MDPI

Article

A Probabilistic Framework for Forecasting Cryptographic Security Under Quantum and Classical Threats

José R. Rosas-Bustos [1,2,3,*], Mark Pecen [3], Jesse Van Griensven Thé [1,2,3], Roydon Fraser [1,3], Nadeem Said [1], Sebastian Ratto Valderrama [5] and Andy Thanos [4]

[1] Department of MME, University of Waterloo, Waterloo, ON N2L 3G1, Canada
[2] LAKES Environmental Research Inc., Waterloo, ON N2L 3L3, Canada
[3] Applied Quantum Technologies (AQT) Initiative, Columbia, MD 21046, USA
[4] Cisco Systems, Inc., San Jose, CA 95134, USA
[5] Department of ECE, University of Waterloo, Waterloo, ON N2L 3G1, Canada
* Correspondence: jrosasbu@uwaterloo.ca

Abstract

This paper presents a probabilistic, multi-layered framework designed to forecast the longevity and security of cryptographic systems under the dual pressures of classical and quantum computational threats. The model integrates thermodynamic decay analogies, stochastic transitions via Hidden Markov Models, and an adapted financial option pricing method to quantify cryptographic degradation, strategic risk, and transition readiness. This model can guide standardization roadmaps, cipher retirement, or quantum-migration planning.

Keywords: cryptographic security; post-quantum cryptography; hidden Markov model; exponential decay; real options; binomial lattice; quantum threat modeling

Received:
Accepted:
Published:

Citation: Rosas-Bustos, J.R.; Pecen, M.; Van Griensven Thé, J.; Fraser, R.; Said, N.; Ratto Valderrama, S.; Thanos, A. Forecasting Cryptographic Security under Quantum and Classical Threats. *Symmetry* **2025**, *1*, 0. https://doi.org/

1. Introduction

The transition toward quantum-resilient cryptographic standards requires predictive tools that not only model technical degradation but also capture probabilistic events, such as the discovery of new quantum algorithms or hardware breakthroughs. We introduce a composite framework that unifies multiple analytical domains to forecast cryptographic security trajectories. Traditional assessments often treat cryptographic security as static or binary (e.g., "secure" vs. "broken"), lacking structures to capture evolving threats or decision-relevant metrics under uncertainty.

2. Model Components

2.1. Exponential Security Decay Model

We adopt a first-order exponential decay model, commonly used in thermodynamics, information theory, and reliability engineering, to represent the gradual erosion of cryptographic strength over time. The model assumes that the rate of decline in effective security $S(t)$ is proportional to its current value:

$$\frac{dS}{dt} = -k\,S(t), \tag{1}$$

where $k > 0$ is a decay constant representing the rate at which entropy or resistance to attack deteriorates under ambient conditions (e.g., hardware improvements, algorithmic refinements).

Solving the differential equation yields the general solution:

$$S(t) = S_0 e^{-kt}, \tag{2}$$

where S_0 is the initial security strength at $t = 0$. This form captures the asymptotic weakening of ciphers over time and can be empirically calibrated using historical cryptanalytic data.

2.2. Discrete-State Security Transition Modeling via Hidden Markov Processes

We represent the cryptographic system as a discrete-time, discrete-state stochastic process whose latent security condition evolves over time. Due to the partially observable nature of cryptographic robustness (e.g., undisclosed attacks or unknown algorithmic advancements), a Hidden Markov Model (HMM) is well-suited to characterize the transition dynamics between different security states.

Let the set of hidden states be defined as:

$$\mathcal{S} = \{S_1 : \text{Highly Secure},\ S_2 : \text{Moderately Secure},\ S_3 : \text{At Risk}\}.$$

The model assumes that at each time step t, the system is in some hidden state $s_t \in \mathcal{S}$, and transitions are governed by the Markov property:

$$P(s_{t+1} \mid s_t, s_{t-1}, \ldots, s_0) = P(s_{t+1} \mid s_t), \tag{3}$$

This yields a transition probability matrix T:

$$T = \begin{bmatrix} P(S_1 \to S_1) & P(S_1 \to S_2) & P(S_1 \to S_3) \\ P(S_2 \to S_1) & P(S_2 \to S_2) & P(S_2 \to S_3) \\ P(S_3 \to S_1) & P(S_3 \to S_2) & P(S_3 \to S_3) \end{bmatrix},$$

Each element $T_{ij} = P(s_{t+1} = S_j \mid s_t = S_i)$ reflects the likelihood of moving from one cryptographic state to another between discrete time steps. These probabilities can be empirically estimated or scenario-driven (e.g., conditioned on a timeline to quantum supremacy, adversarial investment, or regulatory lag).

The observable signal (e.g., key size recommendations, cryptanalytic reports, or algorithm deprecations) is emitted from the hidden state via an emission probability distribution.

This formulation allows us to model both gradual transitions and discontinuous shocks, such as a novel quantum algorithm discovery, or new kind of attack vectors, while maintaining a tractable probabilistic framework; see [1].

2.3. Real Options Approach to Cryptographic Security Valuation (after Pecen)

An option is a financial contract that gives its holder the right to buy or sell an underlying asset at a fixed price within a specified time period. A call option gives the holder the right to buy an underlying asset at a certain price, while a put option gives the holder the right to sell an underlying asset at a certain price. Options are considered a "wasting asset", as their value goes to zero upon expiry. The value of an option is based on the value of the underlying asset along with the exercise price of the option, market volatility and time to expiry. Inspired by Pecen's application of real options theory to intellectual property valuation (M. Pecen, personal communication, September 4, 2025), we model the residual value of a cryptographic system under uncertainty using a discrete-time binomial framework. This approach treats cryptographic security as a real asset whose

future utility can fluctuate due to advances in classical or quantum computational power, breakthroughs in cryptanalysis, or shifts in industry standards.

We begin with the classical Black-Scholes-style representation as a conceptual basis. While Black-Scholes provides theoretical intuition, its assumptions (e.g., continuous trading, log-normal returns, efficient and liquid public markets) do not hold for cryptographic systems. In addition, the Black-Scholes model tends to over-value the options associated with the underlying being far out-of-the-money. This works for financial trading because Black-Sholes is the de-facto standard, and because everyone uses the same model, everyone's valuation results are in line with one another. Hence, for the purpose of evaluating technologies as real-options, we adopt a discrete binomial model for practical evaluation.

$$V = S\,N(d_1) - Xe^{-rt}N(d_2), \tag{4}$$

where:

1. V is the security-adjusted present value of the cipher,
2. S is the current effective security (analogous to the underlying asset value),
3. X is the strike level or the minimum acceptable security threshold,
4. r is the discount or decay rate (e.g., technological obsolescence or adoption pressure),
5. $N(\cdot)$ is the cumulative standard normal distribution function,
6. t is the time horizon over which the cipher is expected to be in use.

However, since cryptographic systems are not traded in efficient markets and changes occur in discrete technological phases, we refine this approach using a **binomial tree model**. Let u and d be the up and down factors per period, and p the risk-neutral probability of an upward move. Over n periods, the security value evolves across a lattice, and the expected option-like value of the cipher is computed by backward induction.

Let:

1. $u = e^{\sigma\sqrt{\Delta t}}$ and $d = e^{-\sigma\sqrt{\Delta t}}$ be multiplicative security shift factors,
2. σ be the volatility of cryptographic risk (e.g., measured from historical compromise timelines),
3. $p = \frac{e^{r\Delta t} - d}{u - d}$ be the risk-neutral transition probability.

At each node, the expected value is:

$$V_{i,j} = e^{-r\Delta t}[p \cdot V_{i+1,j+1} + (1-p) \cdot V_{i+1,j}] \tag{5}$$

with terminal condition:

$$V_{n,j} = \max(S_{n,j} - X, 0) \tag{6}$$

This binomial valuation scheme allows for dynamic decision-making (e.g., retire, reinforce, or transition a cipher) based on real-time threat evolution and security valuation thresholds. Furthermore, the binomial tree approach can easily adapt to differing levels of volatility over time, and can be further extended to even a trinomial, or higher-level, model to evaluate multiple related assets.

3. Notation and Definitions

Table 1. Notation used in the modeling framework.

Symbol	Description
$S(t)$	Effective cryptographic security at time t.
S_0	Initial security value at $t = 0$.
k	Security decay constant.
s_t	Latent security state in HMM.
T_{ij}	Transition probability from s_i to s_j.
$V(t)$	Option-style value of a cipher at time t.
X	Minimum acceptable security threshold (strike).
r	Risk-free or decay-adjusted discount rate.
σ	Volatility of cryptographic risk.
u, d	Up and down factors in binomial model.
p	Risk-neutral probability of an up move.
n	Number of time steps in the binomial tree.

4. Model Integration Architecture

The proposed framework is composed of three interdependent sub-models, each capturing a distinct dimension of cryptographic lifecycle degradation. The integration of these models enables both continuous monitoring and forward-looking valuation under uncertainty. In practical use, they are executed as a pipeline: Stage I produces time-dependent security trajectories, Stage II consumes these trajectories to infer probabilistic risk states, and Stage III maps both quantities into migration-timing and valuation outputs.

4.1. Stage I: Continuous Security Degradation (Exponential Decay)

The security level $S(t)$ is computed as a time-dependent decay function:

$$S(t) = S_0 e^{-kt} \tag{7}$$

where k is a cipher-specific decay rate calibrated to historical cryptanalytic progress, Moore's law (or Neven's Law in quantum), and empirical measurements of implementation erosion (e.g., side-channel leakage growth).

In an operational setting, this stage is instantiated per cipher and per asset class. For each cipher–asset pair (c, a) in an organization's portfolio, the model produces a trajectory

$$S^{(c,a)}(t) = S_0^{(c,a)} e^{-k^{(c,a)} t}, \tag{8}$$

sampled at discrete time points $t = 0, \Delta t, 2\Delta t, \ldots, T$.

Output: The output of Stage I is therefore a matrix of trajectories $\{S^{(c,a)}(t)\}$ that summarise the expected erosion of security for each cipher and asset over time. This matrix is treated as an observable emission signal and passed as input to the HMM in Stage II.

4.2. Stage II: Probabilistic State Transition (Hidden Markov Model)

The time-evolving values $S^{(c,a)}(t)$ are used as observable inputs (emission signals) to infer the latent security state $s_t \in \{S_1, S_2, S_3\}$ using a Hidden Markov Model. For each cipher–asset pair, we construct an observation sequence

$$Y_t^{(c,a)} = S^{(c,a)}(t), \quad t = 0, \Delta t, 2\Delta t, \ldots, T, \tag{9}$$

which captures how its effective security evolves under the assumed classical and quantum threat models.

To connect continuous security levels to discrete risk regimes, we define a *risk-tolerance matrix* R whose entries $R_{a,s}$ encode, for each asset class a and latent state $s \in \{S_1, S_2, S_3\}$, the minimum acceptable security level (e.g., "bits of security"). Ranges of $S^{(c,a)}(t)$ are mapped to provisional labels such as *Highly Secure*, *Moderately Secure*, or *At Risk* using R, providing symbol sequences that initialise or constrain HMM training.

Standard HMM algorithms (e.g., Baum–Welch for parameter estimation and Viterbi decoding for state inference) are then applied to the sequences $\{Y_t^{(c,a)}\}$, yielding both the most likely state $s_t^{(c,a)}$ and posterior state probabilities

$$\pi_s^{(c,a)}(t) = P(s_t = s \mid Y_{0:t}^{(c,a)}) \tag{10}$$

for each cipher–asset pair and time. These probabilities capture, for example, when a classical algorithm is likely to move from *Moderately Secure* to *At Risk* under a given quantum-computing scenario.

Output: The numerical outputs of Stage II are trajectories of inferred states and state probabilities, together with an estimated transition matrix T, which together describe how quickly different ciphers are expected to enter unacceptable risk regimes. These quantities parameterise the valuation model in Stage III.

4.3. Stage III: Security Option Valuation (Binomial Model)

Both $S^{(c,a)}(t)$ and the current inferred state information from Stage II feed into a real-options framework to quantify the residual security value and migration incentives. For each cipher–asset pair, we construct a binomial lattice whose parameters are directly derived from earlier stages:

1. The latest continuous security level $S^{(c,a)}(t)$ represents the underlying value of cryptographic strength in the lattice.
2. The policy-defined minimum acceptable security threshold $X^{(a)}$ (the "strike") is taken from the risk-tolerance matrix R for asset class a.
3. The volatility $\sigma^{(c,a)}$ and effective discount rate $r^{(c,a)}$ are modulated by the HMM output, for example by increasing $\sigma^{(c,a)}$ as the probability $\pi_{S_3}^{(c,a)}(t)$ of being in the *At Risk* state grows, or by raising $r^{(c,a)}$ when adversarial activity is believed to be accelerating.

The binomial option valuation then proceeds using backward induction, yielding an option-style value $V^{(c,a)}(t)$ for maintaining or migrating each cipher–asset pair. Comparing $V^{(c,a)}(t)$ across time and across candidate post-quantum algorithms (e.g., replacing RSA-2048 with Kyber-768 or Dilithium-3) provides a quantitative basis for determining the time window during which proactive migration is economically and operationally justified.

Output: Stage III outputs a time-indexed option value and, by identifying when $V^{(c,a)}(t)$ falls below a policy-defined threshold, an implied "latest safe migration date" for each cipher–asset pair.

4.4. Functional Coupling Summary

Each stage, described above, both enriches the inference of the next and allows policy or organizational decisions to be made with increased granularity and confidence. The coupling can be summarised as follows:

- **Stage I (Security Degradation)** takes as input cipher parameters (e.g., key length, algorithm family, assumed hardware growth curves) and produces calibrated decay parameters (S_0, k) together with trajectories $S^{(c,a)}(t)$ for all cipher–asset pairs.
- **Stage II (State Inference)** ingests these trajectories as emission sequences, combines them with a risk-tolerance matrix and any external threat signals (e.g., new cryptana-

lytic results or standardisation announcements), and outputs the most probable latent security state and state-transition probabilities over time.

- **Stage III (Valuation and Migration Strategy)** uses both the continuous security levels and the discrete risk information to parameterise the binomial lattice—mapping $(S^{(c,a)}(t), \pi_s^{(c,a)}(t))$ into $(\sigma^{(c,a)}, r^{(c,a)}, X^{(a)})$—and returns option values and recommended migration windows.

This modular architecture allows each model to be improved or replaced independently based on future research or data availability, while preserving a clear data and decision flow across the three stages.

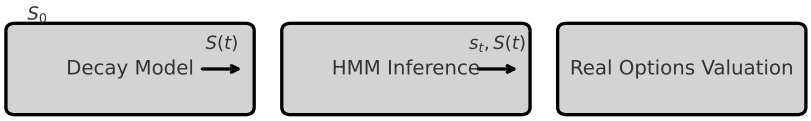

Figure 1. Integration architecture and information flow between the decay model, HMM, and real-options valuation stages.

System Flow:

$$(S_0, k) \xrightarrow{\text{Decay}} S^{(c,a)}(t) \xrightarrow{\text{HMM Emission}} \left(s_t^{(c,a)}, \pi_s^{(c,a)}(t)\right) \xrightarrow{\text{Parameterization}} V^{(c,a)}(t)$$

4.5. Practical Implementation Workflow

To move from conceptual modelling to an implementable tool, we outline a workflow that links the three stages of the framework to concrete data inputs and computational steps. The goal is that a standards body or system operator can start from an inventory of deployed ciphers and derive migration recommendations without modifying the underlying mathematics.

1. **Portfolio and parameter initialization (inputs to Stage I).** The practitioner first compiles a portfolio of cryptographic mechanisms $\mathcal{C} = \{c_1, \ldots, c_m\}$ used across asset classes (e.g., RSA-2048 for TLS key exchange, AES-256-GCM for bulk encryption, and post-quantum candidates such as Kyber-768 or Dilithium-3). For each cipher–asset pair (c, a), the initial security level $S_0^{(c,a)}$ is expressed in "bits of (classical and quantum) security," together with performance and deployment metadata (latency constraints, key-size limits, presence of hardware accelerators). Scenario-specific decay constants $k^{(c,a)}$ are then specified for classical-only, quantum-enabled, and aggressive-quantum threat models.
2. **Running Stage I and constructing observation sequences (output of Stage I → Stage II).** Stage I is evaluated over a discrete horizon to obtain trajectories $S^{(c,a)}(t) = S_0^{(c,a)} e^{-k^{(c,a)} t}$ for all (c, a). These trajectories are sampled at yearly or quarterly resolution to form observation sequences $Y_t^{(c,a)} = S^{(c,a)}(t)$. At this point, the artifact

produced by Stage I is a matrix of security levels over time whose rows are cipher–asset pairs and columns are time steps; this matrix becomes the emission data for the HMM.

3. **Risk-tolerance matrix and HMM calibration (Stage II implementation).** The organization specifies a risk-tolerance matrix R with rows indexed by asset class and columns by latent state $\{S_1, S_2, S_3\}$. Entry R_{a,S_2}, for example, encodes the minimum acceptable security level for asset class a to be considered *Moderately Secure*. Each element of $Y_t^{(c,a)}$ is mapped to a provisional state label using R, producing symbol sequences that are fed into a Hidden Markov Model. Using standard libraries, the practitioner estimates the transition matrix T and obtains posterior state probabilities $\pi_s^{(c,a)}(t)$, which summarise when each cipher–asset pair is likely to become *At Risk* under the specified quantum scenarios.
4. **Mapping to option parameters and valuation (Stage III implementation).** For each cipher–asset pair, Stage III takes as input the latest security level $S^{(c,a)}(t)$ from Stage I and the risk metrics from Stage II. The current underlying in the binomial lattice is set to $S = S^{(c,a)}(t)$; the strike $X^{(a)}$ is taken from the risk-tolerance matrix as the minimum acceptable security for asset class a; and the volatility $\sigma^{(c,a)}$ is calibrated from the variability of $\pi_{S_3}^{(c,a)}(t)$ over the planning horizon. Migration costs—including the expected performance overhead of PQC candidates (e.g., increased ciphertext size or handshake latency) and re-engineering effort—are folded into the effective discount rate $r^{(c,a)}$. Running the binomial option model then yields an option value $V^{(c,a)}(t)$ for migrating from a legacy cipher to a selected PQC algorithm.
5. **Decision outputs and iteration.** The final artifacts of the implementation are: (i) a ranked list of cipher–asset pairs by their option values $V^{(c,a)}(t)$, (ii) recommended migration windows and "latest safe migration dates" for each pair, and (iii) sensitivity analyses under different quantum-threat scenarios. In practice, the workflow is executed periodically as new cryptanalytic results, hardware benchmarks, or PQC performance measurements become available, updating $k^{(c,a)}$, R, T, and the resulting migration roadmap without changing the core structure of the framework.

5. Simulation Results

This section provides an initial simulation to illustrate the framework using synthetic parameters. The simulation mirrors the three-stage pipeline: we first generate a decay trajectory, then derive state transitions from that trajectory, and finally evaluate an option value conditioned on the inferred risk state.

5.1. Exponential Security Decay

Given initial conditions $S_0 = 100$ and $k = 0.02$, the security level over time follows:

$$S(t) = 100e^{-0.02t} \tag{11}$$

At $t = 40$, we compute:

$$S(40) \approx 44.93 \tag{12}$$

The decay curve is shown below:

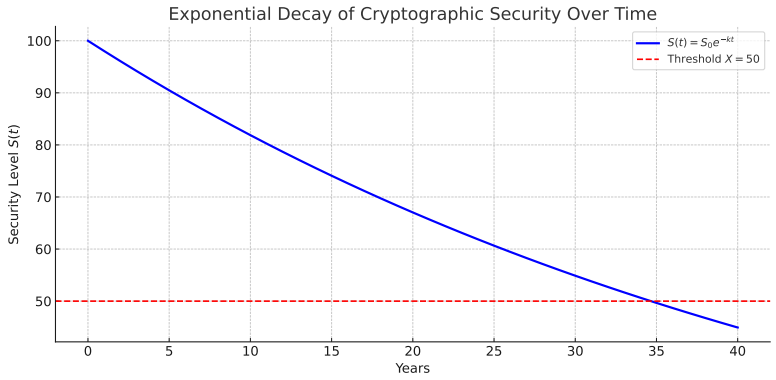

Figure 2. Exponential security decay for a synthetic cipher with $S_0 = 100$ and $k = 0.02$.

5.2. Hidden Markov Transitions

States are defined by thresholding $S(t)$:

1. $S(t) \geq 70$: *Highly Secure*
2. $50 \leq S(t) < 70$: *Moderately Secure*
3. $S(t) < 50$: *At Risk*

From these transitions, the estimated transition matrix T is:

$$T = \begin{bmatrix} 0.944 & 0.056 & 0.000 \\ 0.000 & 0.941 & 0.059 \\ 0.000 & 0.000 & 1.000 \end{bmatrix} \tag{13}$$

This matrix reflects the observed transitions during the decay horizon and serves as an example of how Stage I trajectories can be converted into Stage II transition probabilities.

5.3. Binomial Option Valuation

With parameters:

1. $r = 0.01$ (discount rate)
2. $\sigma = 0.15$ (volatility)
3. $X = 50$ (minimum security threshold)
4. $T = 10$ years
5. $n = 10$ (binomial steps)

we treat the current security level $S(0) = 100$ as the underlying asset and interpret crossing the threshold $X = 50$ as entering an unacceptable risk regime. Applying the binomial option valuation yields an option value of:

$$V(0) \approx 55.36 \tag{14}$$

indicating that, under this synthetic scenario, the cipher retains value above the operational threshold but with limited margin—consistent with a situation in which migration planning should begin.

5.4. Case Study: Migration from RSA-2048 to Kyber-768

To illustrate how the three stages of the framework can support post-quantum migration planning, we consider a simplified case study in which an Internet-facing service currently uses RSA-2048 for TLS key exchange and evaluates Kyber-768 as a post-quantum candidate. The asset under consideration is a latency-sensitive API, and we express security levels in bits of effective security against combined classical and quantum attackers.

5.4.1. Stage I: Security Decay under Quantum Arrival Scenarios

We set the initial security levels to $S_0^{\mathrm{RSA}} = 128$ bits and $S_0^{\mathrm{Kyber}} = 160$ bits. The policy threshold for acceptable security is fixed at $X = 64$ bits. For RSA-2048 we consider three quantum-arrival scenarios in which the algorithm first crosses the threshold X after $t_q \in \{10, 20, 30\}$ years, representing *early*, *baseline*, and *late* availability of large-scale Shor-capable quantum computers.

Using the exponential model of Section 3, the corresponding decay constants are obtained by solving

$$S_0^{\mathrm{RSA}} e^{-k_{\mathrm{RSA}}^{(q)} t_q} = X \;\Rightarrow\; k_{\mathrm{RSA}}^{(q)} = \frac{1}{t_q} \ln\left(\frac{S_0^{\mathrm{RSA}}}{X}\right), \tag{15}$$

which yields $k_{\mathrm{RSA}}^{(10)} \approx 0.0693$, $k_{\mathrm{RSA}}^{(20)} \approx 0.0347$, and $k_{\mathrm{RSA}}^{(30)} \approx 0.0231$. For Kyber-768 we assume a much slower decay, $k_{\mathrm{Kyber}} = 0.005$, capturing incremental algorithmic progress rather than a single catastrophic breakthrough.

For each cipher and scenario, the Stage I trajectories are

$$S^{\mathrm{RSA},(q)}(t) = S_0^{\mathrm{RSA}} e^{-k_{\mathrm{RSA}}^{(q)} t}, \qquad S^{\mathrm{Kyber}}(t) = S_0^{\mathrm{Kyber}} e^{-k_{\mathrm{Kyber}} t}, \tag{16}$$

evaluated on a 40-year horizon with yearly time steps. Figure 3 shows that, while RSA-2048 crosses the threshold X between 10 and 30 years depending on the scenario, Kyber-768 remains comfortably above X throughout the horizon.

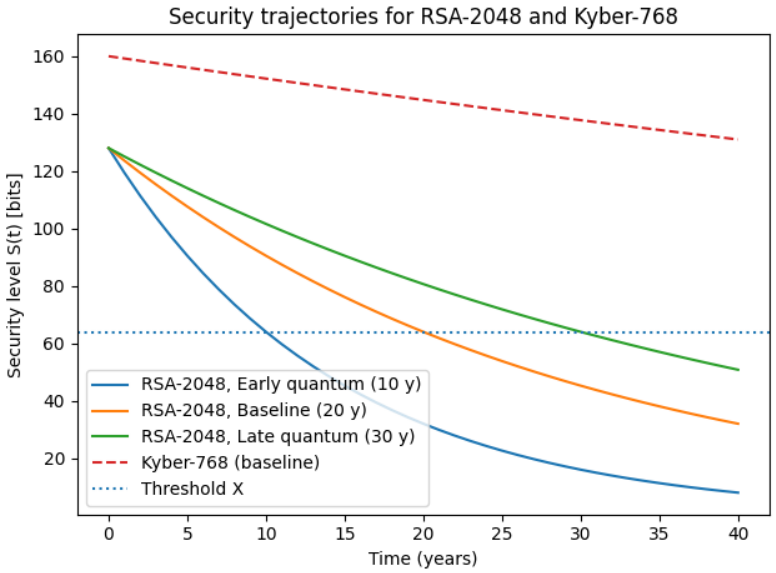

Figure 3. Security trajectories for RSA-2048 and Kyber-768 under early, baseline, and late quantum-arrival scenarios. The horizontal line at $X = 64$ bits denotes the minimum acceptable security level for the TLS service.

5.4.2. Stage II: State Transitions for the Baseline Scenario

For the baseline scenario ($t_q = 20$ years), we discretise $S^{\mathrm{RSA}}(t)$ into the three latent states of Section 3 using the risk-tolerance matrix R for an exposed online service:

$$\begin{aligned} S_1\ (\text{Highly Secure}) &: S(t) \geq 100, \\ S_2\ (\text{Moderately Secure}) &: 70 \leq S(t) < 100, \\ S_3\ (\text{At Risk}) &: S(t) < 70. \end{aligned}$$

Applying these thresholds to the yearly samples of $S^{\mathrm{RSA}}(t)$ produces a state sequence that starts in S_1, spends several years in S_2, and eventually enters S_3.

From this sequence we estimate the empirical transition matrix for RSA-2048 in the baseline scenario as

$$T^{\mathrm{RSA}}_{\mathrm{baseline}} = \begin{bmatrix} 0.875 & 0.125 & 0.000 \\ 0.000 & 0.900 & 0.100 \\ 0.000 & 0.000 & 1.000 \end{bmatrix}, \tag{17}$$

indicating that once the *At Risk* state is reached it is absorbing, and that the probability of remaining in *Highly Secure* or *Moderately Secure* declines over time. Under the same thresholds, the Kyber-768 trajectory remains in S_1 for the full horizon, yielding a degenerate transition matrix with $T^{\mathrm{Kyber}}_{11} = 1$.

5.4.3. Stage III: Option-Style Valuation of Migration Incentives

To capture the economic incentive to migrate, we apply the binomial option model of Section 3 separately to RSA-2048 and Kyber-768 in the baseline scenario. For RSA-2048 we set $S = S^{\mathrm{RSA}}_0 = 128$, strike $X = 64$, horizon $T = 10$ years, volatility $\sigma_{\mathrm{RSA}} = 0.15$, discount rate $r = 0.01$, and $n = 10$ time steps. For Kyber-768 we take $S = S^{\mathrm{Kyber}}_0 = 160$, $X = 80$, $\sigma_{\mathrm{Kyber}} = 0.12$, and the same T, r, and n.

Using backward induction on the corresponding binomial trees, we obtain option-style values

$$V^{\mathrm{RSA}}(0) \approx 70.86, \tag{18}$$

$$V^{\mathrm{Kyber}}(0) \approx 87.84. \tag{19}$$

The higher option value for Kyber-768 reflects both its larger initial security margin and its slower decay under the assumed quantum-threat model. Taken together with the transition matrix above, this case study shows how the framework can prioritise migration from RSA-2048 to Kyber-768: the Stage I trajectories identify when RSA-2048 will fall below policy thresholds under various quantum-arrival scenarios, Stage II quantifies the probability of entering the *At Risk* state, and Stage III converts these technical assessments into a comparative valuation that favours early adoption of the post-quantum cipher.

6. Limitations and Future Work

This model currently operates under several simplifying assumptions:

1. The decay rate k is constant over time, though it may vary in real-world conditions.
2. Hidden Markov transitions are based solely on $S(t)$ thresholds and do not yet incorporate noisy emissions or adversarial signals.
3. The option pricing model assumes a fixed time horizon and deterministic policy threshold X.

Future work will explore:

1. Calibration using real-world cryptanalytic event timelines (e.g., factoring breakthroughs, side-channel disclosures).
2. Monte Carlo simulations across diverse parameter settings and adversarial behaviours.
3. Integration with Standards post-quantum algorithm migration plans and cryptographic lifecycle guidelines.
4. To our knowledge, no existing framework integrates time-dependent security decay, probabilistic transitions, and financial valuation for cryptographic lifecycle forecast-

ing. This model bridges a key gap between theoretical cryptography and actionable cybersecurity risk management.

7. Related Work and Comparative Analysis

The need to anticipate cryptographic degradation has led to various modelling approaches across security, engineering, and economics. However, existing frameworks tend to focus on isolated aspects of risk and lack integration across domains.

NIST Lifecycle Models: NIST's cryptographic standards follow a phased lifecycle (approval, usage, deprecation), yet provide no predictive mechanism for degradation or transition planning.

Bayesian and Markov Models: Prior work has applied Bayesian inference and Markov chains to model intrusion detection and system states. These include probabilistic models of adversarial behaviour, but rarely incorporate continuous entropy decay or economic consequences.

Real Options in Technology and IP: Trigeorgis and Pecen have applied real options theory to R&D and intellectual property strategy. However, such models have not been adapted to forecast cryptographic resilience or transition timing.

Attack Trees and Risk Scoring: Schneier's attack trees and frameworks like FAIR model adversarial pathways or risk impact, but are not optimized for longitudinal entropy decay or valuation thresholds.

Adaptive Cryptography Models: Recent work has proposed adaptive cryptographic protocols that switch ciphers under attack. These models are reactive, lacking predictive degradation or financial valuation.

Table 2. Comparison of prior models and frameworks against the proposed integrated approach.

Model / Framework	Time	Stochastic	Valuation	Quantum	Simulatable	Lifecycle
NIST Lifecycle [4]	No	No	No	Partial	No	Yes
Attack Trees (Schneier) [5]	No	Yes (implicit)	No	No	No	Yes
FAIR Risk Model [6]	No	Yes	No	No	No	Yes
Bayesian/Markov Security Models [1]	Yes	Yes	No	No	Partial	No
Trigeorgis (Real Options) [2]	Yes	No	Yes	No	Yes	Yes
Pecen (IP Valuation)	Yes	No	Yes	No	Yes	Yes
This Work (Proposed)	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**

8. Conclusion

This framework, blending continuous degradation, probabilistic transitions, and financial risk models, offers a predictive approach for cryptographic lifecycle management. It is extensible, data-calibrated, and designed for the evolving landscape of classical and quantum threats.

Author Contributions: Conceptualization, J.R.R.B. and M.P.; methodology, J.R.R.B. and M.P.; investigation, J.R.R.B., M.P. and J.V.G.T.; validation, J.R.R.B., J.V.G.T. and R.A.F.; formal analysis, J.R.R.B.; writing - original draft, J.R.R.B.; writing - review and editing, J.R.R.B., M.P., J.V.G.T., R.A.F., N.S., S.R.V. and A.T.; supervision, R.A.F.

Funding: This research was funded by the Natural Sciences and Engineering Research Council of Canada (NSERC), Discovery Grants Program.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Data available from the corresponding author upon reasonable request.

Acknowledgments: Acknowledgements We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), RGPIN-2023-04513.
Cette recherche a été financée par le Conseil de recherches en sciences naturelles et en génie du Canada (CRSNG), RGPIN-2023-04513.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

HMM	Hidden Markov Model
CDF	Cumulative distribution function
PQC	Post-Quantum Cryptography

References

1. Rabiner, L.R. A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition. *Proc. IEEE* **1989**, *77*, 257–286.
2. Trigeorgis, L. *Real Options: Managerial Flexibility and Strategy in Resource Allocation*; MIT Press: Cambridge, MA, USA, 1996.
3. Shannon, C.E. A Mathematical Theory of Communication. *Bell Syst. Tech. J.* **1948**, *27*, 379–423.
4. National Institute of Standards and Technology (NIST). Post-Quantum Cryptography Standardization. Available online: https://csrc.nist.gov/projects/post-quantum-cryptography (accessed on 27 October 2025).
5. Schneier, B. Attack Trees. *Dr. Dobb's Journal* **1999**, December.
6. Jones, J. An Introduction to Factor Analysis of Information Risk (FAIR). Risk Management Insight LLC, 2005.
7. Kerschbaum, F.; Ochoa, M. Adaptive and Application-Aware Selection of Cryptographic Primitives. In *Proceedings of the 30th Annual Computer Security Applications Conference (ACSAC)*, Los Angeles, CA, USA, 7–11 December 2015.
8. Mun, J. *Real Options Analysis: Tools and Techniques for Valuing Strategic Investments and Decisions with Integrated Risk Management and Advanced Quantitative Decision Analytics*. Wiley Finance, Hoboken, NJ, USA, 2016.
9. Guthrie, G. *Real Options in Theory and Practice*. Oxford University Press, Oxford, UK, 2009.