Quantum finance and stochastic modelling

Encode payoffs and risk events as estimation tasks with explicit model and precision assumptions.

Undergraduate | 21 min lesson | Dr. Charith Jayasekara

By the end, you'll be able to

- Calculate a three-price expected payoff.
- Propagate flag accuracy into payoff units.
- Distinguish threshold probabilities, quantiles and tail means.

Separate the financial model from the estimator

A financial calculation may estimate an expected payoff, a loss probability or a tail-risk quantity under a chosen stochastic model. Quantum methods often target the numerical estimation step using amplitude estimation or related algorithms. They do not decide whether the market model is realistic, whether its parameters are stable or whether the resulting decision is economically appropriate.

For learning, use a small discrete model whose expectation can be calculated exactly. This makes the encoding and error propagation visible before discussing large stochastic simulations. A quantum circuit should reproduce the same mathematical quantity as the classical baseline, with the same distribution, payoff conventions and accuracy target.

Review probability amplitudes, sampling uncertainty and oracle costs. The central computational idea is to turn a bounded payoff or event into a measurable flag probability.

WORKED DIAGRAM · AA25

A bounded payoff becomes a probability



With unit discount, the expected payoff is 5. Dividing by 20 encodes a flag probability 1/4; probability error 0.005 contributes payoff error 0.1 before other errors.

State an expected-payoff setup

Specify the stochastic variables and their distribution, payoff function, discount factor if used, discretisation and target error. Assume a coherent implementation of distribution preparation and reversible payoff evaluation. The output is a scalar estimate with confidence, not a full description of every possible future scenario.

The model's data and parameters are inputs. A quantum routine does not automatically learn them from market observations. If calibration or scenario generation is expensive, include it in the workflow. State whether those costs are paid once and reused across many estimates or repeated for every requested calculation.

Work a three-price payoff example

Let a future price take values 80, 100 and 120 with probabilities one quarter, one half and one quarter. Consider a payoff equal to the positive part of price minus 100, and use a discount factor of one for this teaching example. The payoffs are zero, zero and 20, so the expected payoff is five.

Divide the payoff by 20 and encode it as a flag probability. Only the price-120 branch sets the flag with probability one, so the overall flag probability is one quarter. Estimating that probability and multiplying by 20 recovers five. A probability error of 0.005 corresponds to payoff error 0.1, before allowing for any model or arithmetic error.

The state preparation must use amplitudes whose squares are one quarter, one half and one quarter. Equal amplitudes over the three labels would estimate a different model. The payoff comparison and subtraction must also be reversible if used inside coherent amplitude estimation. A classical formula written in a notebook is not yet a complete oracle implementation.

The coherent flag targets the expected bounded payoff under the stated price model.

$$V(S)=\max(S-100,0),\quad \mathbb{E}[V]=\tfrac{1}{4}(0)+\tfrac{1}{2}(0)+\tfrac{1}{4}(20)=5,\qquad p_{\text{flag}}=\mathbb{E}[V/20]=\frac{1}{4},\qquad |\delta\mathbb{E}[V]|=20|\delta p|.$$

S takes values $80, 100, 120$ with the chapter's probabilities, and the teaching discount factor is one. Probability error 0.005 becomes payoff error 0.1; the statement concerns this toy model and does not validate any market assumptions.

Risk measures are different output tasks

A loss probability asks how often a threshold is exceeded. A quantile asks where the cumulative distribution reaches a chosen confidence level. A conditional tail expectation asks for an average over a defined tail, with care needed when probability mass lies exactly at the boundary. These quantities should not be merged under one generic "risk estimate."

For a simple loss distribution, suppose losses zero, ten and one hundred have probabilities 0.8, 0.15 and 0.05. The probability of a loss of at least one hundred is 0.05. Estimating it to additive error 0.01 is a much weaker relative guarantee than estimating it to error 0.001. Rare events require accuracy specifications matched to their scale.

Finding a quantile can use repeated threshold-probability estimates and a search over thresholds. Tail means add payoff-weighted estimation and normalisation. Count those repeated calls, and propagate their uncertainties. A single amplitude-estimation theorem does not mean every risk measure costs one identical circuit.

The algorithmic proposal

Woerner and Egger develop amplitude-estimation-based risk-analysis constructions and study trade-offs between circuit depth and convergence in small models. Their result motivates a coherent estimation approach under explicit assumptions; it is not evidence of a universal present-day advantage for production financial workloads. Read the quantum risk-analysis paper.

The potential query improvement comes from coherent access to the scenario generator and its inverse. Classical Monte Carlo uses many independent shallow samples. Quantum estimation may use fewer high-level queries but deeper reversible computations. Compare the cost of those queries rather than assuming they are equally cheap.

Follow an end-to-end workflow

Choose the economic quantity and stochastic model. Validate the classical implementation on small cases. Decide how to discretise or truncate variables and bound that error. Build distribution preparation and payoff arithmetic, then choose an amplitude-estimation variant and confidence target. Convert the output and error back into original units, and compare against a strong classical estimator at the same total accuracy.

Classical variance reduction, quasi-Monte Carlo, analytic approximations and model-specific simplifications can be powerful. A comparison against naive independent sampling may overstate the relevant application advantage. If the model has a closed-form expectation, a deep coherent estimator may be unnecessary even though it is mathematically valid.

Precision and hardware resources

Distribution preparation can dominate for complicated correlated scenarios. Payoff arithmetic can require many bits, temporary registers and uncomputation. Long amplification sequences multiply those costs and need sufficient logical reliability. Fault-tolerant gate counts, ancillas, runtime and repeated measurements should be estimated before drawing an engineering conclusion.

The total uncertainty also includes calibration and model risk. A tiny numerical confidence interval under a fixed model does not establish that future losses follow that model. Keep numerical estimation error separate from uncertainty about the assumptions. Faster calculation is useful only when the quantity and model are appropriate to the decision.

Practice with answers

Exercise 1. If the payoff example uses a discount factor of 0.95, what is the discounted expectation and how does probability error propagate?

<details>
<summary>Check the worked answer</summary>

Multiply the undiscounted expectation five by 0.95 to obtain 4.75. The probability-to-payoff scale becomes 19, so a probability error of 0.005 contributes at most 0.095 in discounted payoff units, excluding other errors.
</details>

Exercise 2. Why is estimating a five-percent tail probability to one-percent additive error potentially inadequate?

<details>
<summary>Check the worked answer</summary>

The additive error is one fifth of the probability itself, a 20-percent relative scale. A decision sensitive to small tail changes may need much tighter accuracy. State the required risk precision and confidence before choosing the estimator budget.
</details>

Check your understanding

PAUSE & THINK

What is the discounted expectation if the example's discount factor is 0.95?

- ☐ 4.75
- ☐ 5
- ☐ 19

Answer: 4.75. The undiscounted mean is 5, so multiplying by 0.95 gives 4.75. The model and estimator remain separate sources of uncertainty.

Sources

- Read the quantum risk-analysis paper