# The True Cost of Factoring: Magic and Number-Theoretic Complexity in Shor's Algorithm

**Matteo Seclì**[1,2]
Alessio Paviglianiti[1,2], Emanuele Tirrito[1,2], and Vincenzo Savona[1,2]

[1]*Institute of Physics, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland*
[2]*Center for Quantum Science and Engineering, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland*

matteo.secli@epfl.ch

Shor's factoring algorithm [1] is typically benchmarked by asymptotic gate counts and register sizes, yet these metrics do not directly capture which intrinsically quantum resources, and in what amount, must be created and maintained for the computation to succeed.

Here we formulate a resource-theoretic characterization of period finding, using non-stabilizerness ("magic") as the central currency [2–4]. Magic quantifies how much of a quantum state lies outside the efficiently simulable stabilizer/Clifford regime. It is therefore a natural choice for such a characterization, because stabilizer (Clifford) dynamics admit efficient classical simulation, whereas non-stabilizer resources are required to go beyond that regime and enable universal fault-tolerant quantum computation [2,3].

Within this framework we track how non-stabilizerness is generated and redistributed across the execution of Shor's routine, and we mathematically relate this behaviour to the instance structure of the modular arithmetic problem being solved. When factoring an integer $N$, in fact, Shor's algorithm [1] randomly selects an integer $a$ with $\gcd(a, N) = 1$ and reduces factorization to estimating the period $r$ of the function $f(x) = a^x \bmod N$. The number-theoretic complexity of factoring $N$ is then strongly influenced by the size of $r$, and by the rarity of coprimes $a$ with small periods. Here we show that this difficulty is reflected in the magic budget required by the corresponding quantum states.

Furthermore, we connect QFT precision, which is the well-known practical lever for maintaining Shor's success probability [1], to resource-theoretic constraints: reduced-precision QFTs restrict the accessible non-stabilizer structure of the evolving state, which in turn limits the reliability of period inference.

Our results provide a concise operational link between what must be learned about a specific factoring instance and what must be spent in non-stabilizer resources to learn it, and are therefore directly relevant to fault-tolerant realizations of Shor's algorithm. In leading fault-tolerant schemes, the dominant overhead is typically associated with supplying non-Clifford operations via resource states (magic states), making the magic budget a practical proxy for cost. By expressing instance- and precision-dependence at the level of non-stabilizerness, our approach complements standard circuit-cost analyses [1–4] with a resource metric that is naturally aligned with the real bottlenecks of fault-tolerant quantum computation.

## References

[1] P. W. Shor, SIAM Journal on Scientific and Statistical Computing, 26 (1997) 1484
[2] V. Veitch, et al., New Journal of Physics, 16 (2014) 013009
[3] M. Howard and E. Campbell, Physical Review Letters, 118 (2017) 090501
[4] L. Leone et al., Physical Review Letters, 128 (2022) 050402