Quantum Amplitude Estimation for Expected Loss Computation under a Discretized Latent-Factor Model

Arghya Ghosh
VP, J.P. Morgan Chase, Quantitative Research
Mumbai, India

***Abstract*—We present a quantum algorithm for estimating the expected loss (EL) of a credit portfolio driven by a latent systematic risk factor. We consider a single latent factor model here. The method discretizes a continuous latent variable, encodes its probability distribution into quantum amplitudes, embeds the loss function via controlled rotations on an ancilla qubit, and applies Grover-style amplitude amplification. The resulting state admits a two-subspace decomposition enabling estimation of EL using Quantum Amplitude Estimation(QAE) or a maximum-likelihood estimator (MLE) based on Grover-power measurements, yielding a quadratic speedup over classical Monte Carlo.**

***Index Terms*—Quantum amplitude estimation, Grover operator, credit risk, expected loss.**

I. Introduction

Estimating expected loss (EL) is central in credit risk management for credit portfolios. Sometimes estimating the expected loss with simulation can be time consuming and computationally expensive in the classical domain Classical Monte Carlo (MC) approximates expectations with sampling complexity $O(\varepsilon^{-2})$ for error ε. Quantum Amplitude Estimation (QAE) [1] provides a quadratic improvement, reducing the complexity of the query to $O(\varepsilon^{-1})$ [3].

II. Problem Setting

Let $Z \sim \mathcal{N}(\mu, \sigma^2)$ be a systematic latent risk factor for the PD model. Define the loss

$$L(Z) = \mathrm{EAD} \cdot \mathrm{LGD} \cdot \mathrm{PD}(Z), \tag{1}$$

where $\mathrm{PD}(Z) \in [0,1]$ is a conditional probability of default given Z.
The target quantity is the expected loss

$$\mathbb{E}[L] = \mathbb{E}_Z[L(Z)] = \int_{\mathbb{R}} L(z)\, \varphi_{\mu,\sigma}(z)\, dz, \tag{2}$$

with $\varphi_{\mu,\sigma}$ the Gaussian pdf. This is a simplistic assumption. In the current study we have assumed it to be Gaussian, however, we can deal with different distributions.

[0]The views expressed in this paper are solely those of the author and do not necessarily reflect the views or policies of JPMorgan Chase Co. This work was conducted in a personal capacity and does not contain proprietary or confidential information.

III. Discretization of the Latent Variable

Using n qubits, let $N = 2^n$ and discretize Z on a grid over N grid points $[\mu - 3\sigma, \mu + 3\sigma]$. Recall that within $[\mu - 3\sigma, \mu + 3\sigma]$ it covers almost 99 percent of the grid mass. Somebody can try making changes to the values of $k\sigma$ for different values of k to see sensitives around k.:

$$z_i = \mu - 3\sigma + i\Delta, \quad i = 0, 1, \ldots, N-1, \tag{3}$$

$$\Delta = \frac{6\sigma}{N}. \tag{4}$$

Define discrete probabilities $\{p_i\}_{i=0}^{N-1}$ approximating $\mathbb{P}(Z \in$ bin $i)$. A common robust choice is bin-integrated mass:

$$p_i = \Phi\Big(\frac{z_i + \Delta/2 - \mu}{\sigma}\Big) - \Phi\Big(\frac{z_i - \Delta/2 - \mu}{\sigma}\Big), \qquad \sum_{i=0}^{N-1} p_i \approx 1, \tag{5}$$

where Φ is the standard normal CDF.
The discretized EL becomes

$$\mathbb{E}[L] \approx \sum_{i=0}^{N-1} p_i\, L(z_i). \tag{6}$$

IV. Uncertainty State Preparation

Before we start to work with quantum algorithm it is important that we first encode the current problem in the quantum framework. Define the uncertainty state over the discretized latent factor:

$$|\psi_Z\rangle = \sum_{i=0}^{N-1} \sqrt{p_i}\, |i\rangle\,. \tag{7}$$

The above equation magically captures the information of all loss realization in one equation. This is the essence of quantum superposition. There exists a unitary U(for example, $H^{\otimes n}$) such that

$$U\, |0\rangle^{\otimes n} = |\psi_Z\rangle\,. \tag{8}$$

V. Loss Encoding via Controlled Rotations

Once we are done with encoding the problem, next we need to bring EL information in the equation. Amplitude estimation process requires mapping the expectation to the probability of measuring an ancilla in $|1\rangle$.

A. Normalization

To ensure a valid rotation, define a normalized loss $\tilde{L}(z) \in [0, 1]$. For example, choose a bound $L_{\max} > 0$ and set

$$\tilde{L}(z) = \frac{L(z)}{L_{\max}} \in [0, 1]. \tag{9}$$

Then

$$\mathbb{E}[L] = L_{\max}\, \mathbb{E}[\tilde{L}]. \tag{10}$$

B. Construction of Oracle U_f

We introduce the oracle using rotation gate and each basis states, which is defined as:

$$U_f = \sum_{i=0}^{N-1} |i\rangle\langle i| \otimes R_y(\theta_i), \tag{11}$$

where $R_y(\theta_i)$ is the controlled single qubit rotation gate on the ancilla qubit given by

$$R_y(\theta_i) = \begin{pmatrix} \cos\frac{\theta_i}{2} & -\sin\frac{\theta_i}{2} \\ \sin\frac{\theta_i}{2} & \cos\frac{\theta_i}{2} \end{pmatrix}. \tag{12}$$

Applying U_f to the joint state gives the following

$$U_f\big(\,|\psi\rangle \otimes |0\rangle\,\big) = \sum_{i=0}^{N-1} \sqrt{p_i}\, U_f\big(\,|i\rangle \otimes |0\rangle\,\big) \tag{13}$$

$$= \sum_{i=0}^{N-1} \sqrt{p_i}\,\big(\,|i\rangle \otimes R_y(\theta_i)\,|0\rangle\,\big). \tag{14}$$

Since

$$R_y(\theta_i)\,|0\rangle = \cos\frac{\theta_i}{2}\,|0\rangle + \sin\frac{\theta_i}{2}\,|1\rangle, \tag{15}$$

we obtain

$$|\psi^*\rangle = U_f(|\psi\rangle \otimes |0\rangle) = \sum_{i=0}^{N-1} \sqrt{p_i}\;|i\rangle \otimes \left(\cos\frac{\theta_i}{2}\,|0\rangle + \sin\frac{\theta_i}{2}\,|1\rangle\right). \tag{16}$$

Rearranging terms,

$$|\psi^*\rangle = \left(\sum_{i=0}^{N-1} \sqrt{p_i}\cos\frac{\theta_i}{2}\,|i\rangle\right)|0\rangle + \left(\sum_{i=0}^{N-1} \sqrt{p_i}\sin\frac{\theta_i}{2}\,|i\rangle\right)|1\rangle. \tag{17}$$

Notice one thing how the extra qubit $|0\rangle$ has transformed the problem into a two dimensional space, so that we introduce the Grovers operator. Define

$$|\psi_1\rangle = \sum_{i=0}^{N-1} \sqrt{p_i}\sin\frac{\theta_i}{2}\,|i\rangle. \tag{18}$$

By Born's rule, the probability of measuring the ancilla in state $|1\rangle$ is

$$a = \|\,|\psi_1\rangle\,\|^2 = \sum_{i=0}^{N-1} p_i \sin^2\frac{\theta_i}{2}. \tag{19}$$

Now let us consider the controlled Y-rotation on the ancilla with angle

$$R_y\left(2\arcsin\sqrt{\tilde{L}(z_i)}\right), \tag{20}$$

controlled on $|i\rangle$ (using comparators / multi-controlled rotations like X gates).
Define the state-preparation operator for amplitude estimation:

$$A = U_f\,(U \otimes I). \tag{21}$$

VI. Amplitude Interpretation and Two-Subspace Decomposition

Applying A to the all-zero input yields

$$A\,|0\rangle^{\otimes n}\,|0\rangle = \sum_{i=0}^{N-1} \sqrt{p_i}\,|i\rangle\left(\sqrt{1-\tilde{L}(z_i)}\;|0\rangle + \sqrt{\tilde{L}(z_i)}\;|1\rangle\right)$$

$$= \sqrt{1-a}\;|\Psi_0\rangle\,|0\rangle + \sqrt{a}\;|\Psi_1\rangle\,|1\rangle, \tag{22}$$

where the (normalized) states $|\Psi_0\rangle, |\Psi_1\rangle$ live on the n latent-factor qubits, and

$$a = \sum_{i=0}^{N-1} p_i\,\tilde{L}(z_i) \approx \mathbb{E}[\tilde{L}]. \tag{23}$$

Thus, by Born's rule,

$$\Pr(\text{ancilla} = 1) = a, \tag{24}$$

and the desired expected loss is

$$\mathbb{E}[L] = L_{\max}\, a. \tag{25}$$

so the idea is if we can measure the measure the probability ancillary qubit $|1\rangle$ then we can find the EL. The whole problem statement now has been transformed into two dimensional subspaces spanned by $|\Psi_0\rangle\,|0\rangle\, and\, |\Psi_1\rangle\,|1\rangle$. Let θ be the angle between the basis vectors for some $\theta \in [0, \pi/2]$.

VII. Grover Operator and Phase Amplification

Define reflections (phase flips)

$$S_\chi = I - 2\big(I \otimes |1\rangle\langle 1|\,\big), \tag{26}$$

$$S_0 = I - 2\,|0\cdots 0\rangle\langle 0\cdots 0|\,, \tag{27}$$

where $|0\cdots 0\rangle$ is on all $(n+1)$ qubits (latent register + ancilla).
The Grover operator is

$$Q = A\,S_0\,A^\dagger\,S_\chi. \tag{28}$$

In the two-dimensional subspace spanned by

$$|\Psi\rangle := A\,|0\rangle^{\otimes n}\,|0\rangle\,, \qquad |\Psi^\perp\rangle\,, \tag{29}$$

the operator Q acts as a rotation by angle 2θ. After m Grover iterations,

$$Q^m\,|\Psi\rangle = \sin\big((2m+1)\theta\big)\,|\Psi_1\rangle\,|1\rangle + \cos\big((2m+1)\theta\big)\,|\Psi_0\rangle\,|0\rangle. \tag{30}$$

Therefore,

$$\Pr(\text{ancilla} = 1 \mid m) = \sin^2\big((2m+1)\theta\big). \tag{31}$$

VIII. ESTIMATION METHODS

A. Quantum Amplitude Estimation (QAE)

The standard QAE uses phase estimation on Q to learn eigenphases $\pm 2\theta$ and hence $a = \sin^2 \theta$ with query complexity $O(\varepsilon^{-1})$ [1], [3].

B. Maximum Likelihood Estimation (MLE) from Grover Powers

To avoid deep phase estimation circuits, one may run experiments at several Grover powers $m \in \mathcal{M}$.
For each m, execute the circuit N_m times and observe k_m "successes" (ancilla measured as 1). Then

$$k_m \sim \text{Binomial}\big(N_m,\, p_m(\theta)\big), \quad p_m(\theta) = \sin^2\big((2m+1)\theta\big). \tag{32}$$

The likelihood is

$$\mathcal{L}(\theta) = \prod_{m \in \mathcal{M}} \binom{N_m}{k_m} \big(p_m(\theta)\big)^{k_m} \big(1 - p_m(\theta)\big)^{N_m - k_m}. \tag{33}$$

Equivalently, maximize the log-likelihood

$$\ell(\theta) = \sum_{m \in \mathcal{M}} \left[k_m \log p_m(\theta) + (N_m - k_m)\log(1 - p_m(\theta))\right], \tag{34}$$

over $\theta \in (0, \pi/2)$.

The MLE is

$$\hat{\theta} = \arg \max_{\theta \in (0,\pi/2)} \ell(\theta), \qquad \hat{a} = \sin^2(\hat{\theta}), \qquad \widehat{\mathbb{E}[L]} = L_{\max}\,\hat{a}. \tag{35}$$

IX. COMPLEXITY AND RESOURCES

Classical MC requires $O(\varepsilon^{-2})$ samples to achieve error ε. Quantum methods based on amplitude estimation achieve $O(\varepsilon^{-1})$ oracle queries [3], [4]. The required logical qubits are $n+1$ plus the work qubits for state loading and controlled rotations.

X. NUMERICAL RESULTS (SYNTHETIC STUDY)

We consider a synthetic latent-factor model

$$Z \sim \mathcal{N}(0, 1), \tag{36}$$

with probability of default

$$\text{PD}(Z) = \Phi(\alpha Z + \beta), \tag{37}$$

where Φ is the standard normal CDF, $\alpha = 0.8$, and $\beta = -0.3$. We set $\text{EAD} = \text{LGD} = 1$, so $L(Z) = \text{PD}(Z)$. The true expected loss is computed via high-precision numerical quadrature.

A. Error Scaling

Classical Monte Carlo estimation is performed using N independent samples. Quantum amplitude estimation is simulated by ideal Grover queries and exact amplitude extraction. of the sample size or oracle queries on a log–log scale.

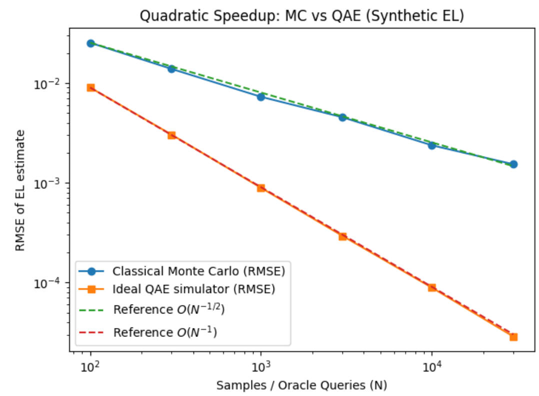

Fig. 1. Root-mean-square error (RMSE) of expected loss estimation versus number of samples (classical Monte Carlo) and oracle queries (quantum amplitude estimation) on a log–log scale. Classical Monte Carlo exhibits $O(N^{1/2})$convergence, while quantum amplitude estimation achieves $O(N^1)$scaling, demonstrating a quadratic speedup.

XI. CONCLUSION

We formulated expected loss estimation under a discretized latent-factor model as an amplitude estimation problem. By encoding the loss into ancilla amplitudes and applying Grover-style amplification, the EL reduces to estimating an amplitude a, enabling quadratic speedups over classical Monte Carlo using QAE or MLE-from-Grover-power measurements.

REFERENCES

[1] G. Brassard, P. Høyer, M. Mosca, and A. Tapp, "Quantum amplitude amplification and estimation," *Contemporary Mathematics*, vol. 305, pp. 53–74, 2002.
[2] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proc. 28th Annual ACM Symposium on Theory of Computing*, 1996.
[3] A. Montanaro, "Quantum speedup of Monte Carlo methods," *Proc. Royal Society A*, vol. 471, no. 2181, 2015.
[4] S. Woerner and D. J. Egger, "Quantum risk analysis," *npj Quantum Information*, vol. 5, no. 15, 2019.