &= (1-4p)\sqrt{1-p}\,|\psi_0\rangle_{\text{in}} + (3-4p)\sqrt{p}\,|\psi_1\rangle_{\text{in}}\ .\end{aligned} \tag{9}$$

Fig. 1 shows the multipliers $(1-4p)^2$ and $(3-4p)^2$ which increase/decrease the probabilities of measuring $|0\rangle_{\text{rm}}$ and $|1\rangle_{\text{rm}}$, respectively.[6] A key property of $\mathcal{Q}$ is that it leaves the states:

$$|\psi_\pm\rangle_{\text{in}} = \frac{1}{\sqrt{2}}\left(\,|\psi_1\rangle_{\text{in}} \pm i\,|\psi_0\rangle_{\text{in}}\,\right), \tag{10}$$

unchanged when applied k times, but introduces a global phase (see App. B.2):

$$\mathcal{Q}^k\,|\psi_\pm\rangle_{\text{in}} = e^{\pm ik\theta}\,|\psi_\pm\rangle_{\text{in}}\ . \tag{11}$$

Phase kickback We can exploit this property by implementing a controlled gate $\mathcal{Q}$ acting on $|\psi\rangle_{\text{in}}$, which we express as (see App. B.3):

$$\begin{aligned}&|\psi\rangle_{\text{in}} \\ &= -i\frac{1}{\sqrt{2}}\left(e^{i\theta/2}\,|\psi_+\rangle_{\text{in}} - e^{-i\theta/2}\,|\psi_-\rangle_{\text{in}}\right).\end{aligned} \tag{12}$$

The target of this control gate is the output qubits, after first rotating their initial state, $|0...00\rangle_{\text{out}}$, with the H operator such that it becomes $|+...++\rangle_{\text{out}}$. This is equivalent to a quantum Fourier transformation, which we can represent as:

[6]For $p < 1/2$, the probability of measuring the state $|1\rangle_{\text{rm}}$ is amplified and that of the state $|0\rangle_{\text{rm}}$ deamplified, and vice versa for $p > 1/2$.

$$|0\rangle_{\text{out}}^{\otimes n} \equiv \boxed{\text{QFT}} \equiv |+\rangle_{\text{out}}^{\otimes n}$$

and mathematically express as:

$$\text{QFT}\,|0\rangle_{\text{out}}^{\otimes n} = \bigotimes_{l=0}^{n-1} H\,|0\rangle_{\text{out}}^{l} = |+\rangle_{\text{out}}^{\otimes n}\ . \tag{13}$$

Then, we can leverage phase kickback to imprint the multiples of the angle θ onto their phases. Based on Eq. (11), the phase of an output qubit $|+\rangle_{\text{out}}^{l} = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ will become $\frac{1}{\sqrt{2}}(|0\rangle + e^{ik\theta}\,|1\rangle)$ for $|\psi_+\rangle_{\text{in}}$, and $\frac{1}{\sqrt{2}}(|0\rangle + e^{-ik\theta}\,|1\rangle)$ for $|\psi_-\rangle_{\text{in}}$, where we set $k = 2^l$. Namely, $k = 2^0$ for the first output qubit, $k = 2^1$ for the second, ..., and $k = 2^{n-1}$ for the last (n-th). The repeated application of the controlled gate $\mathcal{Q}$ can be represented as:

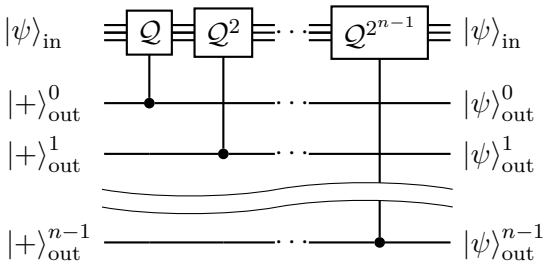

and the state of the output qubits can be written as (see App. B.4):

$$\begin{aligned}&\prod_{l=0}^{n-1}\mathcal{Q}^{2^l}\,|\psi\rangle_{\text{in}}\,|+\rangle_{\text{out}}^{\otimes n} = \\ &= |\psi\rangle_{\text{in}}\left\{\bigotimes_{l=0}^{n-1}\left[\frac{1}{\sqrt{2}}\left(|0\rangle + e^{\pm i2^l\theta}\,|1\rangle\right)\right]\right\} \\ &= |\psi\rangle_{\text{in}}\left(\frac{1}{2^{n/2}}\sum_{x=0}^{2^n-1} e^{\pm ix\theta}\,|x\rangle\right),\end{aligned} \tag{14}$$

where $x = \sum_{l=0}^{n-1} 2^l b_l$ is the binary number $b_{n-1}...b_1b_0$ expressed in the decimal number system.

Interference The state $\sum_{x=0}^{2^n-1} a_x\,|x\rangle$, where $a_x = (1/2^{n/2})e^{\pm ix\theta}$, can be transformed to the standard basis $\sum_{z=0}^{2^n-1} a_z\,|z\rangle$ with an inverse quantum Fourier transform, which depicted by the gate $\text{QFT}^\dagger$:

$$\textstyle\sum_x a_x\,|x\rangle \equiv \boxed{\text{QFT}^\dagger} \equiv \sum_z a_z\,|z\rangle$$

and mathematically expressed by (see App. B.5):

$$|\psi\rangle_{\text{out}} = \text{QFT}^\dagger \frac{1}{2^{n/2}} \sum_{x=0}^{2^n-1} e^{\pm ix\theta} |x\rangle$$
$$= \frac{1}{2^n} \sum_{z=0}^{2^n-1} \sum_{x=0}^{2^n-1} e^{ix(\pm\theta - 2\pi z/2^n)} |z\rangle \,, \quad (15)$$

which consists of a superposition of the states $|z\rangle$, where z is an integer in the decimal number system.

Measurement The measurement of the state $\sum_z a_z |z\rangle$ will result in the collapse of the superposition and will give one of the possible z values, z_0. This can be represented by:

$\sum_z a_z |z\rangle$ ⟹ $|z_0\rangle$

If θ has a value such that an integer $z_0 \in [0, 2^{n-1}]$ exists that makes $\pm\theta - 2\pi z/2^n$ a multiple of 2π, namely, $z_0 = 2^n\theta/2\pi$ or $z_0 = 2^n(2\pi-\theta)/2\pi$, then Eq. (15) implies that the measured state is one of (see App. B.6):

$$|z_0\rangle = |2^n\theta/2\pi\rangle \,, \quad (16)$$
$$|z_0\rangle = |2^n(2\pi-\theta)/2\pi\rangle \,. \quad (17)$$

Then, it is straightforward to calculate the probability p by substituting θ in Eq. (3) with the measured value of z_0:

$$p = \sin^2\left(\frac{\theta}{2}\right) = \sin^2\left(\frac{z_0}{2^n}\pi\right) . \quad (18)$$

If there is no z_0 such that $\pm\theta - 2\pi z_0/2^n$ is a multiple of 2π, the closest integer z_0 is measured with a probability $|a_{z_0}|^2 \sim 20\%$ for either θ or $2\pi-\theta$, which gives a total probability of approximately 40% to get the closest value of p (see App. B.7).

The precision with which θ is estimated is $\delta\theta \simeq 2\pi/2^n$; therefore, the precision of p is (see App. B.8):

$$\delta p \simeq \sin\theta \frac{\pi}{N} \propto \frac{1}{N} \,, \quad (19)$$

where $N = 2^n$ is the total number of possible outcomes. Notably, the error in QAE decreases proportionally to $1/N$ (Eq. 19), much faster than the $1/N^{1/2}$ scaling of classical algorithms (Eq. 2).

Apart from the basic QAE algorithm described above which is based on quantum phase estimation (QPE), recent studies have explored other variants of the QAE family that do not require QPE [33–35].

2 Quantum gates and circuits for scenario generation

2.1 Equity risk factors

Consider an equity price, S_t, that follows the stochastic differential equation:

$$dS_t = \mu S_t dt + \sigma S_t dW_t \,, \quad (20)$$

where μ is the drift, σ the volatility, t the time, and W_t a Wiener process. Itô's lemma implies that

$$d\ln S_t = \left(\mu - \frac{\sigma^2}{2}\right) dt + \sigma dW_t \,, \quad (21)$$

and thus S_t is log-normally distributed,

$$S_{t+dt} = S_t e^{\left(\mu-\sigma^2/2\right)dt + \sigma dW_t} \,, \quad (22)$$

with expected value and variance:

$$E(S_{t+dt}) = S_t e^{\mu dt} \,, \quad (23)$$
$$Var(S_{t+dt}) = E(S_{t+dt}^2) - E(S_{t+dt})^2$$
$$= S_t^2 e^{2\mu dt + \sigma^2 dt} - S_t^2 e^{2\mu dt} \,. \quad (24)$$

The path of the price can be modelled with a binomial tree with $m+1$ nodes at times $t = \{0, \delta t, 2\delta t, ..., T\}$, respectively, where $\delta t = T/m$ is the time interval. At each node, the price S_t can either go up by a factor u, $S^{\text{u}}_{t+\delta t} = S_t u$, or down by a factor d, $S^{\text{d}}_{t+\delta t} = S_t d$, with probabilities q and $1-q$, respectively.

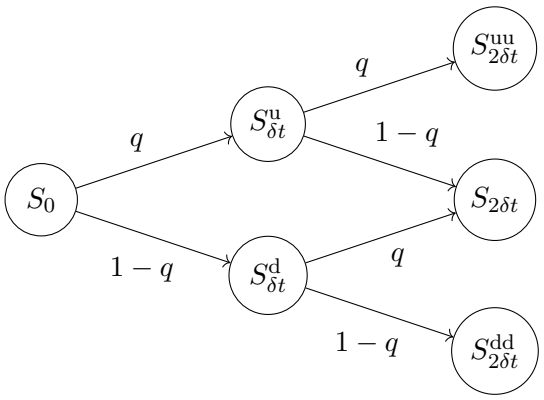

By requiring the discrete model to have the same mean and variance as the continuous model, we obtain the following expressions for u and q, respectively (see App. C.1):

$$q = \frac{ue^{\mu\delta t} - 1}{u^2 - 1} \,, \quad (25)$$
$$u = \frac{1}{d} = e^{\sigma\sqrt{\delta t}} \,. \quad (26)$$

Binomial trees are common in pricing equity derivatives [36].

2.1.1 $\mathcal{D}_{\text{eq}}$: the distribution $P(S_t)$

Since at each timestep the stock price has two possible outcomes, we can model a transition from time period t to $t+\delta t$ with a qubit in superposition, $|\psi\rangle_{\text{rf}}^{t+\delta t}$, such that the states $|0\rangle$ and $|1\rangle$ represent downwards and upwards moves, respectively. Therefore, a scenario consisting of m timesteps can be modelled with m qubits:

$$|\psi\rangle_{\text{rf}} = |\psi\rangle_{\text{rf}}^{\delta t} |\psi\rangle_{\text{rf}}^{2\delta t} \dots |\psi\rangle_{\text{rf}}^{m\delta t} , \tag{27}$$

where the superscripts δt, $2\delta t$, ..., $m\delta t$ are labels of individual qubits. For $m = 2$, the binomial tree is:

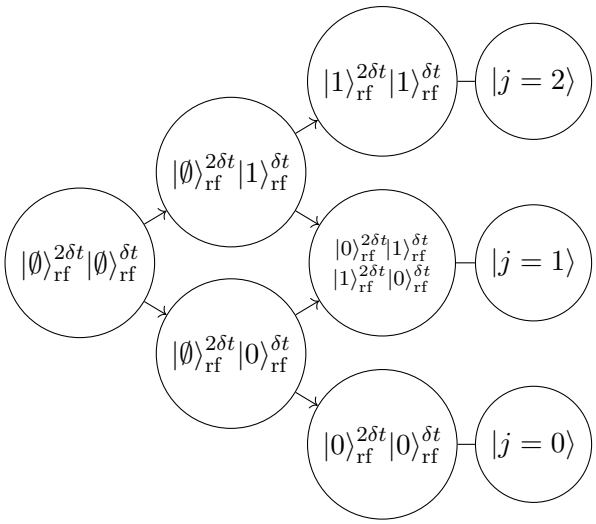

where we use the notation $|\emptyset\rangle = |0\rangle$ to denote a qubit that has not yet been through a gate.

At each time t, the probability q can be encoded into the angle θ_{u} of a "risk factor" qubit with a y-rotation gate, $R_y(\theta_{\text{u}})$, i.e. $|\psi\rangle_{\text{rf}}^{t} = \cos(\theta_{\text{u}}/2)\,|0\rangle + \sin(\theta_{\text{u}}/2)\,|1\rangle$ with $q = \sin^2(\theta_{\text{u}}/2)$:

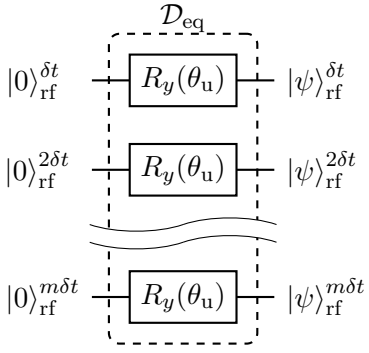

The inverse of this gate, $\mathcal{D}_{\text{eq}}^{\dagger}$, consists of the same rotation gates but with a negative angle, $-\theta_{\text{u}}$:

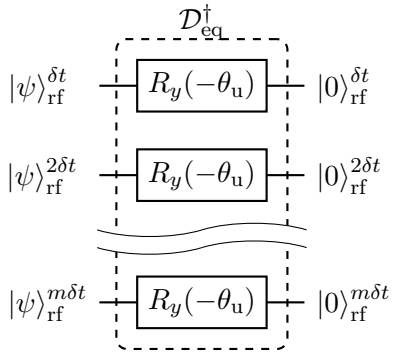

2.1.2 $\mathcal{M}_{\text{eq}}$: risk measures $F(S_t)$

The choice of the risk measure depends on the use case; here, we consider as examples the probabilities of observing the maximum and minimum values.[7] For some risk measures, m_{anc} additional ("ancilla") qubits are required to store the intermediate values of the calculation. We ensure that all assembled gates leave the state of these qubits to their initial state $|0\rangle_{\text{anc}}^{\otimes m_{\text{anc}}}$;[8] therefore, we do not explicitly show them when not needed.

Maximum The gate that calculates the probability of S_T taking its maximum value, $F(S_T) = P(S_{\max}) = P(|1\rangle_{\text{rf}}^{\otimes m}) = q^m$, can be assembled with a sequence of AND (Toffoli) gates, such that the "risk measure" qubit flips to the state $|1\rangle_{\text{rm}}$ if all "risk factor" qubits are in the state $|1\rangle_{\text{rf}}^{\otimes m}$; namely, $|\psi\rangle_{\text{rm}} = (1-q^m)\,|0\rangle + q^m\,|1\rangle$. In its simplest form, an additional $m-2$ "ancilla" qubits are required to store the result of the AND operators, which are then applied a second time to revert the "ancilla" qubits back to their original $|0\rangle_{\text{anc}}^{\otimes m_{\text{anc}}}$ state.

[7] While min and max are not typical risk measures of equity price distributions, such functions appear in the payoff of equity options [31] and thus their quantum implementation is important for pricing derivatives.

[8] Reverting the ancilla qubits back to their original states with a quantum gate rather than a "reset" instruction ensures the gate is reversible. This is because the QAE algorithm relies on the application of $\prod \mathcal{Q}$, which is based on the operator of Eq. (30) (discussed later on) that includes the inverse of both the $\mathcal{D}$ and $\mathcal{M}$ gates.

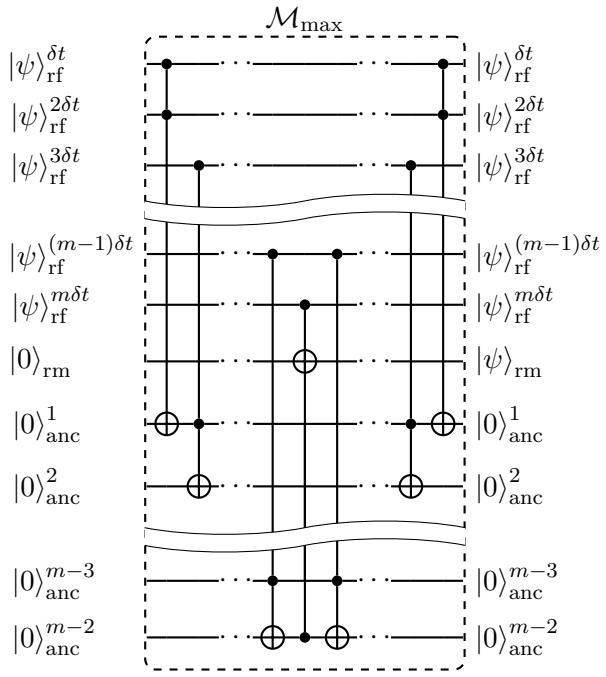

Note that if we apply $\mathcal{M}_{\max}$ twice, the "risk measure" qubit returns to its initial state $|0\rangle_{\rm rm}$, while the states of all other qubits remain unchanged: $\mathcal{M}_{\max}^{\dagger}\mathcal{M}_{\max} = \mathbb{1}$; thus, $\mathcal{M}_{\max}^{\dagger} = \mathcal{M}_{\max}$.

Minimum The probability of measuring the minimum value, $F(S_T) = P(S_{\min}) = P(|0\rangle_{\rm rf}^{\otimes m}) = (1-q)^m$, can be constructed following the same logic. We first flip all "risk factor" qubits and then apply AND operators such that only the state $|0\rangle_{\rm rf}^{\otimes m}$ will lead to the state $|1\rangle_{\rm rm}$ of the "risk measure" qubit: $|\psi\rangle_{\rm rm} = [1-(1-q)^m]\,|0\rangle + (1-q)^m\,|1\rangle$. The corresponding quantum gate can be assembled by leveraging $\mathcal{M}_{\max}$:

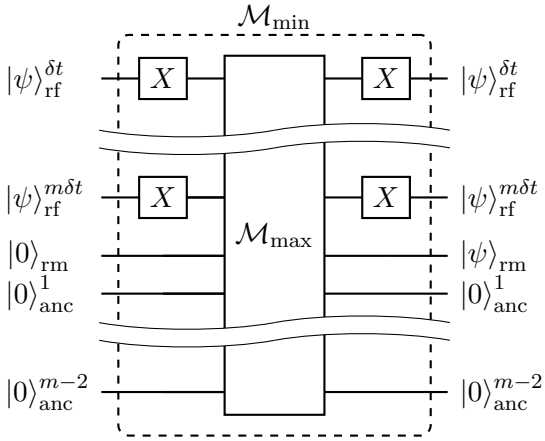

Here too, we notice that the inverse gate is $\mathcal{M}_{\min}^{\dagger} = \mathcal{M}_{\min}$.

Other risk measures Gates for other risk measures can be assembled following similar logic, e.g. for value-at-risk and expected shortfall see Refs. [8, 9]. In Sect. 2.3 we show an example of probability distributions that involve inequalities, e.g. $P(S_T \geq u^{j_T} S_0)$, where j_T is a specified number of upwards moves. Generally, any risk measure can be calculated by using quantum gates for arithmetic operations and comparisons [37, 38].

2.1.3 Phase estimation

To assemble the gate $\mathcal{Q} = \mathcal{Q}_{\psi}\mathcal{Q}_{\psi 0}$, we represent the operator $\mathcal{Q}_{\psi 0} = \mathbb{1} - 2\,|\psi_0\rangle_{\rm in}\,\langle\psi_0|_{\rm in}$ with a gate that flips the sign of the state $|\psi_0\rangle_{\rm in} = |\psi_0\rangle_{\rm rf}\,|0\rangle_{\rm rm}$:

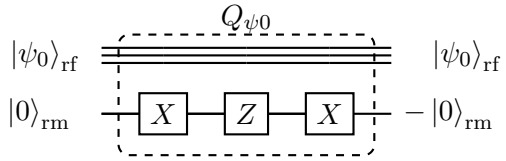

$$\mathcal{Q}_{\psi 0}\,|\psi_0\rangle_{\rm in} = -\,|\psi_0\rangle_{\rm in}\,, \tag{28}$$
$$\mathcal{Q}_{\psi 0}\,|\psi_1\rangle_{\rm in} = |\psi_1\rangle_{\rm in}\,. \tag{29}$$

For $\mathcal{Q}_{\psi}$, we decompose the operator as (see App. C.2):

$$\mathcal{Q}_{\psi} = \mathcal{M}\mathcal{D}Q_{00}\mathcal{D}^{\dagger}\mathcal{M}^{\dagger}\,, \tag{30}$$

where $\mathcal{Q}_{00} = \mathbb{1} - 2\,|0\rangle_{\rm in}\,\langle 0|_{\rm in}$ is a reflection of the initial state $|0\rangle_{\rm in} = |0\rangle_{\rm rf}^{\otimes m}\,|0\rangle_{\rm rm}$:

$$\mathcal{Q}_{00}\,|0\rangle_{\rm in} = -\,|0\rangle_{\rm in}\,, \tag{31}$$
$$\mathcal{Q}_{00}\,|\psi \neq 0\rangle_{\rm in} = |\psi\rangle_{\rm in}\,. \tag{32}$$

To apply a reflection only if all input qubits are in the state $|0\rangle$, we first flip them with NOT gates, we then operate with a sequence of AND gates storing the result in the last "ancilla" qubit, and finally we apply a controlled Z gate to flip the sign of the "risk measure" qubit if the last "ancilla" qubit was in the state $|1\rangle$. Operating with the AND and NOT gates again brings all qubits to their initial state (apart from the sign of the "risk measure" qubit depending on the initial state).

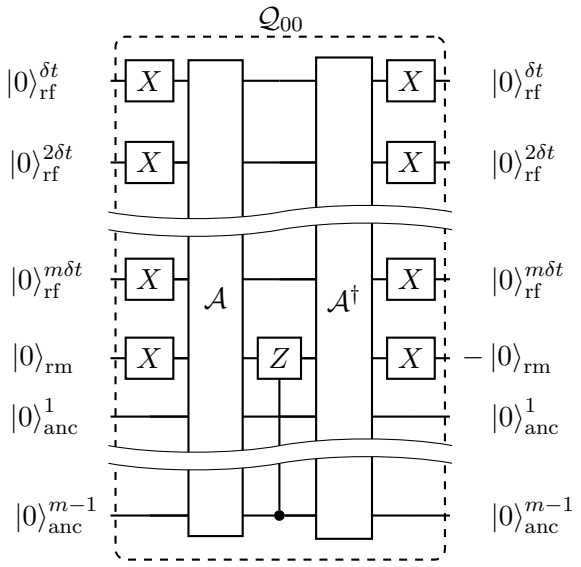

where the $\mathcal{A}$ and $\mathcal{A}^\dagger$ gates consist of the AND gates:

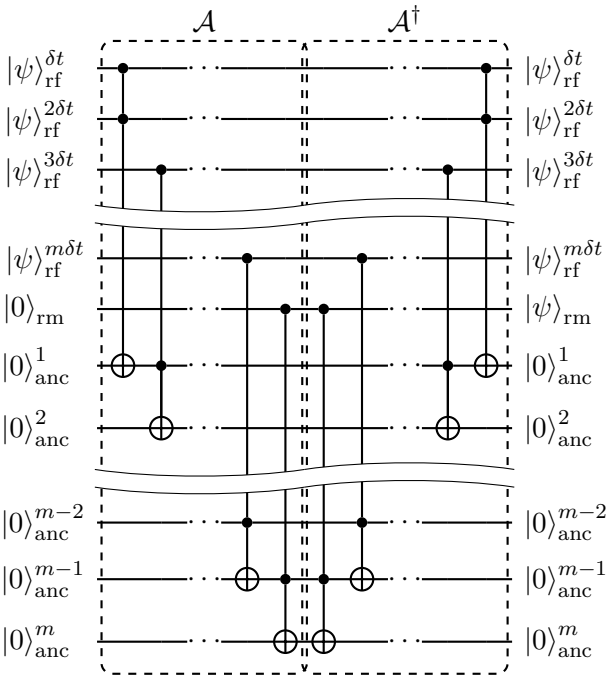

2.1.4 Measurement

The QFT gate prepares the output qubits in the state $|+\rangle$:

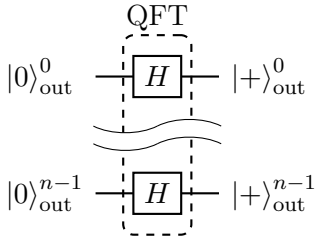

Parameter	Value
m	6
n	1-9
T	1
μ	8%
σ	20%
u	∼1.09
q	∼0.56
$\theta_{\rm u}\frac{180^\circ}{\pi}$	∼97.1°

Table 2: List of the binomial tree parameters for the equity risk factor evolution.

The QFT† gate converts the phases $\pm k\phi$ to an integer $z \in [0, 1, ..., 2^n - 1]$; Fig. 2 shows the implementation based on [39], where the controlled rotations between qubits l_1 and l_2 are:

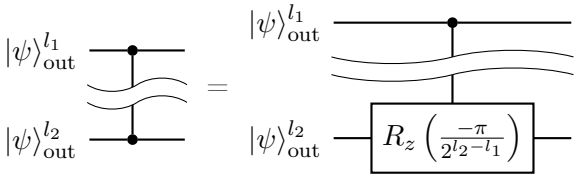

See App. C.3 for an example of a complete quantum circuit that estimates $P(S_{\max})$ for $m = 2$ and $n = 3$. The accompanying figures show a visualisation of the qubit states and their transformations on the Bloch sphere, as well as the corresponding probability distributions when measured.

2.1.5 Results

To analyse the convergence of the estimate of the risk measure we assemble 9 quantum circuits, each one with a different number of output qubits, $n \in [1, 2, ..., 9]$, and adopt the parameters listed in Table 2 (the value of S_0 is not needed). The top panel of Fig. 3 shows the probability distribution $P(S_T)$ obtained after applying the gate $\mathcal{D}_{\rm eq}$ on the "risk factor" qubits. The distribution is represented by the states $|j\rangle$ (see Eq. 6), which consist of all possible 2^m combinations of the "risk factor" qubits when j of them are in the state $|1\rangle$ (i.e. j upwards price moves). Since $q > 1/2$, the distribution is positively skewed: $P(S_T \geq u^2 S_0) > P(S_T \leq d^2 S_0)$. The gates $\mathcal{M}_{\max}$ and $\mathcal{M}_{\min}$ essentially measure the probabilities of the states $|6\rangle$ and $|0\rangle$, respectively,

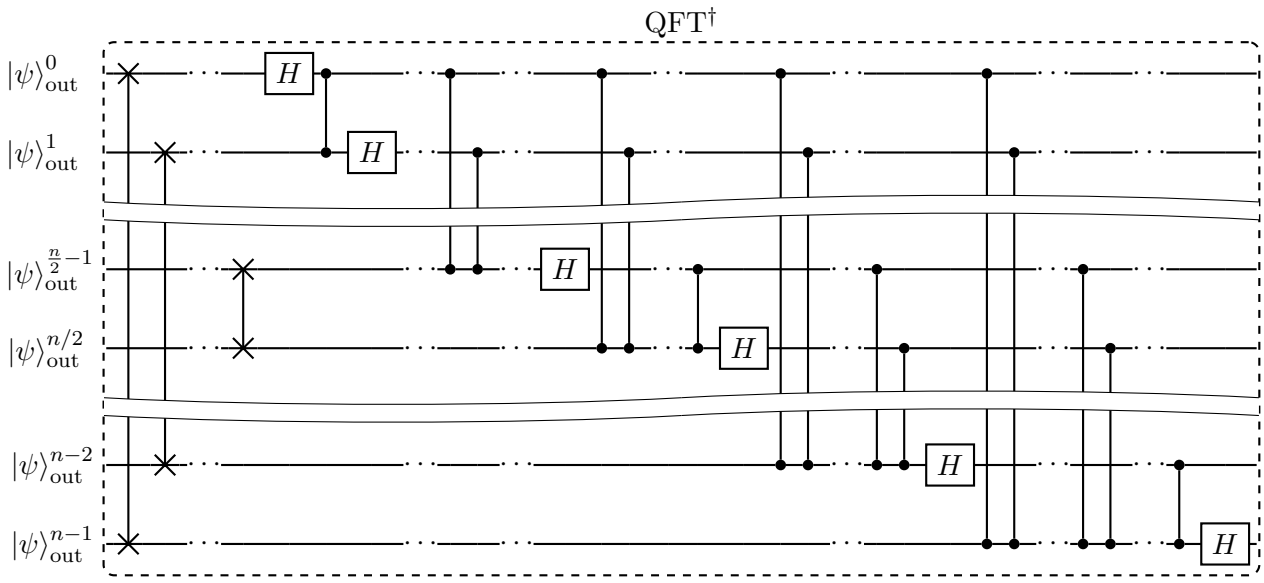

Figure 2: The inverse quantum Fourier transform gate, QFT†.

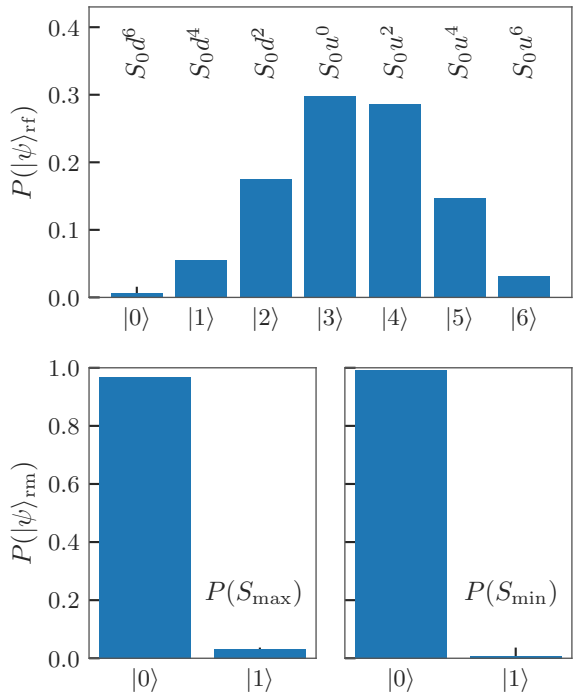

Figure 3: Top: the distribution of S_T for $m = 6$ qubits as generated with the operator $\mathcal{D}_{\text{eq}}$. Bottom: the probabilities $P(S_{\max})$ (left) and $P(S_{\min})$ (right) as encoded with the operators $\mathcal{M}_{\max}$ and $\mathcal{M}_{\min}$, respectively, into the "risk measure" qubit.

which are encoded in the "risk measure" qubit; see bottom panel of Fig. 3.

The bar charts of Fig. 4 show the distribution $P(|z\rangle)$ obtained by measuring the output qubits (here $n = 4$ and 10,000 shots) for the two risk measures, $P(S_{\max})$ (top) and $P(S_{\min})$ (bottom). The left and right peaks correspond to $z = 2^n\theta/2\pi$ and $z = 2^n(2\pi - \theta)/2\pi$, respectively, originating from the two signs of the phase kickback, $\pm k\theta$. The line charts of Fig. 4 show the convergence of the estimated probability p (Eq. 18), $P(S_{\max})$ (top) and $P(S_{\min})$ (bottom), and error δp (Eq. 19), as a function of the number of output qubits, n. Here, the measured state of the output qubits is obtained from one shot. The expected values of the probabilities are $P(S_{\max}) = q^m$ and $P(S_{\min}) = (1 - q)^m$, respectively.

Equity risk factors can also be based on trinomial trees, the implementation of which is presented in Sect. 2.2.

2.2 Interest rate risk factors

The evolution of interest rates can be simulated with short-rate mean reversion models, the most basic of which is the Vasicek model [40]:

$$dr_t = a(b - r_t)dt + \sigma dW_t\,, \tag{33}$$

where r_t is the instantaneous interest rate at time t, b the long-term mean, a the speed of reversion, σ the volatility, and W_t a Wiener process. The expected value and variance are asymptotically

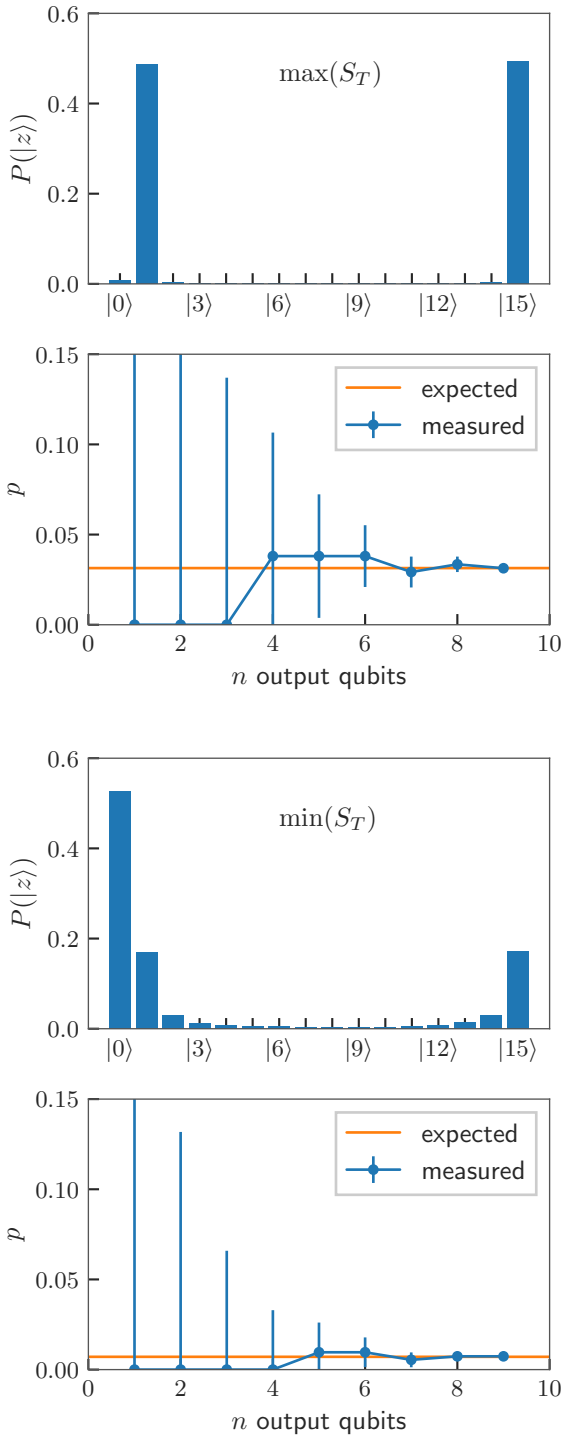

Figure 4: Top two panels: the probability distribution $P(S_{\max})$ (top) when measuring z_0 with 10,000 shots using $n = 4$ qubits, and the convergence of $P(S_{\max}) \to p = q^m$ (with $\delta p \to 0$, bottom), as a function of the number of output qubits. The measured value of p (blue dots) is calculated using Eq. (18) after measuring the result of one shot, the error δp (blue lines) is calculated from Eq. (19), and the expected value is shown with an orange line. Bottom two panels: the probability distribution $P(S_{\min})$ (top) and $P(S_{\min}) \to p = (1-q)^m$ (bottom).

constant for $\delta t \to \infty$, whereas for finite δt the expected value depends on r_t:

$$E(r_{t+\delta t}) = r_t e^{-a\delta t} + b\left(1 - e^{-a\delta t}\right), \tag{34}$$

$$Var(r_{t+\delta t}) = \frac{\sigma^2}{2a}\left(1 - e^{-2a\delta t}\right). \tag{35}$$

Such models are often discretised with trinomial trees. For example, consider a simple tree that is bounded at both low and high interest rates, and, at each timestep, r_t can take one of three possible values: $r_t^{\rm h} = b + \delta r$ (high), $r_t^{\rm m} = b$ (mid), and $r_t^{\rm l} = b - \delta r$ (low), with δr a constant.

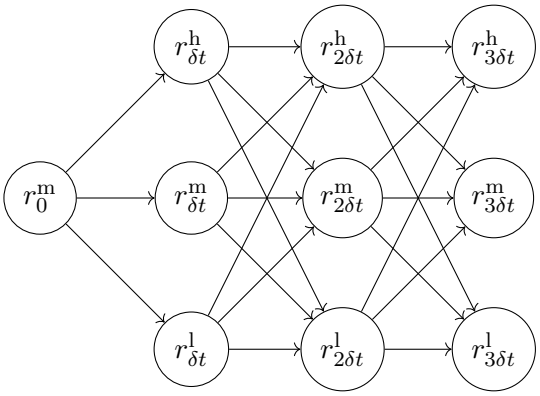

The transition probabilities $q_{t\to t+\delta t}$ from a node r_t to a node $r_{t+\delta t}$ depend both on the t and $t+\delta t$ nodes, and thus $q_{t\to t+\delta t}$ is an array of $3 \times 3 = 9$ values. These probabilities can be derived by equating the expected value and variance of the continuous and discrete models, respectively (see App. D).

2.2.1 $\mathcal{D}_{\rm ir}$: the distribution $P(r_t)$

Given that at each timestep t there are 3 possible outcomes we need two qubits, $|0\rangle^t_{\rm rf_1}|0\rangle^t_{\rm rf_0}$, to encode the probabilities q; therefore, m timesteps require $2m$ "risk factor" qubits. At each timestep, we use the first qubit ($\rm rf_0$) from the pair to represent whether the interest rate transitions to the mid level: $r_t \to r^m_{t+\delta t}$. Specifically, the state $|0\rangle^t_{\rm rf_0}$ describes a transition to the mid node and the state $|1\rangle^t_{\rm rf_0}$ a transition to a different node. In the latter case, we use the second qubit ($\rm rf_1$) to model whether r_t transitions to the high ($r_t \to r^h_{t+\delta t}$, described by $|1\rangle^t_{\rm rf_1}$) or low ($r_t \to r^l_{t+\delta t}$, described by $|0\rangle^t_{\rm rf_1}$) levels, respectively. Because the transition probabilities depend on the level of the interest rate, the "risk factor" state needs to be read at each timestep before encoding the probabilities for the next transition. This can be facilitated by

including 3 additional "state" qubits — alongside the $2m$ "risk factor" qubits — to store the value of r_t at each timestep; i.e. $|001\rangle^t_{\rm st}$, $|010\rangle^t_{\rm st}$, and $|100\rangle^t_{\rm st}$, which represent the high, mid, and low interest rate levels, respectively. Assuming r_0 is at the mid level, the tree is:

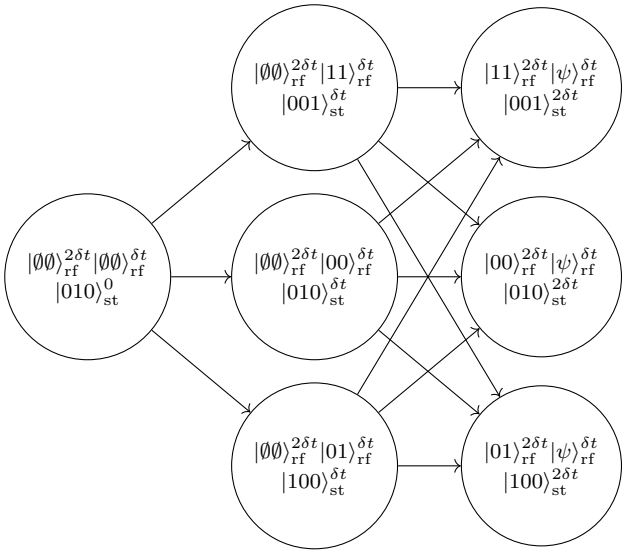

The quantum state at the first timestep is:

$$|\psi\rangle^{\delta t}_{\rm rf}|\psi\rangle^{\delta t}_{\rm st} = \sqrt{q_{\rm mh}}\,|11\rangle^{\delta t}_{\rm rf}|001\rangle^{\delta t}_{\rm st} + \sqrt{q_{\rm mm}}\,|00\rangle^{\delta t}_{\rm rf}|010\rangle^{\delta t}_{\rm st} + \sqrt{q_{\rm ml}}\,|01\rangle^{\delta t}_{\rm rf}|100\rangle^{\delta t}_{\rm st}\,, \tag{36}$$

where the first and second subscripts of the transition probabilities denote the start and end nodes. To assemble the gate $\mathcal{D}_{\rm ir}$ we put together a "read" operator $\mathcal{R}_t = \mathcal{R}^{\rm l}_t\mathcal{R}^{\rm m}_t\mathcal{R}^{\rm h}_t$, which reads the state $|\psi\rangle^t_{\rm st}$ and encodes the transition probabilities to $|\psi\rangle^t_{\rm rf}$. As an example, the operator $\mathcal{R}^{\rm m}_{\delta t}$ is:

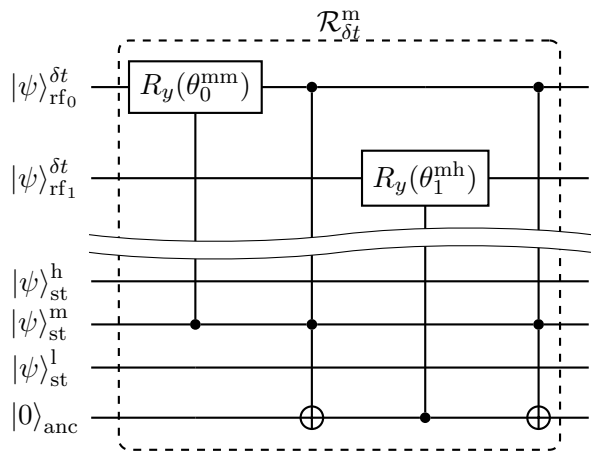

where

$$\theta_0^{\rm mm} = 2\arcsin\sqrt{q_{\rm mm}} \tag{37}$$

$$\theta_1^{\rm mh} = 2\arcsin\sqrt{q_{\rm mh}/(1-q_{\rm mm})}\,. \tag{38}$$

Here, the controlled $R_y(\theta_0^{\rm hm})$ gate encodes the probability $q_{\rm mm}$ — namely the likelihood that the rate will remain at the mid level $r_{\delta t} = r_0 = b$ — into $|\psi\rangle^{\delta t}_{{\rm rf}_0}$. The second gate, AND, checks whether both $r_0 = b$ and $r_{\delta t} \neq b$ are true — namely whether the interest rate changed — and writes the result to an "ancilla" qubit. The third gate, controlled $R_y(\theta_1^{\rm mh})$, checks the result stored in the "ancilla" qubit, and if it is $|1\rangle$ it encodes the conditional probability of $r_{\delta t}$ transitioning to the high interest rate value given that it did not remain at the mid level: $P(r_{\delta t} = b+\delta r | r_{\delta t} \neq b) = P(r_{\delta t} = b+\delta r)/P(r_{\delta t} \neq b) = q_{\rm mh}/(1-q_{\rm mm})$. The fourth gate, AND, ensures the "ancilla" qubit is in its original state. The logic for the $\mathcal{R}^{\rm h}_{\delta t}$ and $\mathcal{R}^{\rm l}_{\delta t}$ gates is similar, see their decomposition in Fig. 5 which shows an example of the entire "read" gate at the first timestep, $\mathcal{R}_{\delta t}$.

The "write" gates $\mathcal{W}_t$ read the "risk factor" qubit pair $|\psi\rangle^t_{\rm rf}$ and write the result to the "state" qubits, $|\psi\rangle^t_{\rm st}$:

$$\mathcal{W}_t\,|11\rangle^t_{\rm rf}\,|000\rangle_{\rm st} = |11\rangle^t_{\rm rf}\,|001\rangle^t_{\rm rf}\,, \tag{39}$$

$$\mathcal{W}_t\,|00\rangle^t_{\rm rf}\,|000\rangle_{\rm st} = |00\rangle^t_{\rm rf}\,|010\rangle^t_{\rm rf}\,, \tag{40}$$

$$\mathcal{W}_t\,|01\rangle^t_{\rm rf}\,|000\rangle_{\rm st} = |01\rangle^t_{\rm rf}\,|100\rangle^t_{\rm rf}\,, \tag{41}$$

see Fig. 5 for its decomposition of $\mathcal{W}_t$ into X and AND operators. Before writing to the "state" qubits, the previous state needs to be erased; therefore, the application of $\mathcal{W}_t$ is always preceded by the gate $\mathcal{W}^\dagger_{t-\delta t}$ which resets the "state" qubits back to $|000\rangle_{\rm st}$. The inverses of the "read" and "write" gates, $\mathcal{R}^\dagger_t$ and $\mathcal{W}^\dagger_t$, consist of the components of $\mathcal{R}_t$ and $\mathcal{W}_t$ put in inverse order. Finally, the operator $\mathcal{D}_{\rm ir}$ can be assembled by a sequence of "read" and "write" operators for each timestep as shown in Fig. 6.

2.2.2 $\mathcal{M}_{\rm ir}$: risk measures $F(r_t)$

For the implementation of the trinomial tree we consider as risk measure the probability of the interest rate being equal to its long-term mean level after m timesteps. In this case, the gate $\mathcal{M}_{\rm m}$ is simply:

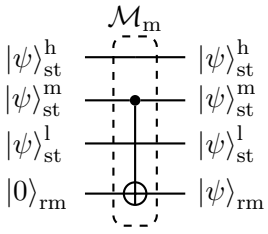

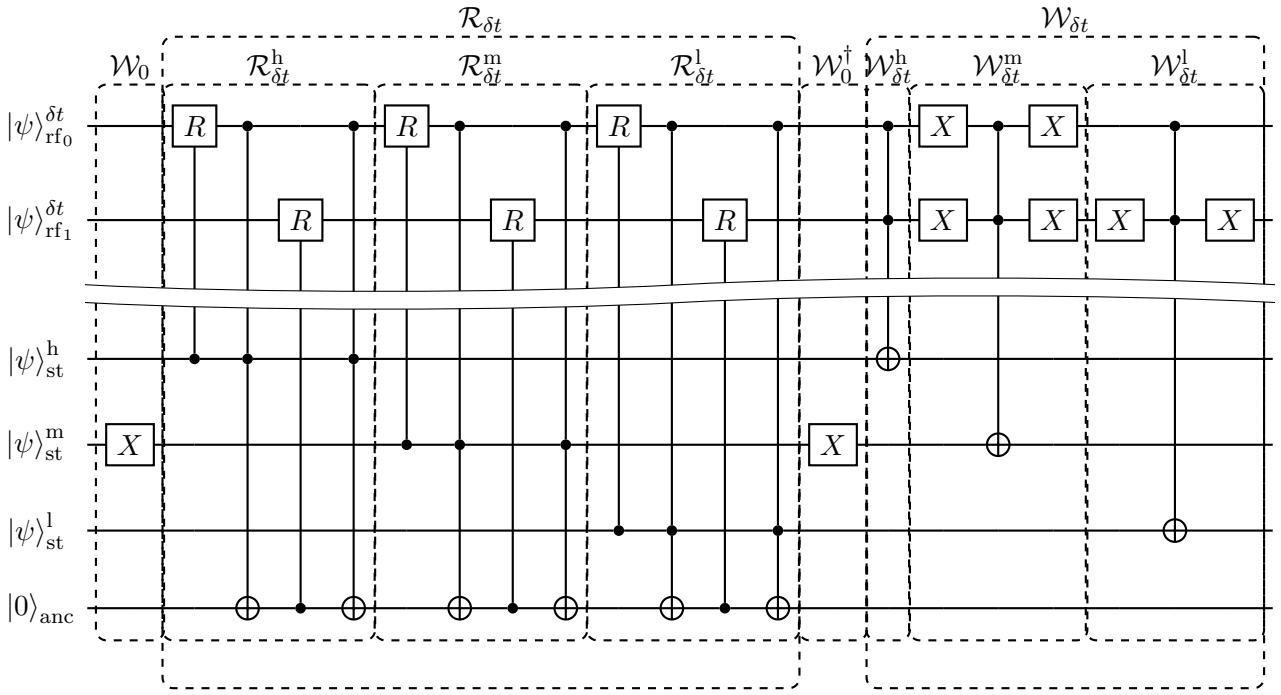

Figure 5: The decomposition of the "read" $\mathcal{R}_t$ and "write" $\mathcal{W}_t$ operators for the first timestep, assuming that $r_0 = b$ (i.e. $\mathcal{W}_0 |000\rangle_{\mathrm{st}} = |010\rangle^0_{\mathrm{st}}$). From left to right the R gates are $R_y(\theta_0^{\mathrm{hm}})$, $R_y(\theta_1^{\mathrm{hh}})$, $R_y(\theta_0^{\mathrm{mm}})$, $R_y(\theta_1^{\mathrm{mh}})$, $R_y(\theta_0^{\mathrm{lm}})$, and $R_y(\theta_1^{\mathrm{lh}})$.

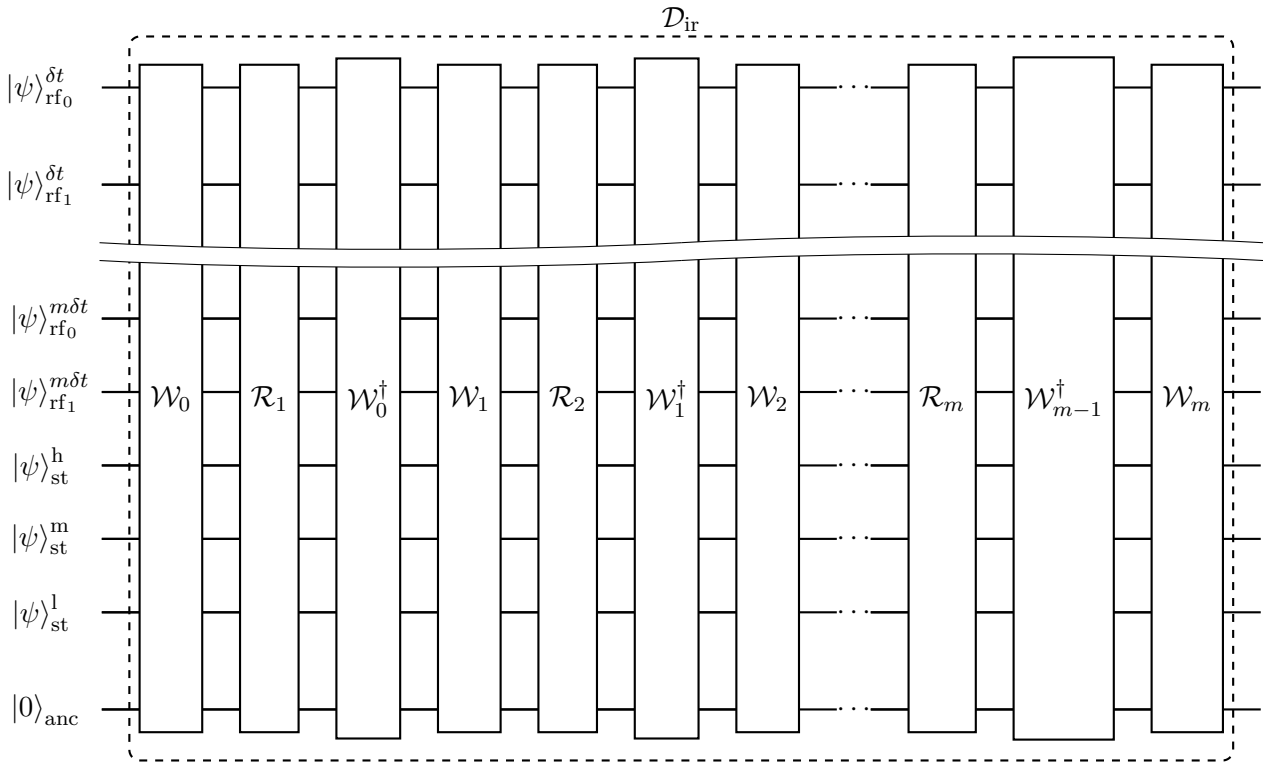

Figure 6: Gate $\mathcal{D}_{\mathrm{ir}}$ consists of a sequence of "read" and "write" gates.

Parameter	Value
m	3
n	1-9
δr	$\sqrt{3Var}$
δt	$3a/12$

Table 3: List of the trinomial tree parameters for the interest rate risk factor evolution.

$q_{t\to t+\delta t}$	$r_{t+\delta t}$		
r_t	high	mid	low
high	$\frac{19}{24}$	$\frac{4}{24}$	$\frac{1}{24}$
mid	$\frac{4}{24}$	$\frac{16}{24}$	$\frac{4}{24}$
low	$\frac{1}{24}$	$\frac{4}{24}$	$\frac{19}{24}$

Table 4: The transition probabilities of the trinomial tree.

with $\mathcal{M}_{\rm m}^\dagger = \mathcal{M}_{\rm m}$.

2.2.3 Results

Table 3 lists the choice of parameters for a quantum circuit that implements trinomial trees for interest rate evolution (due to our parametrization, we do not need to choose values for a and b). For each one of the $m = 3$ timesteps we need 2 "risk factor" qubits, a total of 6: $|\psi\rangle_{\rm rf}$. We need an additional 3 qubits to store the interest rate state $|\psi\rangle_{\rm st}$, 1 qubit for the risk measure $|\psi\rangle_{\rm rm}$, and 9 "ancilla" qubits $|\psi\rangle_{\rm anc}$ for the $\mathcal{A}$ gates. We can freely choose δr, which we set to a multiple of the standard deviation. Note that the definition $0 \leq q_{t\to t+\delta t} \leq 1$ constrains the choice of δt (see App. D). Having defined δr and δt, we can calculate the transition probabilities, see Table 3.

When starting from $r_0 = b$, the probability of measuring the value $r_{3\delta t} = b$ after 3 timesteps can be calculated by adding up the probabilities of all possible paths:

$$P(r_{3\delta t} = b) = \sum_{s_1} \left[q_{{\rm m}s_1} \left(\sum_{s_2} q_{s_1 s_2}\, q_{s_2 {\rm m}} \right) \right] , \quad (42)$$

where s_1, $s_2 \in \{{\rm h}, {\rm m}, {\rm l}\}$ are the nodes of the first and second timesteps, respectively.

Figure 7 shows the probability distribution of the "state" qubits (top panel) and the distribution of the "risk measure" qubit (bottom panel)

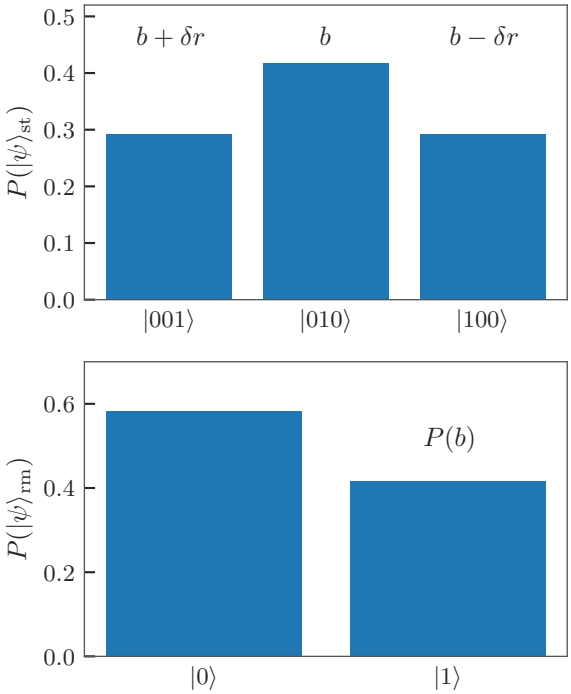

Figure 7: The distribution $P(r_{3\delta t})$ (top panel) and superposition of the "risk factor" qubit that encodes $P(r_{3\delta t} = b)$ (bottom panel).

at $t = 3\delta t$ when measuring these qubits directly with 10,000 shots.

Similar to the equity risk factor analysis, Fig. 8 shows the distribution of the output qubits for $n = 4$ (top panel) and the convergence of the measured value of $p = P(r_{3\delta t} = b)$ as a function of the number of output qubits (bottom panel).

2.3 Credit risk factors

This section focuses on simulating default probabilities of issuers. Among the common approaches are structural and reduced-form credit risk models, as well as credit rating migration.

2.3.1 Structural credit risk models

In structural credit risk models, e.g. the Merton model [41], default takes place when the assets of a company become less than its liabilities. As a simple example, we assume that the liabilities D_t have maturity at time T, and that the value of the assets A_t follows the stochastic process:

$$dA_t = \mu A_t dt + \sigma A_t dW_t \,, \quad (43)$$

where μ is the mean, σ the volatility, and W_t a Wiener process. Default occurs if at time $t = T$

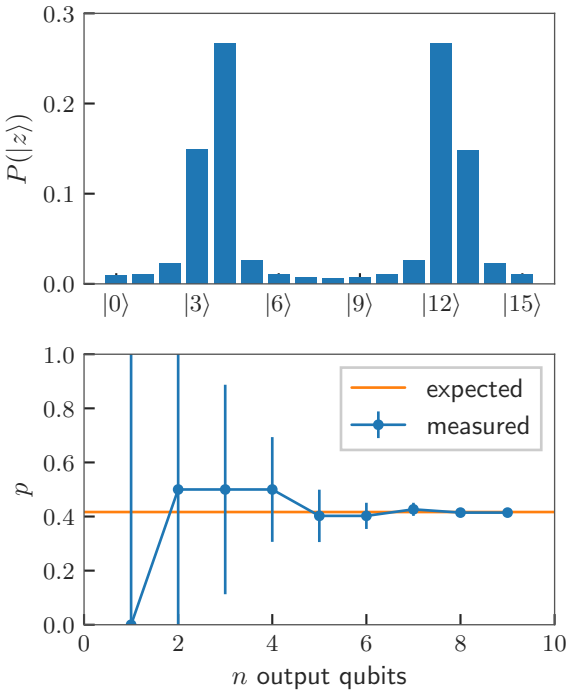

Figure 8: Top panel: the distribution $P(|z\rangle)$ of the measured output state for $n = 4$ and $10,000$ shots. Bottom panel: the measured value of p as a function of the number of output qubits for 1 shot.

the value of the assets is less than, or equal to, the liabilities, $A_T \leq D_T$; therefore, the probability of default is $P(A_T \leq D_T)$.

$\mathcal{D}_{\text{def}}$: the distribution $P(A_t)$ Since this problem has a similar stochastic process to that of equity risk factors, we can model the evolution with a binomial tree using the gate $\mathcal{D}_{\text{def}} = \mathcal{D}_{\text{eq}}$ described in Sect. 2.1. As an example, we consider the same binomial tree parameters as those listed in Table 2 and set $D_T = A_0 d^4$. Based on the top panel of Fig. 3, default occurs if $j = 0$ or $j = 1$ (a "down" price move at all timesteps or a single "up" move before time T).

$\mathcal{M}_{\text{def}}$: probability of default The risk measure here is the probability of default: $P(A_T \leq D_T)$. While in the equity risk factor examples the gates $\mathcal{M}_{\text{max}}$ and $\mathcal{M}_{\text{min}}$ pick up a single value of $|\psi\rangle_{\text{rf}}$ ($|11...1\rangle_{\text{rf}}$ and $|00...0\rangle_{\text{rf}}$, respectively), here we assemble the gate $\mathcal{M}_{\text{def}}$ such that it flips the "risk factor" qubit for any value of $|\psi\rangle_{\text{rf}}$ that satisfies $A_T \leq D_T$; namely, if/when the "risk factor" qubits takes one of the following 7 values: $|000000\rangle$, $|000001\rangle$, $|000010\rangle$, ..., $|100000\rangle$.

For the implementation of the gate we also include "count" qubits, $|\psi\rangle_{\text{c}} = |j\rangle_{\text{c}}$, in order to store the binary value of j. For 6 "risk factor" qubits we have $j \in \{0, ..., 6\}$, and thus 3 "count" qubits are needed to encode these 7 values. We also include "state" qubits, $|\psi\rangle_{\text{st}}$, to represent each of the values of j that correspond to default; since there are two such values here, $j = 0$ and $j = 1$, we include two "state" qubits.

The gate $\mathcal{M}_{\text{def}}$ can be decomposed into three parts. The first part counts the number of up moves in $|\psi\rangle_{\text{rf}}$, i.e. it computes j, and writes it to $|\psi\rangle_{\text{c}}$. The second part reads the state $|\psi\rangle_{\text{c}}$ and flips the corresponding "state" qubit when that value is $j = 0$ or $j = 1$. The third part consists of an OR gate than takes as input the two "state" qubits and flips the "risk measure" qubit if any of them is in the state $|1\rangle$.

The first part can be achieved with the gate $\mathcal{C} = \mathcal{C}_m ... \mathcal{C}_2 \mathcal{C}_1$:

$$|j\rangle_{\text{c}} = \mathcal{C} \, |000\rangle_{\text{c}} \, , \tag{44}$$

where each operator $\mathcal{C}_l$ increments the count by one if the "risk factor" qubit of timestep l is an "up" move: $|\psi\rangle_{\text{rf}}^{l\delta t} = |1\rangle$.

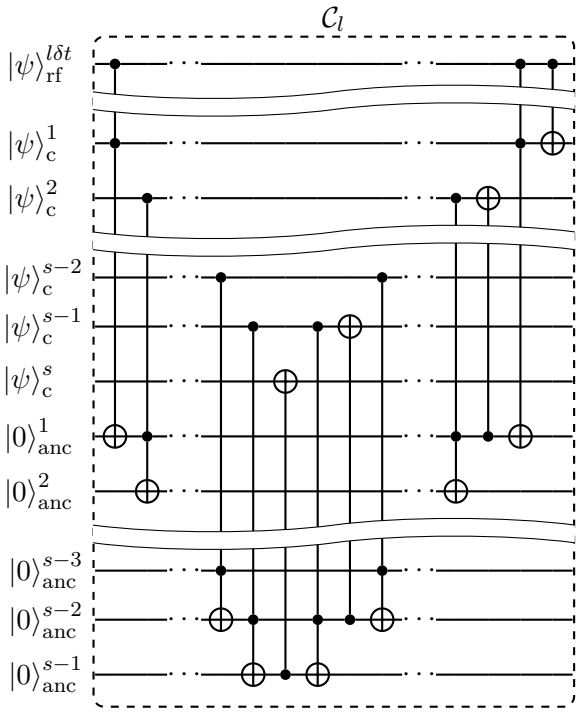

After the operator $\mathcal{C}$ counts the "up" moves (j in total) and writes them as a binary number to $|\psi\rangle_{\text{c}}$, the operator $\mathcal{J} = \mathcal{J}_m ... \mathcal{J}_1 \mathcal{J}_0$ reads j and flips the corresponding qubit of the "state" qubits $|\psi\rangle_{\text{st}}$, see Fig. 9. For our example, we are inter-

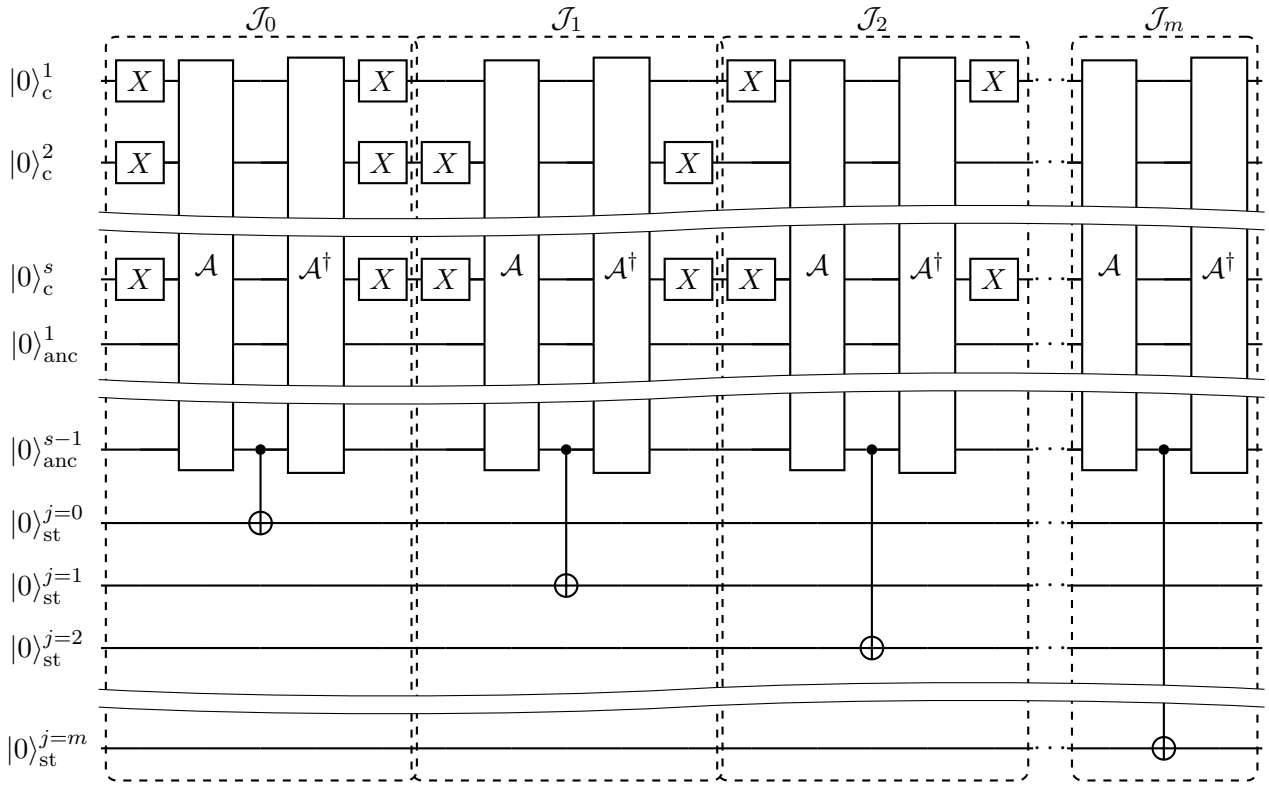

Figure 9: The gates $\mathcal{J}_l$ that flip the "state" qubit $|0\rangle_{\rm st}^{j=l}$ if $|\psi\rangle_{\rm c} = |l\rangle_{\rm c}$

ested in the values $j = 0$ and $j = 1$ so we only include gates $\mathcal{J}_0$ and $\mathcal{J}_1$. Therefore, the components of gate $\mathcal{M}_{\rm def}$ are:

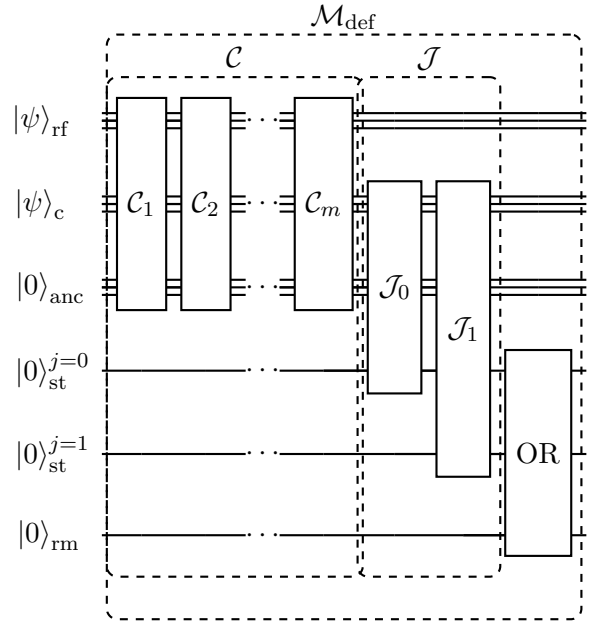

Results The top panel of Fig. 10 shows the probability distribution of the "count" qubits and the bottom panel that of the "risk measure" qubit. Since the binomial tree parameters are the same

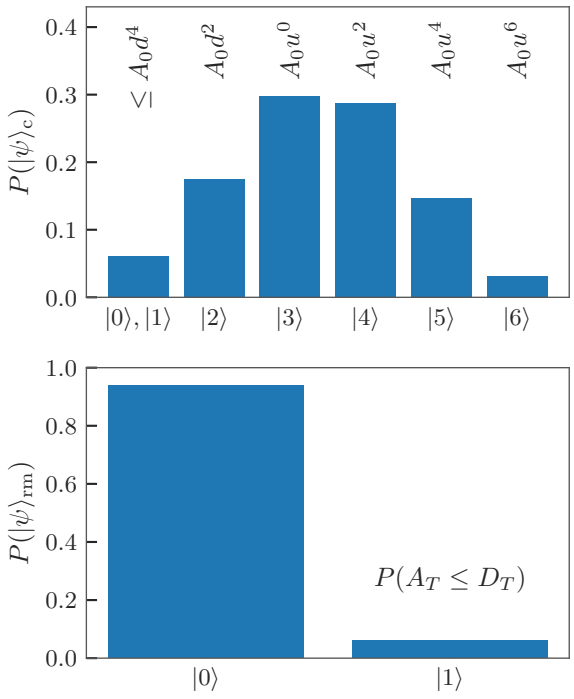

Figure 10: Top panel: the distribution of A_T as generated from 10,000 shots; the first bar represents the two "default" states. Bottom panel: the distribution of the "risk measure" qubit representing the probability of default.

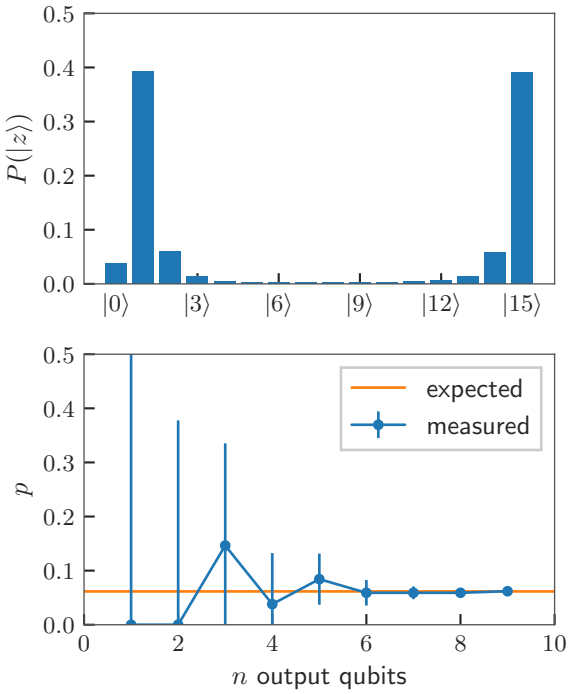

Figure 11: The measured probability distribution $P(|z\rangle)$ (top panel) and the convergence of the estimated probability of default when using one shot (bottom panel).

as in the equity risk factor example, the top panel is identical to Fig. 3 apart from the stacking of the "default" states $|0\rangle$ and $|1\rangle$ (the probability of which is encoded in the "risk measure" qubit shown in the bottom panel). Figure 11 shows the measurement of the output qubits and the convergence of the probability distribution when assembling all gates into a quantum circuit. The expected value of the probability is

$$\begin{aligned} P(j=0)+P(j=1) &= \\ &= \sum_{j=0}^{1} \frac{m!}{(m-j)!j!} q^j (1-q)^{m-j} \\ &= (1-q)^m + mq(1-q)^{m-1} \end{aligned} \tag{45}$$

2.3.2 Reduced-form credit risk models

Another approach to estimate probabilities of default is reduced-form credit risk models, in which default is modelled as a statistical process. The survival probability from time t_0 to time t is given by:

$$P(t_0+t) = e^{-t/T_{\text{def}}}\,, \tag{46}$$

where T_{def} is a characteristic timescale, the inverse of which is called the hazard rate, $\lambda = 1/T_{\text{def}}$.

Parameter	Value
m	6
n	1-9
T	1
q_{def}	2%
$\theta_{\text{def}}\frac{180^\circ}{\pi}$	∼16.3°

Table 5: List of parameters for the reduced-form credit risk model.

To calculate the survival probability at time T we can discretise the time interval $t \in [0, T]$ with m timesteps, $\delta t = T/m$, such that the survival probability is:

$$P(T) = e^{-\sum \delta t/T} = \prod_{j=1}^{m} P(\delta t)\,. \tag{47}$$

A visualisation of this process is:

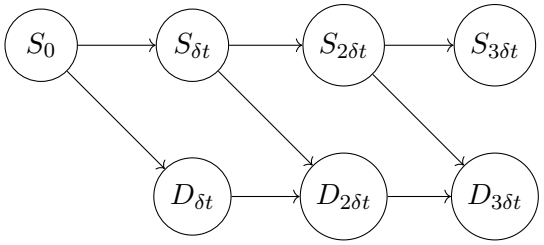

where S_t and D_t denote the survival or default states at time t, respectively.

$\mathcal{D}_{\text{surv}}$: the distribution $P(T)$ At each of the m timesteps we can represent survival with one "risk factor" qubit, such that $|0\rangle$ represents survival and $|1\rangle$ default. Thus, if the company has survived until time $t = (l-1)\delta t$, then $|\psi\rangle_{\text{rf}}^{l\delta t} = \sqrt{1-q_{\text{def}}}\,|0\rangle + \sqrt{q_{\text{def}}}\,|1\rangle$, where $q_{\text{def}} = 1 - P(\delta t)$. If the company has defaulted at a previous timestep, then: $|\psi\rangle_{\text{rf}}^{l\delta t} = |1\rangle$. For the example of $m = 3$:

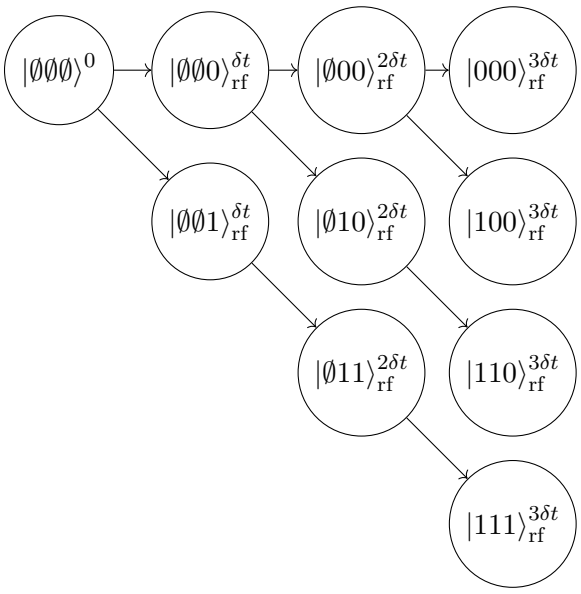

Gate $\mathcal{D}_{\mathrm{surv}}$ can be assembled with rotation gates $R_y(\theta_{\mathrm{def}})$ (shown below as θ_{d} for brevity) obtained from $q_{\mathrm{def}} = \sin(\theta_{\mathrm{def}}/2)$. In addition, in order to ensure that in the event of default at time t the company will also be in default at $t + \delta t$, we include a controlled gate $R_y(\theta^{\mathrm{c}}_{\mathrm{def}})$, where $\theta^{\mathrm{c}}_{\mathrm{def}} + \theta_{\mathrm{def}} = \pi$: this breaks the superposition of the qubit representing $t + \delta t$ and sets it to the state $|1\rangle$. Therefore, the gate $\mathcal{D}_{\mathrm{surv}}$ is:

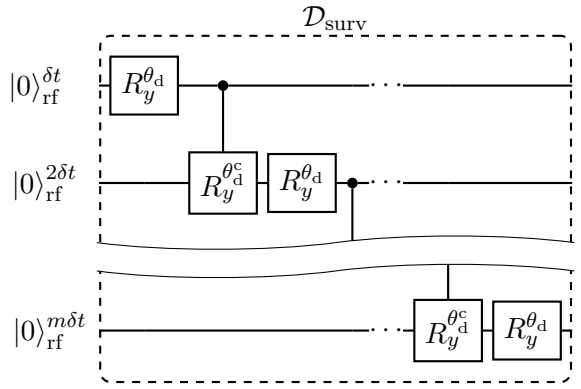

The inverse gate, $\mathcal{D}^{\dagger}_{\mathrm{surv}}$, has the rotation gates in reverse order and the signs of the rotation angles flipped.

$\mathcal{M}_{\mathrm{surv}}$: survival probability The probability of survival at $t = T$ is represented by the state $|0...00\rangle^{T}_{\mathrm{rf}}$. Therefore, we can use the gate $\mathcal{M}_{\mathrm{surv}} = \mathcal{M}_{\mathrm{min}}$ decribed in Sect 2.1, which flips the "risk measure" qubit if all "risk factor" qubits are in the state $|0\rangle$. The expected value of the probability of survival at time T is $(1 - q_{\mathrm{def}})^m$.

Results Table 5 lists the parameters of the quantum circuit, and Fig. 12 the survival ($|1\rangle_{\mathrm{rm}}$) and default ($|0\rangle_{\mathrm{rm}}$) probabilities encoded in the "risk measure" qubit. Figure 13 shows the measurement of the output state and its convergence similarly to the previous use cases.

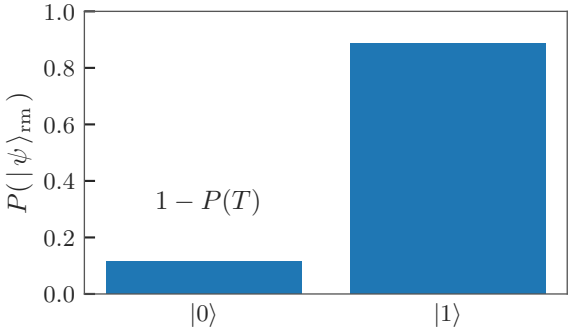

Figure 12: The probabilities of the "risk measure" qubit states: $|1\rangle$ denotes survival up to $t = T$ and $|0\rangle$ denotes default at any timestep before reaching T.

2.3.3 Credit rating migration

A more elaborate approach to simulate probabilities of default is with credit rating migration matrices. For example, consider the following three ratings: A (investment grade), B (high yield), and D (defaulted), and assume that, at $t = 0$, an issuer has the rating A_0. The evolution of the credit rating can be described with multinomial trees, such as this:

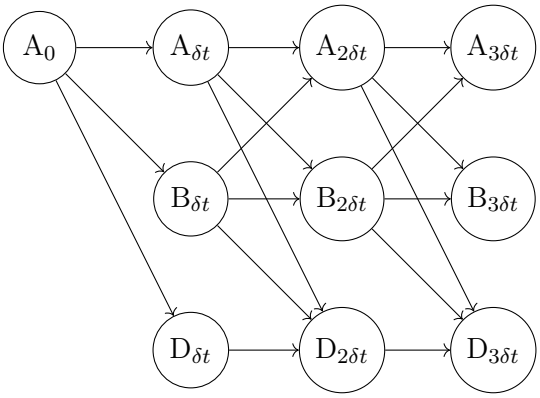

where each transition, e.g. $q_{\mathrm{A}\to\mathrm{B}}$ has a predefined probability which is calibrated with historical data.

Since we have already described the implementation of such trees in Sect. 2.2, we do not provide a more detailed example here. The credit rating evolution and probability of default at time t can be modelled with the gate $\mathcal{D}_{\mathrm{migr}} = \mathcal{D}_{\mathrm{ir}}$ (with a minor adjustment to set the initial rating to A,

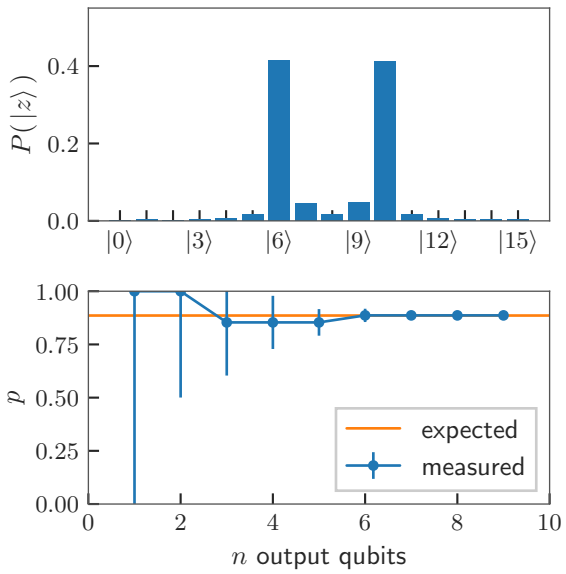

Figure 13: Top panel: the probability distribution of the output qubits for $m = 4$ and 10,000 shots. Bottom panel: the estimated probability (one shot) and its convergence when increasing the number of output qubits.

which would correspond to the high interest rate level), and the default probability can be encoded with $\mathcal{M}_{\text{def}} = \mathcal{M}_{\text{l}}$ (corresponding to measuring the low interest rate level).

3 Discussion

We proceed to estimate the number of qubits required, and the resulting circuit depth, for typical financial risk use cases. We limit the analysis to scenario generation and ignore the additional qubits and gates that would be needed when, for example, the pricing of a derivative is also included in the circuit.

Since the number of "risk factor" qubits, m_{rf}, determines the number of time steps that discretises a time period, we can consider m_{rf} to represent "risk model accuracy". Similarly, because the number of output qubits, n, determines the error of the estimate, we consider n to represent "measurement precision". The higher the number of input and output qubits is, the higher the accuracy and precision are, respectively.

Financial risk models ignore intra-day changes, thus the smallest timestep is one business day. For scenario generation over long time horizons, such as those that span decades, timesteps up to one month can be a sufficient as well as an efficient choice.

Equity price scenarios are used in pricing exchange-traded equity derivatives, most of which expire in less than 12 months, with some extending up to 2 years. With ∼260 business days in a year, the $\mathcal{D}_{\text{eq}}$ gate would not require more than ∼520 "risk factor" qubits. Equity scenarios are also used in equity indices derivatives, structural credit risk models, convertible bonds, and over-the-counter equity derivatives, all of which can have time horizons from few months to several years. Even for periods as large as 50 years, and by considering a timestep of one month, the $\mathcal{D}_{\text{eq}}$ gate has an upper limit of 600 qubits. Because the number of "ancilla" qubits, m_{anc}, for the gates $\mathcal{M}_{\text{min}}$ and $\mathcal{M}_{\text{max}}$ scales linearly with m_{rf}, these gates require roughly: $m_{\text{anc}} \simeq m_{\text{rf}}$. This is not the case for the gate $\mathcal{M}_{\text{def}}$, which only needs few additional qubits ("count", "state", and "ancilla") as m_{rf} increases; e.g., only 8 qubits are needed to count from 1 to 256. Overall, the total number of input qubits ("risk factor", "risk measure", "ancilla", "count", "state") for typical use cases of equity risk factor scenarios has an upper limit of $\sim 1,200$.

Reduced-form credit risk models can be used to simulate the survival probability of an entity over several years, whether for credit default swaps or counterparty credit risk use cases. Similar to structural credit risk models, a timestep of one month can be sufficient for long time horizons; this limits the number of "risk factor" qubits of the $\mathcal{D}_{\text{surv}}$ gate to be a few hundred at most (e.g. 600 for 50 years). Because $\mathcal{M}_{\text{surv}}$ needs a number of "ancilla" qubits similar to m_{rf}, the total number of input qubits for this type of circuits would not exceed $\sim 1,200$ either.

For interest rate scenarios, the number of "risk factor" qubits has a linear dependence on the number of timesteps as in the case of equity scenarios. Because interest rates do not fluctuate as much on a daily basis as equities, a timestep of one month is often adequate to model their evolution. On the other hand, by increasing the number of nodes per timestep, the number of qubits needed to model the discrete interest rate values also increases (based on the implementation of the $\mathcal{D}_{\text{ir}}$ gate here). The scaling is sublinear though, as each new additional qubit can double the amount of interest rate values modelled. Moreover, the $\mathcal{M}_{\text{m}}$ gate of our example

only needs one additional qubit, the "risk measure" one. While we refrain from a more detailed estimate, modelling a 50-year period would not require more than a thousand input qubits. The same analysis also applies to multinomial trees for the modelling of credit rating migration.

In terms of precision, we can require that $\delta p \sim 1\,\text{bp} = 0.01\% = 10^{-4}$,[9] which can be achieved with $n \simeq 14$ output qubits (based on Eq. 19).

Therefore, when accounting for both input and output qubits, scenario generation for typical financial risk use cases would overall require around or a little more than a thousand qubits for high-accuracy, high-precision results.

Next, we proceed to estimate the circuit depth. With current quantum computers, the higher the number of gates the more prone the computations are to errors due to noise. We calculate the depth numerically by decomposing the gates or circuits of Sect. 2 into their constituent gates, a process that we repeat until no further decomposition is possible.[10] The left panels of Fig. 14 show the depth of gates $\mathcal{D}$ and $\mathcal{M}$ as a function of the number of input qubits m (model accuracy); the depth is found to scale linearly with m. Since for $m \simeq 10$ the depth is $\lesssim 500$, for $m \sim 1,000$ we expect the gate depth to be on the order of $\lesssim 50,000$.

The right panels of Fig. 14 show the total circuit depth (it includes the gates $\mathcal{D}$, $\mathcal{M}$, $\prod \mathcal{Q}$, QFT, and $\text{QFT}^\dagger$) in log scale as a function of the number of output qubits n (model precision). The depth increases exponentially with respect to n because every extra output qubit doubles the number of times the $\mathcal{Q}$ gate is applied (see Eq. 14). For $n = 14$, the total depth is projected to be on the order of $\lesssim 10^8$. Estimates from pricing exotic derivatives report qubit numbers and circuit depths that are on the same order of magnitude, $\sim 10^3$-10^4 and $\sim 10^8$, respectively [14, 26].

While quantum computers are expected to reach $\sim 1,000$ qubits in the next few years, the estimated circuit depth prevents such use cases from being feasible without significant progress on fault tolerance. Moreover, note that all the above estimates apply when modelling a single risk factor. For a portfolio that depends on multiple risk factors, we would need that many times of input qubits (parallel run), or to generate scenarios sequentially (serial run) [7]. Hardware limitation make the latter option more likely in the near term.

4 Summary and conclusions

QAE algorithms have been shown to provide a quadratic speedup over their classical counterparts, a result that has motivated several recent papers on financial risk applications. Most of these papers start from pre-computed risk factor distributions and focus on the calculation of common risk measures, distribution loading, and QAE optimisations. In this paper, we extend these studies by consistently integrating scenario generation — which we call QMC — into QAE quantum circuits. Specifically, we assemble quantum gates that implement stochastic risk models for the evolution of equity, interest rate, and credit risk factors. For equities, we generate scenarios by discretising a geometric Brownian motion with binomial trees, and for interest rates, by discretising mean-reversion stochastic differential equations with bounded trinomial trees. For credit risk factors, we calculate the default probability from structural models based on binomial trees and the survival probability from reduced-form models based on a Poisson process. We also describe how multinomial tress can be used for the implementation of credit rating migration matrices. Moreover, we assemble quantum gates to encode generic risk measures, such as the probabilities of measuring the minimum, maximum, and the tail of a distribution. For each use case, we build end-to-end QMC/QAE circuits that incorporate: the generation of risk factor scenarios (QMC), the encoding of the risk measure, and the estimation of the risk measure value (QAE). We then validate the quantum computation with simulated runs, and demonstrate that the measured value converges to the expected value and that the error goes to zero as the number of output qubits increases.

For the typical model accuracy and measure-

[9] "bp" stands for basis points.

[10] For the numerical calculation of the gate or circuit depth we use the Qiskit functions `decompose()` and `depth()` [42]. For example, the depth of a 2-qubit circuit in which each qubit is going through an H gate, $H|0\rangle \otimes H|0\rangle = |{+}{+}\rangle$, has a gate count of 2 and a circuit depth of 1 because the two gates act in parallel. An example of decomposition is breaking down the Toffoli gate into the basic gates it consists of, which gives a gate depth of 11.

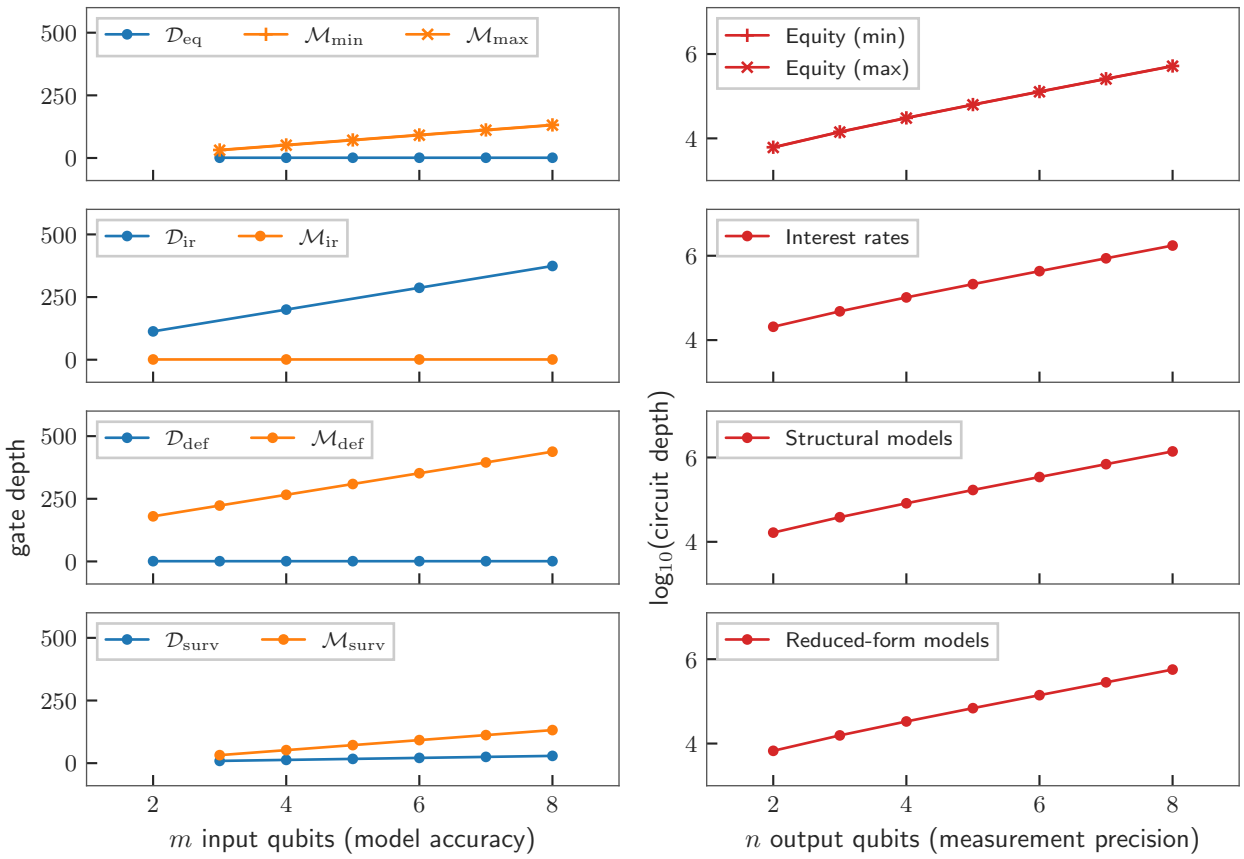

Figure 14: Left panels: depth of the gates $\mathcal{D}$ and $\mathcal{M}$ as a function of the number of input qubits, m, which represents model accuracy. Right panels: depth (in log scale) of the entire quantum circuit as a function of the number of output qubits, n, which represents measurement precision.

ment precision of realistic financial use cases, we estimate that risk factor evolution requires $\lesssim 1,200$ qubits, which is within the capabilities of the quantum computers expected in the next few years. However, we estimate that the typical circuit depth is on the order of $\lesssim 10^8$, which poses the biggest challenge until fault-tolerant quantum devices are available.

The probabilistic output of quantum gates combined with the quantum property of superposition provide a natural framework for the implementation of stochastic risk models. On the one hand, quantum gates can model the output states of a random variable by encoding the probabilities of the outcomes. On the other hand, superposition enables the quantum circuit to simultaneously model all possible paths of a time-dependent random variable, eliminating the classical-computation constraint of iterating over paths. We conclude that quantum financial risk applications can benefit from consistently incorporating scenario generation as part of QMC/QAE simulations, further reducing their dependency on classical computers.

Acknowledgements

We are grateful to three anonymous referees for their detailed reports that helped improve the results and presentation of the paper. The quantum computations were performed on a classical computer using a zero-noise simulator from Qiskit (version `0.24.2`), an open-source framework for quantum computing [42].

References

[1] Román Orús, Samuel Mugel, and Enrique Lizaso. "Quantum computing for finance: Overview and prospects". Reviews in Physics **4**, 100028 (2019).

[2] Daniel J. Egger, Claudio Gambella, Jakub Marecek, Scott McFaddin, Martin Mevissen, Rudy Raymond, Andrea Simonetto, Stefan Woerner, and Elena Yndurain. "Quantum computing for finance: State-of-the-art and future prospects". IEEE Transactions on Quantum Engineering **1**, 1–24 (2020).

[3] Andrés Gómez, Alvaro Leitao Rodriguez, Alberto Manzano, Maria Nogueiras, Gustavo Ordóñez, and Carlos Vázquez. "A survey on quantum computational finance for derivatives pricing and var". Archives of Computational Methods in Engineering **29**, 4137–4163 (2022).

[4] Dylan Herman, Cody Googin, Xiaoyuan Liu, Alexey Galda, Ilya Safro, Yue Sun, Marco Pistoia, and Yuri Alexeev. "A Survey of Quantum Computing for Finance" (2022). arXiv:2201.02773.

[5] Sascha Wilkens and Joe Moorhouse. "Quantum computing for financial risk measurement". Quantum Information Processing**22** (2023).

[6] Philip Intallura, Georgios Korpas, Sudeepto Chakraborty, Vyacheslav Kungurtsev, and Jakub Marecek. "A Survey of Quantum Alternatives to Randomized Algorithms: Monte Carlo Integration and Beyond" (2023). arXiv:2303.04945.

[7] Alexander M. Dalzell, Sam McArdle, Mario Berta, Przemyslaw Bienias, Chi-Fang Chen, András Gilyén, Connor T. Hann, Michael J. Kastoryano, Emil T. Khabiboulline, Aleksander Kubica, Grant Salton, Samson Wang, and Fernando G. S. L. Brandão. "Quantum algorithms: A survey of applications and end-to-end complexities" (2023). arXiv:2310.03011.

[8] Stefan Woerner and Daniel J. Egger. "Quantum risk analysis". npj Quantum Information **5**, 15 (2019).

[9] D. J. Egger, R. Garcia Gutierrez, J. Cahue Mestre, and S. Woerner. "Credit risk analysis using quantum computers". IEEE Transactions on ComputersPages 1–1 (5555).

[10] Kazuya Kaneko, Koichi Miyamoto, Naoyuki Takeda, and Kazuyoshi Yoshino. "Quantum speedup of Monte Carlo integration with respect to the number of dimensions and its application to finance". Quantum Information Processing **20**, 185 (2021).

[11] Patrick Rebentrost, Brajesh Gupt, and Thomas R. Bromley. "Quantum computational finance: Monte carlo pricing of financial derivatives". Phys. Rev. A **98**, 022321 (2018).

[12] Nikitas Stamatopoulos, Daniel J. Egger, Yue Sun, Christa Zoufal, Raban Iten, Ning Shen, and Stefan Woerner. "Option Pricing using Quantum Computers". Quantum **4**, 291 (2020).

[13] Almudena Carrera Vazquez and Stefan Woerner. "Efficient state preparation for quantum amplitude estimation". Phys. Rev. Appl. **15**, 034027 (2021).

[14] Shouvanik Chakrabarti, Rajiv Krishnakumar, Guglielmo Mazzola, Nikitas Stamatopoulos, Stefan Woerner, and William J. Zeng. "A Threshold for Quantum Advantage in Derivative Pricing". Quantum **5**, 463 (2021).

[15] João F. Doriguello, Alessandro Luongo, Jinge Bao, Patrick Rebentrost, and Miklos Santha. "Quantum algorithm for stochastic optimal stopping problems with applications in finance" (2021). arXiv:2111.15332.

[16] Hao Tang, Anurag Pal, Lu-Feng Qiao, Tian-Yu Wang, Jun Gao, and Xian-Min Jin. "Quantum Computation for Pricing the Collateralized Debt Obligations" (2020). arXiv:2008.04110.

[17] Javier Alcazar, Andrea Cadarso, Amara Katabarwa, Marta Mauri, Borja Peropadre, Guoming Wang, and Yudong Cao. "Quantum algorithm for credit valuation adjustments". New Journal of Physics **24**, 023036 (2022).

[18] Jeong Yu Han and Patrick Rebentrost. "Quantum advantage for multi-option portfolio pricing and valuation adjustments" (2022). arXiv:2203.04924.

[19] Nikitas Stamatopoulos, Guglielmo Mazzola, Stefan Woerner, and William J. Zeng. "Towards Quantum Advantage in Financial Market Risk using Quantum Gradient Algorithms". Quantum **6**, 770 (2022).

[20] John Preskill. "Quantum Computing in the NISQ era and beyond". Quantum **2**, 79 (2018).

[21] Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. "Quantum amplitude amplification and estimation". Quantum Computation and InformationPages 53–74 (2002).

[22] Lov Grover and Terry Rudolph. "Creating superpositions that correspond to efficiently integrable probability distributions" (2002). arXiv:quant-ph/0208112.

[23] Steven Herbert. "No quantum speedup with grover-rudolph state preparation for quantum monte carlo integration". Phys. Rev. E **103**, 063302 (2021).

[24] Christa Zoufal, Aurélien Lucchi, and Stefan Woerner. "Quantum generative adversarial networks for learning and loading random distributions". npj Quantum Information **1**, 103 (2019).

[25] Junxu Li and Sabre Kais. "A universal quantum circuit design for periodical functions". New Journal of Physics **23**, 103022 (2021).

[26] Nikitas Stamatopoulos and William J. Zeng. "Derivative Pricing using Quantum Signal Processing" (2023). arXiv:2307.14310.

[27] Sam McArdle, András Gilyén, and Mario Berta. "Quantum state preparation without coherent arithmetic" (2022). arXiv:2210.14892.

[28] Ashley Montanaro. "Quantum speedup of monte carlo methods". Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences **471**, 20150301 (2015).

[29] Michael B. Giles. "Multilevel monte carlo methods". Acta Numerica **24**, 259–328 (2015).

[30] Dong An, Noah Linden, Jin-Peng Liu, Ashley Montanaro, Changpeng Shao, and Jiasu Wang. "Quantum-accelerated multilevel Monte Carlo methods for stochastic differential equations in mathematical finance". Quantum **5**, 481 (2021).

[31] John C. Hull. "Options, futures, and other derivatives". Pearson. (2021). 11th ed., pearson global ed. edition.

[32] Lov K. Grover. "A fast quantum mechanical algorithm for database search". In Gary L. Miller, editor, Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996. Pages 212–219. ACM (1996).

[33] Yohichi Suzuki, Shumpei Uno, Rudy Raymond, Tomoki Tanaka, Tamiya Onodera, and Naoki Yamamoto. "Amplitude estimation without phase estimation". Quantum Information Processing**19** (2020).

[34] Dmitry Grinko, Julien Gacon, Christa Zoufal, and Stefan Woerner. "Iterative quantum amplitude estimation". npj Quantum Information**7** (2021).

[35] Kirill Plekhanov, Matthias Rosenkranz, Mattia Fiorentini, and Michael Lubasch.

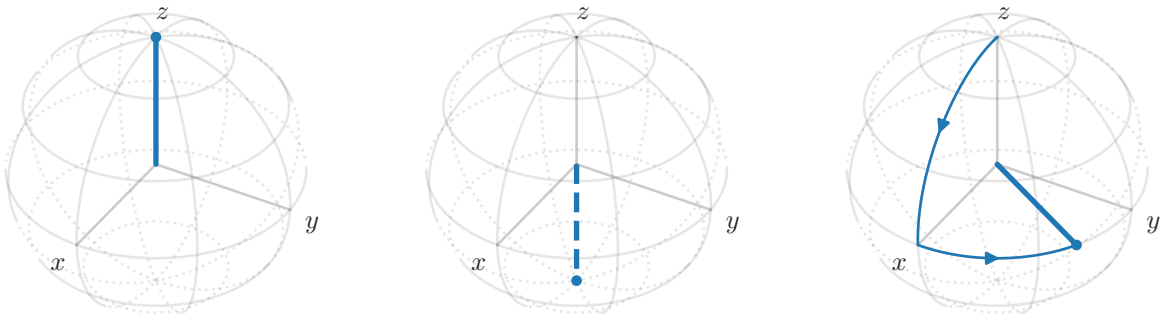

$|0\rangle$ $X|0\rangle = |1\rangle$ $R_z^{\pi/3} R_y^{\pi/2} |0\rangle = |\psi\rangle$

Figure 15: Left: the quantum state $|0\rangle$ on the Bloch sphere. Middle: the quantum state $|1\rangle$, which can be obtained by acting on $|0\rangle$ with the NOT operator, X. The dashed line indicates that the point is on the non-visible part of the sphere's surface. Right: the superposition $|\psi\rangle = \cos(\pi/4)\,|0\rangle + e^{i\pi/3}\sin(\pi/4)\,|1\rangle$ obtained by the operation of two rotation gates.

"Variational quantum amplitude estimation". Quantum **6**, 670 (2022).

[36] John C. Cox, Stephen A. Ross, and Mark Rubinstein. "Option pricing: A simplified approach". Journal of Financial Economics **7**, 229–263 (1979).

[37] Vlatko Vedral, Adriano Barenco, and Artur Ekert. "Quantum networks for elementary arithmetic operations". Phys. Rev. A **54**, 147–153 (1996).

[38] David Oliveira and Rubens Ramos. "Quantum bit string comparator: Circuits and applications". Quantum Computers and Computing**7** (2007).

[39] Various authors. "Qiskit textbook". Github. (2023). url: github.com/Qiskit/textbook.

[40] Oldrich Vasicek. "An equilibrium characterization of the term structure". Journal of Financial Economics **5**, 177–188 (1977).

[41] Robert C. Merton. "On the pricing of corporate debt: the risk structure of interest rates". The Journal of Finance **29**, 449–470 (1974).

[42] "Qiskit: An open-source framework for quantum computing" (2021).

[43] John C Hull and Alan D White. "Numerical procedures for implementing term structure models i". The Journal of Derivatives **2**, 7–16 (1994).

A Qubits, gates, and quantum circuits

Quantum computers consist of qubits, which are the equivalent of classical bits. The quantum states $|0\rangle$ and $|1\rangle$, which represent the values of 0 and 1, are shown on the Bloch sphere in the left and middle panels of Fig. 15, respectively.[11] The key property of qubits is that they can exist in any superposition of these two states, i.e. $|\psi\rangle = a_0\,|0\rangle + a_1\,|1\rangle$, where a_0, a_1 are complex numbers satisfying the condition $|a_0|^2 + |a_1|^2 = 1$. When such a qubit is measured, it will give either the state $|0\rangle$ or $|1\rangle$ with probabilities a_0^2 and a_1^2, respectively.[12] Due to the probabilistic nature of the measured outcome, repeated experiments — called shots — are often needed to improve the precision of a result. Quantum states of single qubits can be conveniently written using the real-valued parameters $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi]$:

$$|\psi\rangle = \cos(\theta/2)\,|0\rangle + e^{i\phi}\sin(\theta/2)\,|1\rangle\ , \tag{48}$$

where θ determines the probability of each state and ϕ is the relative phase (we have ignored the global phase).

[11] $|0\rangle$ and $|1\rangle$ are orthonormal, i.e. $\langle 0|0\rangle = 1$, $\langle 1|1\rangle = 1$, and $\langle 0|1\rangle = \langle 1|0\rangle = 0$.

[12] Measurement collapses the quantum state to a single state, destroying the superposition. If the qubit were to be measured again, it would be found in the same state.

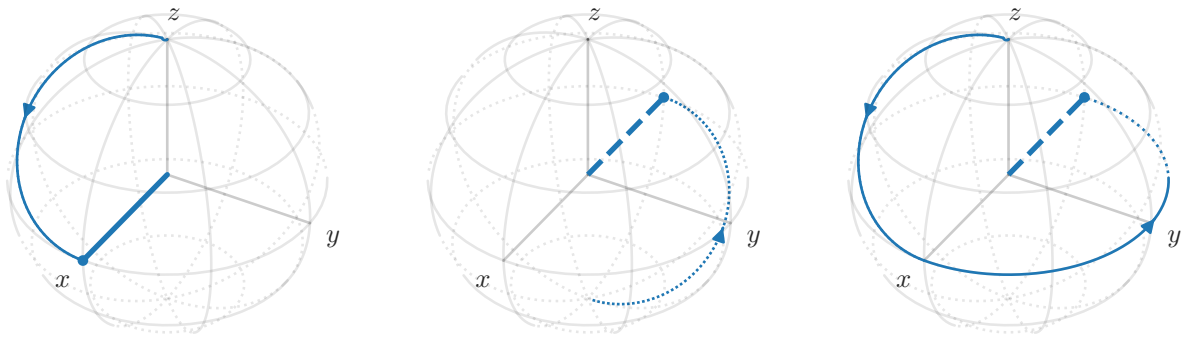

Figure 16: Left and middle: the quantum states $|+\rangle = H|0\rangle$ and $|-\rangle = H|1\rangle$, where H is the Hadamard gate rotating the qubits by π along the $x = z$ axis. Right: the quantum state $|-\rangle = ZH|0\rangle$, where Z rotates the qubit by π along the z-axis, changing its state from $|+\rangle$ to $|-\rangle$.

Quantum circuits are composed of quantum gates operating on qubits. Because qubits can be put in a superposition of states, quantum gates process multiple states simultaneously, when classical gates can only process one at a time.[13] Single-qubit gates modify the state of a qubit — for instance, see the middle panel of Fig. 15 that shows the X, or NOT, operator flipping $|0\rangle$ to output $|1\rangle$. Generally, single-qubit gates modify the parameters θ and ϕ. An example is shown in the right panel of Fig. 15, where the operator $R_y^{\pi/2}$ rotates the qubit $|0\rangle$ by $\theta = \pi$ around the y-axis, and then $R_z^{\pi/3}$ rotates it by $\phi = \pi/3$ around the z-axis. The quantum circuits for the three qubits of Fig. 15 can be written as

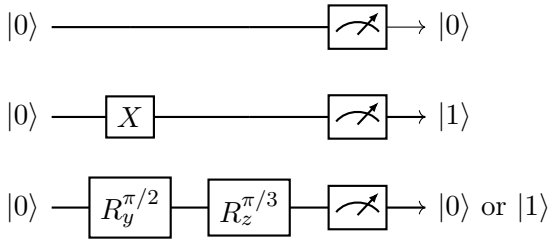

where the third qubit has a 50% chance to be found in either the $|0\rangle$ or $|1\rangle$ state.

Two-qubit gates leverage quantum entanglement, a property that binds the states of two qubits together, and may alter the qubits depending on their initial states. This mechanism allows to implement logical operators similar to an "if" statement, e.g. the CNOT (controlled NOT) gate below:

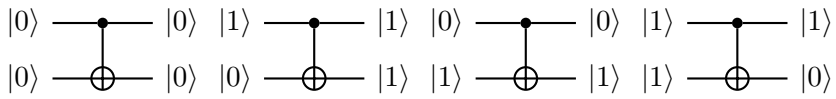

When the "control" qubit (top) is in the state $|1\rangle$, the CNOT gate flips the "target" qubit (bottom). If the "control" qubit is in the state $|0\rangle$, the "target" qubit remains unchanged. Another example with the CNOT gate is the "phase kickback" mechanism, where the phase of the target is "kicked back" to the controller, e.g.

[13] This is one of the key properties that can make quantum computers more efficient than classical ones.

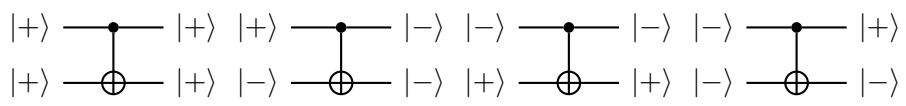

where $|+\rangle$ and $|-\rangle$ is another basis, with the qubits along the x-axis, see Fig. 16.[14]

The CNOT gate is a building block of the three-qubit AND gate (also known as Toffoli or CCNOT gate):

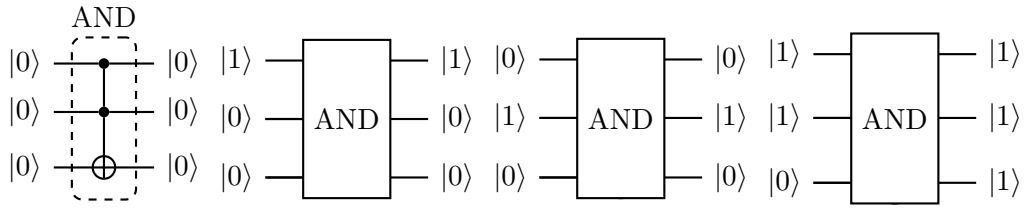

With the help of the NOT and AND gates we can build the logic for an OR gate:

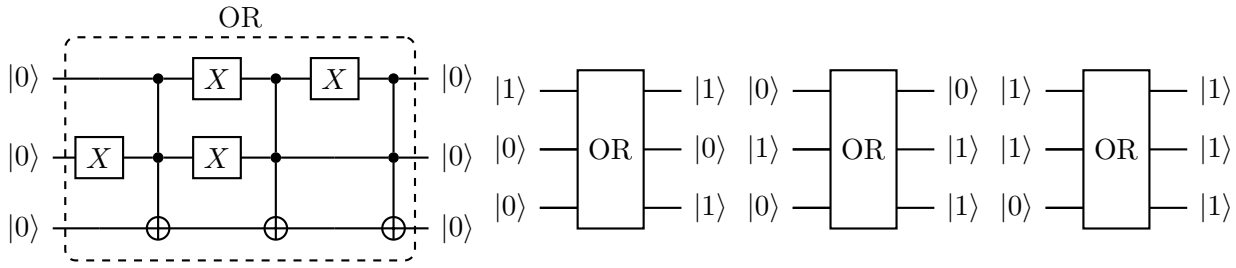

Other multi-qubit gates can be assembled as sub-circuits by combining basic two-qubit gates, however, their depth is currently limited by hardware noise.

B Calculations

In the below we drop the subscript "in" for brevity.

B.1 Calculation of $\mathcal{Q}|\psi\rangle$

From Eq. (7) notice that:

$$\langle\psi_0|\psi_1\rangle = \langle\psi_0|\psi_1\rangle_{\mathrm{rf}}\langle 0|1\rangle_{\mathrm{rm}} = 0 \tag{49}$$
$$\langle\psi_1|\psi_0\rangle = \langle\psi_1|\psi_0\rangle_{\mathrm{rf}}\langle 1|0\rangle_{\mathrm{rm}} = 0 \tag{50}$$

and also

$$\begin{aligned} \langle\psi|\psi\rangle &= 1 \\ (1-p)\langle\psi_0|\psi_0\rangle + p\langle\psi_1|\psi_1\rangle &= 1 \\ \langle\psi_0|\psi_0\rangle + p\left(\langle\psi_1|\psi_1\rangle - \langle\psi_0|\psi_0\rangle\right) &= 1 \end{aligned} \tag{51}$$

which, because it should hold for any p, implies:

$$\langle\psi_0|\psi_0\rangle = \langle\psi_1|\psi_1\rangle = 1 \tag{52}$$

[14] The relation of the two bases is: $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and $|-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$.

Applying $\mathcal{Q}$ (Eq. 8) on $|\psi\rangle$ (Eq. 7) gives:

$$\begin{aligned}
\mathcal{Q}\,|\psi\rangle &= Q_\psi(\mathbb{1} - 2\,|\psi_0\rangle\,\langle\psi_0|\,)(\sqrt{1-p}\,|\psi_0\rangle + \sqrt{p}\,|\psi_1\rangle\,) \\
&= \Big[\mathbb{1} - 2(1-p)\,|\psi_0\rangle\,\langle\psi_0| - 2\sqrt{p(1-p)}\,|\psi_0\rangle\,\langle\psi_1| - 2\sqrt{p(1-p)}\,|\psi_1\rangle\,\langle\psi_0| - 2p\,|\psi_1\rangle\,\langle\psi_1|\Big] \\
&\quad \times \Big(-\sqrt{1-p}\,|\psi_0\rangle + \sqrt{p}\,|\psi_1\rangle\Big) \\
&= -\sqrt{1-p}\,|\psi_0\rangle + \sqrt{p}\,|\psi_1\rangle + 2(1-p)\sqrt{1-p}\,|\psi_0\rangle - 2p\sqrt{1-p}\,|\psi_0\rangle \\
&\quad + 2(1-p)\sqrt{p}\,|\psi_1\rangle - 2p\sqrt{p}\,|\psi_1\rangle \\
&= (1-4p)\sqrt{1-p}\,|\psi_0\rangle + (3-4p)\sqrt{p}\,|\psi_1\rangle\ . \qquad (53)
\end{aligned}$$

B.2 Calculation of $\mathcal{Q}\,|\psi_\pm\rangle$

$$\begin{aligned}
\mathcal{Q}\,|\psi_\pm\rangle &= \frac{1}{\sqrt{2}}\mathcal{Q}_\psi(\mathbb{1} - 2\,|\psi_0\rangle\,\langle\psi_0|\,)(\,|\psi_1\rangle \pm i\,|\psi_0\rangle\,) \\
&= \frac{1}{\sqrt{2}}\Big[\mathbb{1} - 2(1-p)\,|\psi_0\rangle\,\langle\psi_0| - 2\sqrt{p(1-p)}\,|\psi_0\rangle\,\langle\psi_1| - 2\sqrt{p(1-p)}\,|\psi_1\rangle\,\langle\psi_0| - 2p\,|\psi_1\rangle\,\langle\psi_1|\Big] \\
&\quad \times (\,|\psi_1\rangle \mp i\,|\psi_0\rangle\,) \\
&= \frac{1}{\sqrt{2}}\Big[|\psi_1\rangle \mp i\,|\psi_0\rangle \pm i2(1-p)\,|\psi_0\rangle - 2\sqrt{p(1-p)}\,|\psi_0\rangle \pm i2\sqrt{p(1-p)}\,|\psi_1\rangle - 2p\,|\psi_1\rangle\Big] \\
&= \frac{1}{\sqrt{2}}\Big[\pm i(1-2p)\,|\psi_0\rangle - 2\sqrt{p(1-p)}\,|\psi_0\rangle + (1-2p)\,|\psi_1\rangle \pm i2\sqrt{p(1-p)}\,|\psi_1\rangle\Big]\ . \qquad (54)
\end{aligned}$$

But

$$1 - 2p = 1 - 2\sin^2(\theta/2) = \cos\theta\,, \qquad (55)$$

$$2\sqrt{p(1-p)} = 2\sin(\theta/2)\cos(\theta/2) = \sin\theta\,. \qquad (56)$$

Therefore,

$$\begin{aligned}
\mathcal{Q}\,|\psi_\pm\rangle &= \frac{1}{\sqrt{2}}(\,\pm i\cos\theta\,|\psi_0\rangle - \sin\theta\,|\psi_0\rangle + \cos\theta\,|\psi_1\rangle \pm i\sin\theta\,|\psi_1\rangle\,) \\
&= \frac{1}{\sqrt{2}}\cos\theta(\,|\psi_1\rangle \pm i\,|\psi_0\rangle\,) + \frac{1}{\sqrt{2}}\sin\theta(\,\pm i\,|\psi_1\rangle - |\psi_0\rangle\,) \\
&= \cos\theta\,|\psi_\pm\rangle \pm i\sin\theta\,|\psi_\pm\rangle \\
&= e^{\pm i\theta}\,|\psi_\pm\rangle\ , \qquad (57)
\end{aligned}$$

which implies that:

$$\mathcal{Q}^k\,|\psi_\pm\rangle = e^{\pm ik\theta}\,|\psi_\pm\rangle\ . \qquad (58)$$

B.3 Expressing $|\psi\rangle$ as a function of $|\psi_\pm\rangle$

We can write $|\psi_0\rangle$ and $|\psi_1\rangle$ as functions of $|\psi_\pm\rangle$:

$$|\psi_0\rangle = -i\frac{1}{\sqrt{2}}(\,|\psi_+\rangle - |\psi_-\rangle\,)\,, \qquad (59)$$

$$|\psi_1\rangle = \frac{1}{\sqrt{2}}(\,|\psi_+\rangle + |\psi_-\rangle\,)\,. \qquad (60)$$

Therefore,

$$
\begin{aligned}
|\psi\rangle &= \sqrt{1-p}\,|\psi_0\rangle + \sqrt{p}\,|\psi_1\rangle \\
&= \frac{1}{\sqrt{2}}\sqrt{p}(\,|\psi_+\rangle + |\psi_-\rangle\,) - i\frac{1}{\sqrt{2}}\sqrt{1-p}(\,|\psi_+\rangle - |\psi_-\rangle\,) \\
&= -i\frac{1}{\sqrt{2}}\left[i\sin(\theta/2)(\,|\psi_+\rangle + |\psi_-\rangle\,) + \cos(\theta/2)(\,|\psi_+\rangle - |\psi_-\rangle\,)\right] \\
&= -i\frac{1}{\sqrt{2}}\left[\Big(\cos(\theta/2) + i\sin(\theta/2)\Big)\,|\psi_+\rangle - \Big(\cos(\theta/2) - i\sin(\theta/2)\Big)\,|\psi_-\rangle\right] \\
&= -i\frac{1}{\sqrt{2}}(e^{i\theta/2}\,|\psi_+\rangle - e^{-i\theta/2}\,|\psi_-\rangle\,)\,. \qquad (61)
\end{aligned}
$$

B.4 The output qubits after the controlled gate $\mathcal{Q}$

We set the output qubits — which we label here with $l = 0, 1, ..., n-1$ — as controls to the gate $\mathcal{Q}$, and apply the gate 2^l times for each qubit, respectively. Their output state is:

$$
\begin{aligned}
\prod_{l=0}^{n-1} \mathcal{Q}^{2^l}\,|\psi\rangle\,|+\rangle_{\text{out}}^{\otimes n} &= |\psi\rangle \bigotimes_{l=0}^{n-1}\left[\frac{1}{\sqrt{2}}\left(|0\rangle + e^{\pm i2^l\theta}\,|1\rangle\right)\right] \\
&= |\psi\rangle\,\frac{1}{2^{n/2}}\bigotimes_{l=0}^{n-1}\left(e^{i2^l\theta\cdot 0}\,|0\rangle + e^{\pm i2^l\theta\cdot 1}\,|1\rangle\right) \\
&= |\psi\rangle\,\frac{1}{2^{n/2}}\bigotimes_{l=0}^{n-1}\left(\sum_{b_l=0}^{1} e^{\pm i2^l b_l\theta}\,|b_k\rangle\right) \\
&= |\psi\rangle\,\frac{1}{2^{n/2}}\sum_{b_0=0}^{1}\sum_{b_1=0}^{1}\cdots\sum_{b_{n-1}=0}^{1}\prod_{l=0}^{n-1} e^{\pm i2^l b_l\theta}\,|b_0\rangle\,|b_1\rangle\,...\,|b_{n-1}\rangle \\
&= |\psi\rangle\,\frac{1}{2^{n/2}}\sum_{b_0=0}^{1}\cdots\sum_{b_{n-1}=0}^{1} e^{\pm\sum_{l=0}^{n-1} i2^l b_l\theta}\,|b_0 ... b_{n-1}\rangle \\
&= |\psi\rangle\,\frac{1}{2^{n/2}}\sum_{x=0}^{2^n-1} e^{\pm ix\theta}\,|x\rangle\,. \qquad (62)
\end{aligned}
$$

B.5 Inverse Quantum Fourier Transform

The inverse quantum Fourier transform of the state of the output qubits gives:

$$
\begin{aligned}
|\psi\rangle_{\text{out}} &= \text{QFT}^\dagger\left[\frac{1}{2^{n/2}}\sum_{x=0}^{2^n-1} e^{\pm ix\theta}\,|x\rangle\right] \\
&= \frac{1}{2^{n/2}}\sum_{x=0}^{2^n-1} e^{\pm ix\theta}\text{QFT}^\dagger\,|x\rangle \\
&= \frac{1}{2^{n/2}}\sum_{x=0}^{2^n-1} e^{\pm ix\theta}\left(\frac{1}{2^{n/2}}\sum_{z=0}^{2^n-1} e^{-i2\pi xz/2^n}\,|z\rangle\right) \\
&= \frac{1}{2^n}\sum_{z=0}^{2^n-1}\sum_{x=0}^{2^n-1} e^{ix(\pm\theta - 2\pi z/2^n)}\,|z\rangle \\
&= \sum_{z=0}^{2^n-1} a_z\,|z\rangle\,. \qquad (63)
\end{aligned}
$$

B.6 The amplitudes a_z when $2^n\theta/2\pi$ or $2^n(2\pi-\theta)/2\pi$ are integers

In the special case where either $2^n\theta/2\pi$ or $2^n(2\pi-\theta)/2\pi$ is an integer equal to z_0, the amplitude a_{z_0} is:

$$a_{z_0} = \frac{1}{2^n}\sum_{x=0}^{2^n-1} 1 = 1\,, \tag{64}$$

and for all other a_z:

$$\begin{aligned} a_z &= \frac{1}{2^n}\sum_{x=0}^{2^n-1} e^{ix(2\pi z_0/2^n - 2\pi z/2^n)} \\ &= \frac{1}{2^n}\sum_{x=0}^{2^n-1} \left(e^{i2\pi(z_0-z)/2^n}\right)^x \\ &= \frac{1}{2^n}\left[\frac{1-\left(e^{i2\pi(z_0-z)/2^n}\right)^{2^n}}{1-e^{i2\pi(z_0-z)/2^n}}\right] \\ &= \frac{1}{2^n}\left(\frac{1-e^{i2\pi(z_0-z)}}{1-e^{i2\pi(z_0-z)/2^n}}\right) \\ &= 0 \end{aligned} \tag{65}$$

where $e^{i2\pi(z_0-z)} = 1$ because $z_0 - z$ is an integer.

B.7 The general case for the amplitudes $|a_z|$

If neither $2^n\theta/2\pi$ nor $2^n(2\pi-\theta)/2\pi$ is an integer, we can find the closest integer such that $z_0+\epsilon = 2^n\theta/2\pi$ or $z_0+\epsilon = 2^n(2\pi-\theta)/2\pi$, where $\epsilon \in (0, 1/2]$. Namely,

$$\begin{aligned} a_z &= \frac{1}{2^n}\sum_{x=0}^{2^n-1} e^{ix(2\pi z_0/2^n - 2\pi z/2^n + 2\pi\epsilon/2^n)} \\ &= \frac{1}{2^n}\sum_{x=0}^{2^n-1} \left(e^{i2\pi(z_0-z+\epsilon)/2^n}\right)^x \\ &= \frac{1}{2^n}\left[\frac{1-\left(e^{i2\pi(z_0-z+\epsilon)/2^n}\right)^{2^n}}{1-e^{i2\pi(z_0-z+\epsilon)/2^n}}\right] \\ &= \frac{1}{2^n}\left(\frac{1-e^{i2\pi(z_0-z)}e^{i2\pi\epsilon}}{1-e^{i2\pi(z_0-z+\epsilon)/2^n}}\right) \\ &= \frac{1}{2^n}\left(\frac{1-e^{i2\pi\epsilon}}{1-e^{i2\pi(z_0-z+\epsilon)/2^n}}\right) \end{aligned} \tag{66}$$

The probability $|a_z|^2$ is

$$\begin{aligned} |a_z|^2 &= \frac{1}{2^{2n}}\left(\frac{1-e^{i2\pi\epsilon}}{1-e^{i2\pi(z_0-z+\epsilon)/2^n}}\right)\left(\frac{1-e^{-i2\pi\epsilon}}{1-e^{-i2\pi(z_0-z+\epsilon)/2^n}}\right) \\ &= \frac{1}{2^{2n}}\left(\frac{2-e^{i2\pi\epsilon}-e^{-i2\pi\epsilon}}{2-e^{i2\pi(z_0-z+\epsilon)/2^n}-e^{-i2\pi(z_0-z+\epsilon)/2^n}}\right) \\ &= \frac{1}{2^{2n}}\left(\frac{1-\cos(2\pi\epsilon)}{1-\cos[2\pi(z_0-z+\epsilon)/2^n]}\right)\,, \end{aligned} \tag{67}$$

and, for large n, the probability of measuring the closest integer ($z = z_0$) is:

$$\begin{aligned}
|a_{z_0}|^2 &= \frac{1}{2^{2n}} \left(\frac{1-\cos(2\pi\epsilon)}{1-\cos(2\pi\epsilon/2^n)} \right) \\
&\simeq \frac{1}{2^{2n}} \left[1 - \left(1 - \frac{(2\pi\epsilon)^2}{2!} + \frac{(2\pi\epsilon)^4}{4!} \right) \right] \left[1 - \left(1 - \frac{(2\pi\epsilon)^2}{2^{2n}2!} \right) \right]^{-1} \\
&= \frac{1}{2^{2n}} \left[\frac{(2\pi\epsilon)^2}{2!} - \frac{(2\pi\epsilon)^4}{4!} \right] \left[\frac{2^{2n}2!}{(2\pi\epsilon)^2} \right] \\
&= 1 - \frac{\pi^2\epsilon^2}{3} \\
&\geq 1 - \frac{\pi^2}{12}\,,
\end{aligned} \tag{68}$$

where in the last expression we took the maximum value of $\epsilon = 1/2$.

B.8 Calculation of δp for QAE

The error in p can be estimated as follows:

$$\begin{aligned}
\delta p &= \delta \sin^2\left(\frac{\theta}{2}\right) \\
&\simeq 2\sin\left(\frac{\theta}{2}\right)\cos\left(\frac{\theta}{2}\right)\frac{\delta\theta}{2} \\
&= \sin\theta\frac{\pi}{2^n}\,,
\end{aligned} \tag{69}$$

because $\delta\theta = 2\pi/2^n$.

C Equity risk factor calculations

C.1 Binomial tree parameters

By matching the mean of the continuous and discrete models we obtain the value of q:

$$\begin{aligned}
E(S_{t+\delta t}) &= qS^{\mathrm{u}}_{t+\delta t} + (1-q)S^{\mathrm{d}}_{t+\delta t} \\
S_t e^{\mu\delta t} &= qS_t u + (1-q)S_t d \\
q &= \frac{e^{\mu\delta t} - d}{u - d} \\
&= \frac{ue^{\mu\delta t} - 1}{u^2 - 1}
\end{aligned} \tag{70}$$

And from the variance, we obtain the value of u:

$$\begin{aligned}
Var(S_{t+\delta t}) &= E(S^2_{t+\delta t}) - E(S_{t+\delta t})^2 \\
S_t^2 e^{2\mu\delta t+\sigma^2\delta t} - S_t^2 e^{2\mu\delta t} &= q(S_t^2u^2) + (1-q)(S_t^2d^2) - S_t^2 e^{2\mu\delta t} \\
e^{2\mu\delta t+\sigma^2\delta t} &= qu^2 + (1-q)d^2 \\
&= u^2\frac{ue^{\mu\delta t}-1}{u^2-1} + \frac{1}{u^2}\frac{u^2 - ue^{\mu\delta t}}{u^2-1} \\
&= \frac{1}{u(u^2-1)}\left(u^4e^{\mu\delta t} - u^3 + u - e^{\mu\delta t}\right) \\
&= \frac{1}{u(u^2-1)}\left[e^{\mu\delta t}(u^4-1) - u(u^2-1)\right] \\
&= \frac{1}{u}\left[e^{\mu\delta t}(u^2+1) - u\right]\,.
\end{aligned} \tag{71}$$

We consider small timesteps, $\sigma\sqrt{\delta t} \ll 1$ and $\mu\delta t \ll 1$, ignore terms higher than $\mathcal{O}(\delta t)$, and try the solution $u = e^{b\sqrt{\delta t}}$. With Taylor expansion, the terms $e^{\mu\delta t+\sigma^2\delta t}$, u, u^2, and u^{-1} are:

$$
\begin{aligned}
e^{2\mu\delta t+\sigma^2\delta t} &\simeq 1 + 2\mu\delta t + \sigma^2\delta t && (72)\\
u &\simeq 1 + b\sqrt{\delta t} + \frac{1}{2}b^2\delta t && (73)\\
u^2 &\simeq \left(1 + b\sqrt{\delta t} + \frac{1}{2}b^2\delta t\right)^2 \\
&\simeq 1 + b\sqrt{\delta t} + \frac{1}{2}b^2\delta t + b\sqrt{\delta t} + b^2\delta t + \frac{1}{2}b^2\delta t \\
&= 1 + 2b\sqrt{\delta t} + 2b^2\delta t && (74)\\
u^{-1} &= e^{-b\sqrt{\delta t}} \\
&\simeq 1 - b\sqrt{\delta t} + \frac{1}{2}b^2\delta t && (75)
\end{aligned}
$$

Therefore,

$$
\begin{aligned}
1 + 2\mu\delta t + \sigma^2\delta t &\simeq \left(1 - b\sqrt{\delta t} + \frac{1}{2}b^2\delta t\right)\left[(1+\mu\delta t)\left(1 + 2b\sqrt{\delta t} + 2b^2\delta t + 1\right) - \left(1 + b\sqrt{\delta t} + \frac{1}{2}b^2\delta t\right)\right] \\
1 + 2\mu\delta t + \sigma^2\delta t &\simeq \left(1 - b\sqrt{\delta t} + \frac{1}{2}b^2\delta t\right)\left(2 + 2b\sqrt{\delta t} + 2b^2\delta t + 2\mu\delta t - 1 - b\sqrt{\delta t} - \frac{1}{2}b^2\delta t\right) \\
1 + 2\mu\delta t + \sigma^2\delta t &\simeq \left(1 - b\sqrt{\delta t} + \frac{1}{2}b^2\delta t\right)\left(1 + b\sqrt{\delta t} + \frac{3}{2}b^2\delta t + 2\mu\delta t\right) \\
1 + 2\mu\delta t + \sigma^2\delta t &\simeq 1 + b\sqrt{\delta t} + \frac{3}{2}b^2\delta t + 2\mu\delta t - b\sqrt{\delta t} - b^2\delta t + \frac{1}{2}b^2\delta t \\
b &\simeq \sigma && (76)
\end{aligned}
$$

C.2 Decomposition of the $\mathcal{Q}_\psi$ operator

The operator that applies a reflection to the state $|\psi\rangle_{\text{in}}$ can be written as:

$$
\begin{aligned}
\mathcal{Q}_\psi &= \mathbb{1} - 2\,|\psi\rangle_{\text{in}}\,\langle\psi|_{\text{in}} \\
&= \mathbb{1} - 2\mathcal{M}\mathcal{D}\,|0\rangle_{\text{in}}\,\langle 0|_{\text{in}}\,\mathcal{D}^\dagger\mathcal{M}^\dagger \\
&= \mathcal{M}\mathcal{D}(\mathbb{1} - 2\,|0\rangle_{\text{in}}\,\langle 0|_{\text{in}}\,)\mathcal{D}^\dagger\mathcal{M}^\dagger \\
&= \mathcal{M}\mathcal{D}Q_{00}\mathcal{D}^\dagger\mathcal{M}^\dagger\,. && (77)
\end{aligned}
$$

C.3 Example circuit with $m = 2$ and $n = 3$

Below is an example of the gates $\mathcal{D}_{\text{eq}}$, $\mathcal{M}_{\max}$, $\mathcal{Q}$, QFT, and QFT† for a quantum circuit that estimates the probability of measuring the maximum value, $S_{\max}$. The parameters adopted for the binomial tree are $q = 0.3827$, $\mu = 0$, $\sigma = 0$, $T = 1$, and for the quantum circuit are $m = 2$ and $n = 3$.

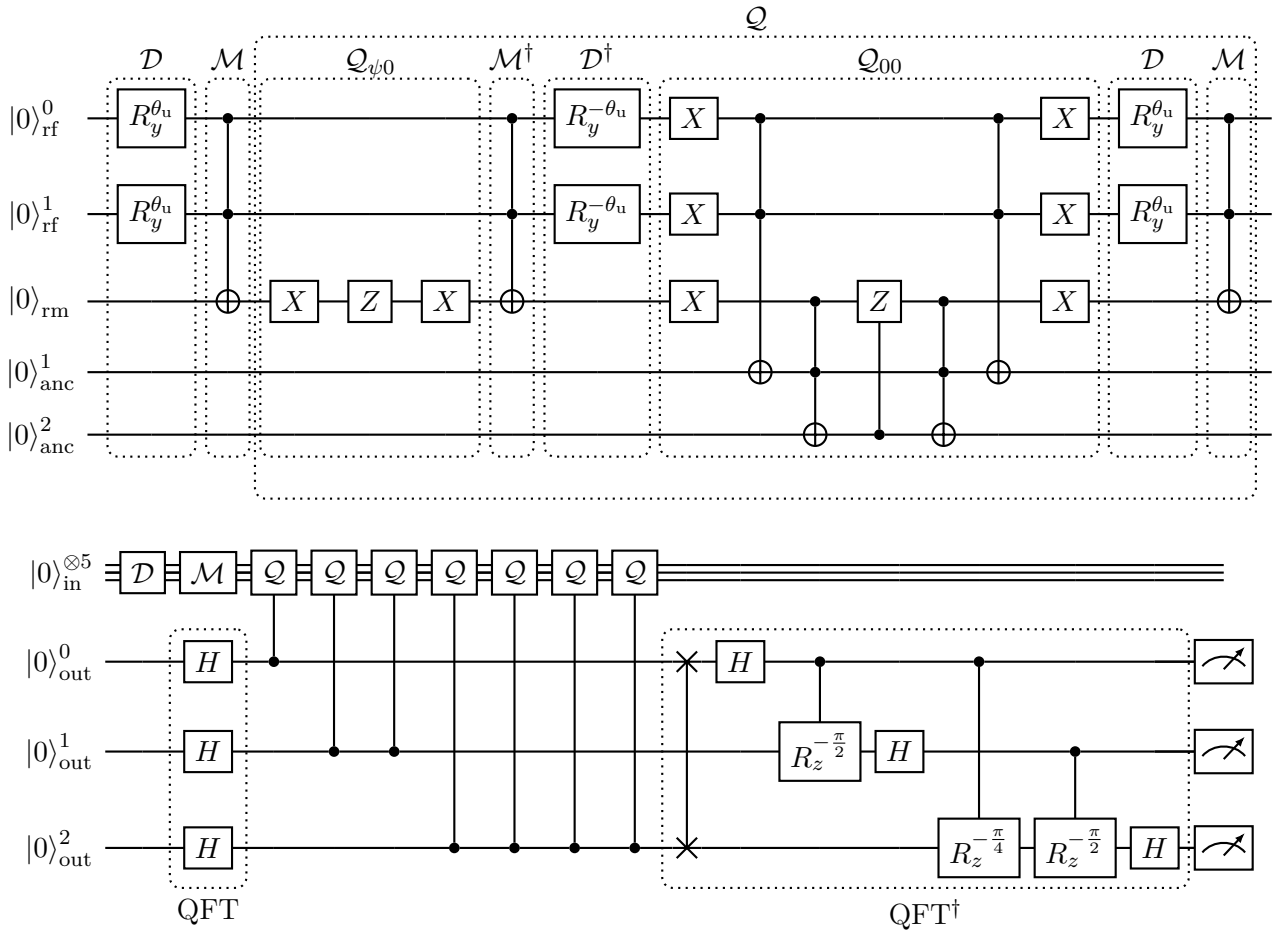

The transformation of the qubit states by the quantum gates is shown in Fig. 17. Specifically, the "risk factor" qubits ($|0\rangle^0_{\rm rf}$ and $|0\rangle^1_{\rm rf}$) before and after the operation of the gate $\mathcal{D}$ are shown in green in the top left panel. The application of gate $\mathcal{M}$ on the "risk measure" qubit ($|0\rangle_{\rm rm}$) is shown in red in the bottom left panel. From left to right, the three vertical panels show the transformation of the output qubits (in blue and orange; $|0\rangle^0_{\rm out}$, $|0\rangle^1_{\rm out}$, and $|0\rangle^2_{\rm out}$) by the QFT, $\prod \mathcal{Q}$, and QFT† gates, respectively.

Figure 18 shows the probability distribution of the "risk factor" (left, in green), "risk measure" (middle, in red), and output (right, in blue/orange) qubits. The distributions of the input qubits are only shown for illustrative purposes, they are not measured in the circuit because this would collapse their states. The measurement of the output qubit gives either the state $|1\rangle$ or $|7\rangle$; from Eqs. (16) and (17), these imply

$$1 = 2^3 \frac{\theta}{2\pi} \tag{78}$$

$$7 = 2^3 \frac{2\pi - \theta}{2\pi} \tag{79}$$

both of which give $\theta = \pi/4$. Therefore, from Eq. (4) we have $p = \sin(\pi/8)^2 = 0.1464$, and since $p = q^2$ (because we require two up moves), we validate the computation by recovering the input value $q = 0.3827$.

D Interest rate risk factor calculations

The sum of the probabilities satisfy: $q_{wu} + q_{wm} + q_{wd} = 1$, so

$$q_{wm} = 1 - q_{wu} - q_{wd}\,. \tag{80}$$

The expected value of the Vasicek model should match that of the trinomial tree. We can parametrise the three possible values of r_t ($b-\delta r$, b, and $b+\delta r$), as $r_t = b + c_t \delta r$, where $c_t = \{-1, 0, 1\}$. Therefore,

$$\begin{aligned} E(r_{t+\delta t}) &= q_{wu}\left(b+\delta r\right) + q_{wm} b + q_{wd}\left(b-\delta r\right) \\ (b + c_t\delta r)e^{-a\delta t} + b(1-e^{-a\delta t}) &= q_{wu} b + q_{wu}\delta r + q_{wm} b + q_{wd} b - q_{wd}\delta r \\ b + c_t e^{-a\delta t}\delta r &= b + (q_{wu} - q_{wd})\delta r \\ q_{wd} &= q_{wu} - c_t e^{-a\delta t} \end{aligned} \tag{81}$$

Similarly, the variance of the continuous and discrete models should match:

$$\begin{aligned} Var(r_{t+\delta t}) &= E(r_{t+\delta t}^2) - E(r_{t+\delta t})^2 \\ \frac{\sigma^2}{2a}(1-e^{-2a\delta t}) &= q_{wu}\left(b+\delta r\right)^2 + q_{wm} b^2 + q_{wd}\left(b-\delta r\right)^2 - (b + c_t e^{-a\delta t}\delta r)^2 \\ \frac{\sigma^2}{2a}(1-e^{-2a\delta t}) &= q_{wu} b^2 + 2q_{wu} b\delta r + q_{wu}\delta r^2 + q_{wm} b^2 + q_{wd} b^2 - 2q_{wd} b\delta r + q_{wd}\delta r^2 \\ &\quad - b^2 - 2bc_t e^{-a\delta t}\delta r - c_t^2 e^{-2a\delta t}\delta r^2 \\ \frac{\sigma^2}{2a}(1-e^{-2a\delta t}) &= 2(q_{wu} - q_{wd}) b\delta r + (q_{wu} + q_{wd})\delta r^2 - 2bc_t e^{-a\delta t}\delta r - c_t^2 e^{-2a\delta t}\delta r^2 \\ \frac{\sigma^2}{2a}(1-e^{-2a\delta t}) &= (q_{wu} + q_{wd} - c_t^2 e^{-2a\delta t})\delta r^2 \end{aligned} \tag{82}$$

By setting $\delta r^2 = 3Var(r_{t+\delta t})$ [e.g. 43] we obtain

$$\begin{aligned} \frac{1}{3} &= q_{wu} + q_{wd} - c_t^2 e^{-2a\delta t} \\ \frac{1}{3} &= 2q_{wu} - c_t e^{-a\delta t} - c_t^2 e^{-2a\delta t} \\ q_{wu} &= \frac{1}{6} + \frac{1}{2} c_t e^{-a\delta t} + \frac{1}{2} c_t^2 e^{-2a\delta t} \\ &\simeq \frac{1}{6} + \frac{1}{2}\left[c_t(1-a\delta t) + c_t^2(1-2a\delta t)\right]. \end{aligned} \tag{83}$$

For $c_t = 0$,

$$q_{mu} = \frac{1}{6} \tag{84}$$

$$q_{mm} = \frac{2}{3} \tag{85}$$

$$q_{md} = q_{wu} \tag{86}$$

For $c_t = 1$,

$$\begin{aligned} q_{uu} &\simeq \frac{1}{6} + 1 - \frac{3}{2} a\delta t \\ &\simeq \frac{7}{6} - \frac{3}{2} a\delta t \end{aligned} \tag{87}$$

$$\begin{aligned} q_{ud} &\simeq q_{wu} - 1 + a\delta t \\ &\simeq \frac{1}{6} - \frac{1}{2} a\delta t \end{aligned} \tag{88}$$

$$\begin{aligned} q_{um} &\simeq 1 - \frac{8}{6} + 2a\delta t \\ &\simeq -\frac{1}{3} + 2a\delta t \end{aligned} \tag{89}$$

And for $c_t = -1$,

$$q_{du} \simeq \frac{1}{6} - \frac{1}{2}a\delta t \tag{90}$$

$$\begin{aligned} q_{dd} &\simeq q_{wu} + 1 - a\delta t \\ &\simeq \frac{7}{6} - \frac{3}{2}a\delta t \end{aligned} \tag{91}$$

$$\begin{aligned} q_{dm} &\simeq 1 - \frac{8}{6} + 2a\delta t \\ &\simeq -\frac{1}{3} + 2a\delta t \end{aligned} \tag{92}$$

When $c_t = \pm 1$ we need to ensure positive probabilities, therefore:

$$0 < \frac{1}{6} - \frac{1}{2}a\delta t < 1 \tag{93}$$

$$0 < \frac{7}{6} - \frac{3}{2}a\delta t < 1 \tag{94}$$

$$0 < -\frac{1}{3} + 2a\delta t < 1 \tag{95}$$

or

$$-\frac{5}{6} < a\delta t < \frac{1}{3} \tag{96}$$

$$\frac{1}{9} < a\delta t < \frac{7}{9} \tag{97}$$

$$\frac{1}{6} < a\delta t < \frac{2}{3} \tag{98}$$

which implies $\frac{1}{6} < a\delta t < \frac{2}{6}$. We choose $\delta t = \frac{3}{12}a$. For $c_t = 1$, this gives

$$q_{uu} \simeq \frac{19}{24} \tag{99}$$

$$q_{ud} \simeq \frac{1}{24} \tag{100}$$

$$q_{um} \simeq \frac{4}{24} \tag{101}$$

and for $c_t = -1$,

$$q_{du} \simeq \frac{1}{24} \tag{102}$$

$$q_{dd} \simeq \frac{19}{24} \tag{103}$$

$$q_{dm} \simeq \frac{4}{24} \tag{104}$$

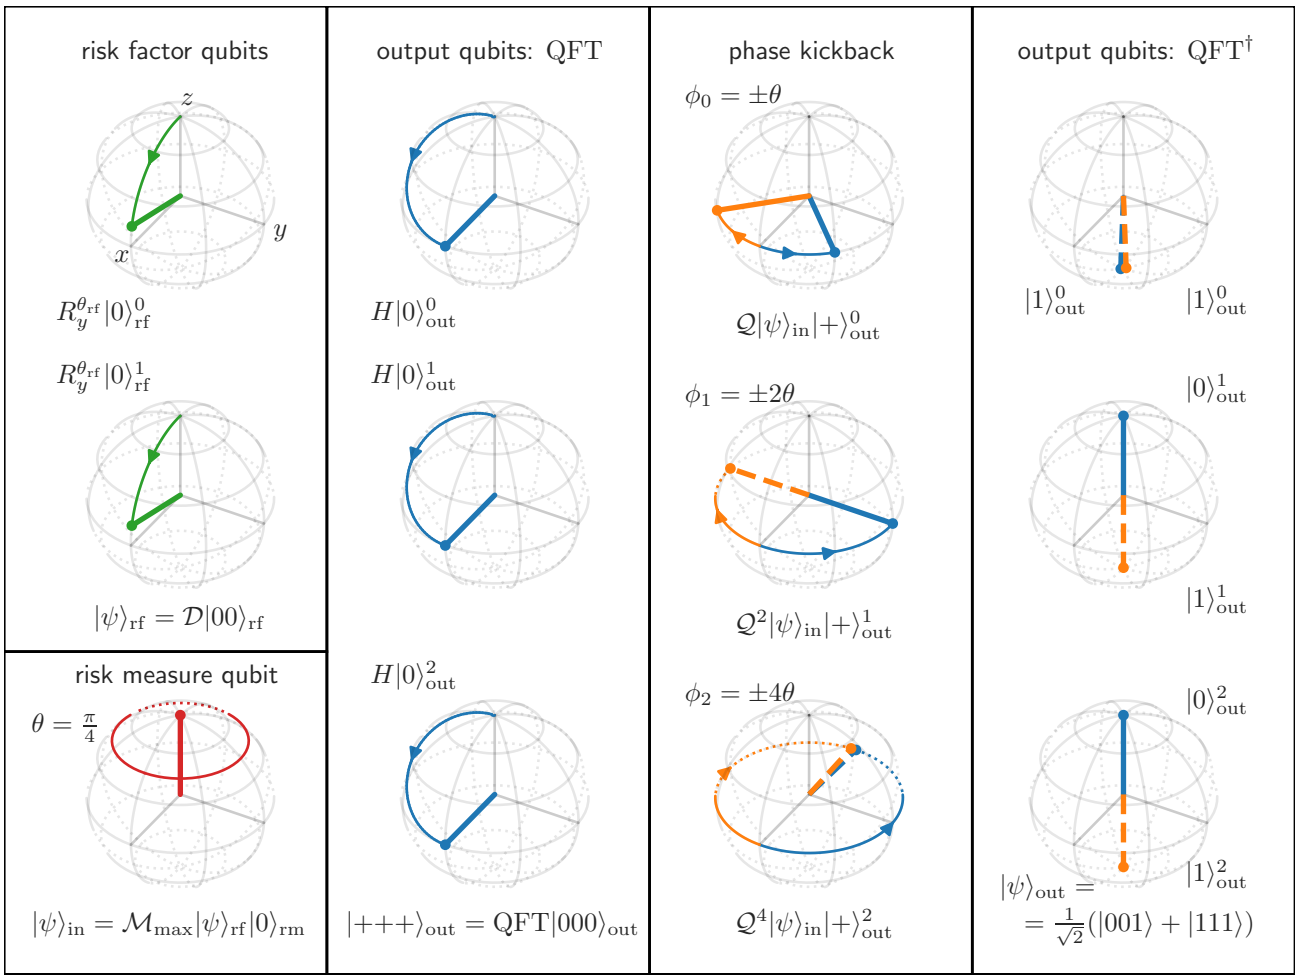

Figure 17: The qubit states before and after the operation of the quantum gates. Top left: gate $\mathcal{D}$ rotates the "risk factor" qubits (green) around the y-axis. Bottom left: the controlled gate $\mathcal{M}$ reads the "risk factor" qubits and sets the "risk measure" qubit in superposition (red), encoding the probability that both "risk factor" qubits are in the state $|1\rangle$ (since no rotation is involved, the phase of the "risk measure" qubit does not have a specific value and we draw the state as a circle). Left vertical panel: the output qubits (blue) before and after the application of the QFT gate, which consists of a Hadamard gate, H, applied to each qubit. Middle vertical panel: the controlled gate $\mathcal{Q}$ imprints the angle θ of the "risk measure" qubit as a positive ($+\phi$, blue) or negative ($-\phi$, orange) phase onto the first output qubit (top); the controlled gate $\mathcal{Q}^2$ imprints the angle 2θ onto the phase of the second qubit (middle); the controlled gate $\mathcal{Q}^4$ imprints the angle 4θ onto the phase of the third qubit (bottom). Right vertical panel: the inverse quantum Fourier transform gate, QFT†, leverages quantum interference to transform the phases of the output qubits to a binary number expressed with $|0\rangle$ and $|1\rangle$ states.

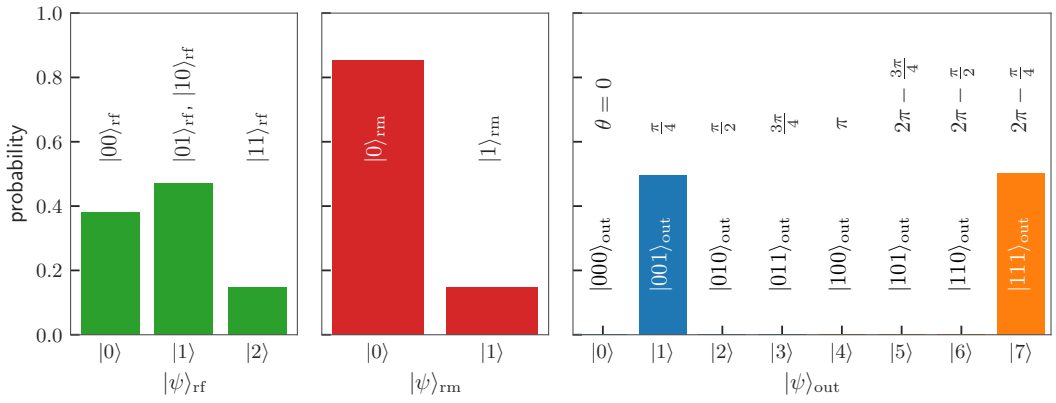

Figure 18: The probability distributions of the risk factor (left), risk measure (middle), and output qubits (right).