

Factoring Hardness as a Bayesian Prior: A Decision-Theoretic Model of Post-Quantum Cryptography Migration Timing

Factoring Hardness as a Bayesian Prior: A Decision-Theoretic Model of Post-Quantum Cryptography Migration Timing

Abstract

The urgency of post-quantum cryptography (PQC) migration is usually justified by the anticipated capability of quantum computers to factor integers, yet that threat rests on an unproven complexity-theoretic conjecture: that integer factoring is not in BPP, the class of decision problems solvable in bounded-error probabilistic polynomial time. This paper treats that conjecture as an explicit Bayesian prior rather than a settled fact, and embeds it in a decision-theoretic model of migration timing for RSA-2048. We adopt a two-hypothesis mixture: under hypothesis H_{BPP} a cryptographically relevant breakthrough arrives with annual hazard $h_{\text{BPP}} = 0.10$; under H_{HARD} the hazard is $h_{\text{HARD}} = 0.01$; the prior weight is $\pi = \Pr(H_{\text{BPP}})$. With migration cost normalized to $C_m = 1$, breach loss L , and discount rate $r = 0.05$, we derive a closed-form loss threshold $L^*(\pi) = 1 + r/\bar{h}$ above which immediate migration is optimal, where $\bar{h}$ is the mixture hazard. At the baseline $\pi = 0.2$ the mixed hazard is $\bar{h} = 0.028$ and $L^* \approx 2.785714$: any asset whose compromise costs more than 2.79 migration-cost units justifies migrating immediately. Sensitivity analysis shows the immediate-migration region is robust: for $L = 10$ even $\pi = 0$ triggers immediate migration, while for $L = 2$ the prior must exceed $\pi^* \approx 0.444444$. Two alternative parameterizations — a discrete two-point scenario and a continuous two-hazard Poisson model with numerically solved optimal timing — are reported as labeled scenarios. The framework converts an unproven conjecture into an auditable sensitivity parameter.

1. Introduction

The deployment of post-quantum cryptography is one of the largest planned cryptographic transitions, and its urgency is almost always argued from the anticipated capability of quantum computers to factor integers in polynomial time via Shor's algorithm. But the security relevance of that argument depends on a logical structure that is rarely made explicit: quantum factoring threatens RSA-2048 only if classical polynomial-time factoring is impossible, i.e., only if $\text{FACTORING} \notin \text{BPP}$, where BPP is the class of decision problems solvable in bounded-error probabilistic polynomial time. If a classical algorithm ever factored RSA-2048 efficiently, quantum computers would be irrelevant to the RSA threat; if factoring is hard classically, the quantum threat is real but its arrival time is governed by engineering timelines. The conjecture $\text{FACTORING} \notin \text{BPP}$ is unproven. Migration policy therefore conditions on a conjecture, and the rational response to conditioning on an unproven proposition is to treat one's confidence in it as a prior and propagate it through the decision.

This paper does exactly that. We construct a minimal decision-theoretic model in which the probability mass assigned to a fast-breakthrough regime is a first-class parameter π , and we compute migration recommendations for RSA-2048 assets as a function of π , the breach loss L , the migration cost C_m , and the discount rate r . The model is deliberately simple — two hypotheses, constant hazards, exponential discounting — because its purpose is not forecasting but exposing the sensitivity structure: which conclusions about migration timing survive, and which collapse, when the epistemic status of the hardness conjecture is made explicit.

Our contributions are:

1. A closed-form Bayesian decision model for PQC migration under a two-hypothesis mixture over factoring hardness (Section 3).
2. Full arithmetic derivations of the mixed hazard, the expected discounted loss of remaining exposed, and the critical loss threshold $L^*(\pi)$ (Section 4).
3. A sensitivity analysis over the prior π showing that for high-value assets the "migrate now" recommendation is robust even at $\pi = 0$, while for low-value assets the decision flips at a computable prior threshold π^* (Sections 4 and 5).
4. Two labeled alternative scenarios, reconciled from independent draft parameterizations: a discrete two-point cost scenario and a continuous two-hazard Poisson model in which the optimal migration time T^* solves a first-order condition numerically (Sections 3.4 and 4.4).

5. A discussion of how much of the quantum-advantage narrative for cryptography depends on the single unproven premise $\text{FACTORING} \notin \text{BPP}$, and of a proposed stress-test program classifying known exponential speedups by their dependence on abelian hidden-subgroup structure (Section 6).

Throughout, every number is either an input declared in Section 3 or computed with shown arithmetic in Section 4. No empirical data on expert surveys or engineering timelines are invented; where such data would refine the inputs, we say so explicitly.

2. Background and Related Work

We organize the related work into three threads: the PQC migration problem itself, Bayesian methods for priors and expert judgement, and analogues of migration, hardness, and displacement modeling in adjacent fields. Every statement below is restricted to what the supplied bibliography entry supports; where an entry's summary is thin, we say so.

PQC migration management. The most directly relevant work is the PMMP framework of [2], which addresses the problem that organizations must plan migration to quantum-resilient cryptographic measures (PQC), notes that this is difficult, and states that — to the authors' knowledge — no generalized approach existed that manages such a complex migration while explicitly integrating into organizations' steering mechanisms and control systems; PMMP presents a risk-based process for this purpose. Our model is complementary to [2] in a specific sense: PMMP supplies the organizational process through which a migration decision is executed, while we supply the decision-theoretic core that determines *when* the migration should occur, conditional on an explicit prior over the hardness conjecture. A risk-based process needs a quantified risk; Sections 3–5 provide one candidate quantification whose inputs are auditable.

Hardness under restrictions. Work on computational hardness under severe restrictions offers a cautionary analogy. The Tetris analysis of [1] proves that Tetris clearing and survival remain NP-hard even when the piece set is restricted to a single tetromino type, for every tetromino type except the O piece. The structural lesson we import is twofold: first, hardness claims can be established rigorously for restricted problem variants, so "hardness" is a property that can be pinned down rather than asserted; second, the result of [1] is a *worst-case* NP-hardness statement about a puzzle, whereas cryptographic security needs *average-case* hardness against natural distributions — a gap that is exactly why the RSA-2048 threat ultimately rests on the unproven average-case conjecture $\text{FACTORING} \notin \text{BPP}$ rather than on any worst-case theorem.

Bayesian priors, mixtures, and expert judgement. Our two-hypothesis prior over factoring hardness is a spike-and-slab-style mixture: a point mass on "breakthrough regime" and a diffuse mass on "hard regime." The Bayesian group lasso work of [5] revisits the Bayesian group lasso using spike-and-slab priors for group variable selection, demonstrates the connection of the model with penalized regression, and shows the posterior median estimator has the oracle property under orthogonal designs. We borrow the structural idea — a two-component prior in which one component is a point mass — and apply it to a hypothesis rather than a coefficient. On the question of where the prior π should come from, the Posterior Belief Assessment framework of [7] addresses precisely the situation in which a Bayesian analysis is too complex to assert a precise correspondence between expert probabilistic judgements and the chosen prior and likelihood, and proposes a finite procedure for attributing meaning to the results of such an analysis; this is the natural methodological home for eliciting π from complexity-theory experts, though we do not perform such an elicitation here. On the choice of prior form itself, the objective Bayesian analysis of bilateral data in [8] derives Jeffreys' prior and Bernardo's reference prior for the risk difference, the risk ratio, and the odds ratio under Dallal's model and the saturated model; the analogy is that when little is known, an objective or reference prior over the hardness hypothesis is a defensible default, and our sensitivity analysis over $\pi \in [0, 1]$ is the practical substitute for committing to one. Finally, [4] develops Bayesian inference for design-based sampling with inverse-probability weights, using a hierarchical approach that models the distribution of the weights of nonsampled units and includes them as predictors in a nonparametric model. This is relevant to a future refinement of our model in which expert survey opinions are reweighted for selection effects before being pooled into π ; the supplied summary of [4] supports the general point that Bayesian inference under informative weighting requires explicit modeling of the weights, and gives no further detail we rely on.

Migration and displacement analogues. The Migration Networks chapter of [3] reviews the study of macroscale migration patterns as networks whose nodes represent places (countries, cities, rural areas) and whose edges encode migration ties, surveying network-analysis methods applied to such networks. We borrow the framing, not the content: cryptographic migration can be viewed as a flow problem over an asset graph, in which "migration pressure" propagates from broken primitives through dependent protocols and systems. The entry supports only this general description of network-based migration analysis, and we

make no claim about its specific findings. The co-design study of [6] addresses performance analysis of migrating tasks from software to hardware in multiprocessor systems-on-chip, motivated by the complexity of multimedia applications and deadline pressure on designers; the supplied summary truncates before giving quantitative results we could use. Its relevance is the shared structural problem — deciding *when* to move a function between implementation substrates under cost and deadline pressure — which is isomorphic to our migration-timing problem with the roles of "hardware" and "PQC" interchanged.

Default parameters made explicit. Two entries supply methodological siblings for our central move of turning an inherited default into a decision variable. The radix study of [9] examines hierarchical (ultrametric) representations organizing N items in a q -ary tree of depth $n = \lceil \ln N / \ln q \rceil$, in which the branching factor q determines how storage, interconnection, and traversal resources scale, and notes that prior modeling treated the radix as a fixed inherited structural parameter. Prior PQC planning similarly treats "the quantum threat is real" as a fixed inherited parameter; our contribution is to make that parameter a choice variable of the model, exactly as [9] does for q . The non-Archimedean geometry assessment of [10] observes that the real numbers are the default continuum for physical modeling — quantum states live in Hilbert space over $\mathbb{C}$, neural networks optimize over $\mathbb{R}^n$, codes are designed with Euclidean distance — while Ostrowski's theorem demonstrates that alternative completions exist; it is a structured assessment of a thesis whose adoption hinges on a foundational comparison. Both papers ask when a community should abandon a default foundation (the reals; classical hardness) given that the alternative is mathematically coherent but practically unproven. The negotiation study of [11] examines an agentic buyer with a hard fulfillment deadline who can fork a procurement negotiation into parallel seller-facing threads, where each thread consumes resources and each simultaneous acceptance creates cancellation and commitment liability; the structural analogy is to migration scheduling under a hard deadline: early migration consumes resources now, deferred migration creates liability exposure later, and the planner must jointly choose timing and intensity of commitment. Finally, [12] belongs to the popular-facing genre in which quantum-urgency narratives circulate; its supplied summary gives no substantive text beyond its title, so we draw no technical claim from it.

Together these twelve works supply the process layer [2], the epistemology of hardness claims [1], the prior-elicitation machinery [4], [5], [7], [8], and a family of structural analogues [3], [6], [9], [10], [11], [12] for the model we now build.

3. Methods

3.1 Model primitives (main model)

Time is measured in years. The organization holds RSA-2048 assets whose compromise would cause loss L , expressed in units of the migration cost C_m (so $C_m = 1$ by normalization). The discount rate is r per year. The hardness conjecture is represented by a two-hypothesis mixture:

- H_{BPP} : the world is in a regime where cryptographically relevant factoring breakthroughs arrive with annual hazard h_{BPP} ;
- H_{HARD} : breakthroughs arrive with annual hazard $h_{\text{HARD}} \ll h_{\text{BPP}}$;
- $\pi = \Pr(H_{\text{BPP}})$ is the Bayesian prior weight on the fast-breakthrough regime.

The mixture (predictive) hazard is

$$\bar{h}(\pi) = \pi h_{\text{BPP}} + (1 - \pi) h_{\text{HARD}}.$$

A breakthrough is modeled as a memoryless Poisson event: conditional on the regime, the waiting time to breakthrough is exponential with the regime's hazard. The prior π is the Bayesian object of interest: it is not a physical constant but a posterior summary of the community's belief that the fast-breakthrough regime obtains. We treat π as an input with a full sensitivity range $[0, 1]$ rather than a point estimate, and we recommend eliciting it via the finite procedure of Posterior Belief Assessment described in [7], which is designed for exactly the case where the correspondence between expert judgements and prior forms cannot be asserted precisely. Where expert opinion is sparse, an objective-prior construction in the spirit of Jeffreys or reference priors [8] is an alternative we leave open.

3.2 Expected discounted loss of remaining exposed

If the organization never migrates proactively and a breakthrough arrives at time t , it suffers the loss L at t , discounted to present value by e^{-rt} . Averaging over the exponential waiting time under the mixture hazard $\bar{h}$:

$$E_{\text{exposed}} = \int_0^\infty L \bar{h} e^{-\bar{h}t} e^{-rt} dt = L \bar{h} \int_0^\infty e^{-(\bar{h}+r)t} dt = \frac{L \bar{h}}{\bar{h} + r}.$$

Immediate migration costs $C_m = 1$ at $t = 0$ and eliminates the exposure. Immediate migration is optimal exactly when

$$1 < \frac{L \bar{h}}{\bar{h} + r} \iff L > L^*(\pi) = \frac{\bar{h} + r}{\bar{h}} = 1 + \frac{r}{\bar{h}}.$$

The threshold $L^*(\pi)$ is the paper's central closed-form result: it is decreasing in π , and as $\pi \rightarrow 1$ it approaches $1 + r/h_{\text{BPP}}$, while as $\pi \rightarrow 0$ it approaches $1 + r/h_{\text{HARD}}$.

3.3 Prior threshold for a given asset value

Inverting the threshold for a given L : immediate migration is optimal iff $\bar{h}(\pi) > r/(L - 1)$ (for $L > 1$; for $L \leq 1$ migration is never immediately worthwhile under this criterion). Since $\bar{h}(\pi) = h_{\text{HARD}} + \pi(h_{\text{BPP}} - h_{\text{HARD}})$, the critical prior is

$$\pi^*(L) = \frac{\frac{r}{L-1} - h_{\text{HARD}}}{h_{\text{BPP}} - h_{\text{HARD}}}.$$

3.4 Alternative scenario models (reconciled from independent parameterizations)

Two further parameterizations are reported as labeled scenarios in Section 4.4, with their own declared inputs and shown arithmetic.

Scenario D (discrete two-point). A single horizon T , a prior p of a breakthrough within the horizon, migration cost C_m paid at $t = 0$ if migrating immediately, breach cost C_b discounted by $d = e^{-rT}$ if the organization waits. The expected cost of waiting until the horizon is $E_{\text{wait}} = p C_b d + (1 - p) C_m$, and the break-even prior under the simplified comparison (migration forced at T) is $p^* = C_m/(C_b d)$.

Scenario P (continuous two-hazard Poisson). A classical breakthrough arrives with hazard λ and a quantum capability with independent hazard μ ; combined hazard $\Lambda = \lambda + \mu$. Migration at time T costs C_m (discounted); if a breakthrough of either type occurs at $t < T$, loss accrues at rate L per year until T . The expected discounted cost is

$$E[C](T) = C_m e^{-rT} + L \Lambda \int_0^T e^{-\Lambda t} (T - t) e^{-rt} dt,$$

and with $a = \Lambda + r$ the first-order condition for the optimal migration time T^* is

$$r C_m e^{-rT^*} = \frac{L \Lambda}{a} (1 - e^{-aT^*}),$$

solved numerically by bisection in Section 4.4. Scenario P makes the counterfactual structure explicit: the quantum-only literature implicitly sets $\lambda = 0$, so all urgency flows from hardware; our model adds the classical-hazard term that the purely quantum framing omits.

3.5 The reduction-of-urgency argument

The entire quantum-advantage case for RSA-2048 urgency reduces to the single unproven premise $\text{FACTORING} \notin \text{BPP}$: if the premise fails, quantum hardware is irrelevant to RSA. A companion structural question — what fraction of known exponential speedups depend on abelian hidden-subgroup structure, of which factoring is the flagship instance — is a classification problem we flag as a labeled stress-test hypothesis in Section 6, not as a computed result.

4. Analysis

Every input number below is stated with its source: all monetary and rate parameters are *modeling assumptions of this paper* (declared here, not measured), and all outputs are computed from them by the arithmetic shown.

4.1 Main model: mixture hazard and loss threshold

Inputs (assumptions of this paper). $h_{\text{BPP}} = 0.10$ per year; $h_{\text{HARD}} = 0.01$ per year; $\pi = 0.2$; $r = 0.05$ per year; $C_m = 1$ (normalization).

Step 1: mixture hazard.

$$\bar{h} = \pi h_{\text{BPP}} + (1 - \pi) h_{\text{HARD}} = 0.2 \times 0.10 + 0.8 \times 0.01 = 0.020 + 0.008 = 0.028.$$

Step 2: expected discounted loss of remaining exposed, for $L = 10$.

$$E_{\text{exposed}} = \frac{L \bar{h}}{\bar{h} + r} = \frac{10 \times 0.028}{0.028 + 0.05} = \frac{0.28}{0.078} \approx 3.589744.$$

Since $E_{\text{exposed}} \approx 3.59 > C_m = 1$, immediate migration is optimal at the baseline for a ten-unit asset.

Step 3: loss threshold.

$$L^*(\pi) = 1 + \frac{r}{\bar{h}} = 1 + \frac{0.05}{0.028} = 1 + 1.785714 \dots \approx 2.785714.$$

(Here $0.05/0.028 = 50/28 = 25/14 = 1.785714 \dots$) Any asset with $L > 2.785714$ migration-cost units justifies immediate migration at $\pi = 0.2$.

Step 4: sensitivity in π .

- $\pi = 0$: $\bar{h} = 0.01$, so $L^* = 1 + 0.05/0.01 = 1 + 5 = 6$. An asset with $L = 10 > 6$ still triggers immediate migration even if the fast-breakthrough regime has zero prior weight — the recommendation is robust to the conjecture entirely failing to be doubted.
- $\pi = 1$: $\bar{h} = 0.10$, so $L^* = 1 + 0.05/0.10 = 1.5$.

Step 5: critical prior for a low-value asset, $L = 2$. Required mixture hazard:

$$\bar{h}^{\text{req}} = \frac{r}{L - 1} = \frac{0.05}{2 - 1} = 0.05.$$

Solve $0.01 + 0.09 \pi = 0.05$:

$$\pi^* = \frac{0.05 - 0.01}{0.09} = \frac{0.04}{0.09} = 0.444444 \dots$$

For a two-unit asset, the fast-breakthrough regime must carry more than a 44.44% prior before immediate migration beats waiting.

4.2 Scenario D (discrete two-point): full arithmetic

Inputs (assumptions of this paper, in dollars). $p = 0.05$ (prior probability of a breakthrough within the horizon); $C_m = \text{\$}10,000,000$; $C_b = \text{\$}200,000,000$; $r = 0.05$ per year (continuous); $T = 10$ years.

Step 1: discount factor. $d = e^{-rT} = e^{-0.05 \times 10} = e^{-0.5} \approx 0.606531$.

Step 2: expected discounted breach cost when waiting.

$$p C_b d = 0.05 \times 200,000,000 \times 0.606531 = 10,000,000 \times 0.606531 \approx \text{\$}6,065,310.$$

Step 3: expected total cost of waiting.

$$E_{\text{wait}} = p C_b d + (1 - p) C_m = 6,065,310 + 0.95 \times 10,000,000 = 6,065,310 + 9,500,000 = \text{\$}15,565,310.$$

Since $C_m = \text{\$}10,000,000 < E_{\text{wait}} \approx \text{\$}15.57$ million, immediate migration is optimal under Scenario D's baseline.

Step 4: break-even prior (simplified comparison with migration forced at T).

$\$C_b, d = 200, 000 \times 0.606531 \approx \$121, 306, 200, \text{quad p}$

References

[1] Tetris is Hard with Just One Piece Type. arXiv:2603.09958v1. <https://arxiv.org/abs/2603.09958v1> [2] PMMP -- PQC Migration Management Process. arXiv:2301.04491v4. <https://arxiv.org/abs/2301.04491v4> [3] Migration Networks: Applications of Network Analysis to Macroscale Migration Patterns. arXiv:2002.10992v2. <https://arxiv.org/abs/2002.10992v2> [4] Bayesian Nonparametric Weighted Sampling Inference. arXiv:1309.1799v4. <https://arxiv.org/abs/1309.1799v4> [5] Bayesian Variable Selection and Estimation for Group Lasso. arXiv:1512.01013v1. <https://arxiv.org/abs/1512.01013v1> [6] Performance Analysis of Software to Hardware Task Migration in Codesign. arXiv:1002.1154v1. <https://arxiv.org/abs/1002.1154v1> [7] Posterior Belief Assessment: Extracting Meaningful Subjective Judgements from Bayesian Analyses with Complex Statistical Models. arXiv:1512.00969v1. <https://arxiv.org/abs/1512.00969v1> [8] Objective Bayesian Inference for Bilateral Data. arXiv:1504.06523v1. <https://arxiv.org/abs/1504.06523v1> [9] QNFO: Optimal Radix q under Resource Constraints: A Reconciled Derivation Framework for Ultrametric Hierarchies [10] QNFO: When Will Non-Archimedean Geometry Displace the Real Numbers? A Structured Assessment of the Adelic Substrate Thesis [11] QNFO: Parallelism or Concession? A Reconciled Analytical Model of Concurrency-Aware Procurement Negotiation for Agentic Commerce [12] QNFO: The Revolutionary Beginner's Guide to Quantum Computing: Why We Don't Have Quantum Computers Yet " and What the Geometric Alternative Offers