

Fundamental accuracy-resolution trade-off for timekeeping devices

Florian Meier,^{1,2,*} Emanuel Schwarzhans,^{1,†} Paul Erker,^{1,3,‡} and Marcus Huber^{1,3,§}

¹*Vienna Center for Quantum Science and Technology,*

Atominstytut, Technische Universität Wien, 1020 Vienna, Austria

²*Institute for Theoretical Physics, ETH Zurich, 8093 Zürich, Switzerland*

³*Institute for Quantum Optics and Quantum Information,
Austrian Academy of Sciences, 1090 Vienna, Austria*

(Dated: November 30, 2023)

From a thermodynamic point of view, all clocks are driven by irreversible processes. Additionally, one can use oscillatory systems to temporally modulate the thermodynamic flux towards equilibrium. Focusing on the most elementary thermalization events, this modulation can be thought of as a temporal probability concentration for these events. There are two fundamental factors limiting the performance of clocks: On the one level, the inevitable drifts of the oscillatory system, which are addressed by finding stable atomic or nuclear transitions that lead to astounding precision of today's clocks. On the other level, there is the intrinsically stochastic nature of the irreversible events upon which the clock's operation is based. This becomes relevant when seeking to maximize a clock's resolution at high accuracy, which is ultimately limited by the number of such stochastic events per reference time unit. We address this essential trade-off between clock accuracy and resolution, proving a universal bound for all clocks whose elementary thermalization events are memoryless.

Clocks dominate our daily lives unlike any other technology – from ordinary things like catching the train in the morning to locating ourselves using the GPS – they are involved everywhere. Physics has provided the theoretical and experimental foundation to develop accurate and stable clocks which has culminated in the 50s with the invention of atomic clocks [1]. Year by year, these clocks have been improving their accuracy, with a state-of-the-art optical clock accumulating less than a hundred milliseconds of error over the lifespan of the sun [2–4]. These advancements beg for the question whether there are any physical principles constraining a clock's performance. Driven by this open problem, quantum clocks are emerging into an independent field of research, unifying approaches from quantum information theory [5–8] and quantum thermodynamics [9–13]. Recent experiments showcase today's technology is capable of exploring these ultimate limits of timekeeping [14–16].

This letter explores fundamental limitations of clocks coming from thermodynamics, because clocks, like all other physical systems, are subject to thermodynamical laws [11]. Even worse, they very much rely on the increase of entropy originating in the second law of thermodynamics. This implies that (1) clocks are witnesses of the macroscopic breaking of time-reversal symmetry because they tick forwards in time. (2) Thus, they must be driven by irreversible processes that drain out-of-equilibrium resources to output temporal information. These processes, though, are thermodynamic, therefore inherently stochastic and never perfectly predictable. (3) We conclude, even an idealized clock could never be perfect simply due to the fact that the clock itself is fundamentally driven by stochastic processes.

We look at clocks that use this flow by counting elementary thermalization events to define *ticks*. Instructive examples of such a process could range from grains of sand that pass through an hourglass to the photons that

are reflected from a pendulum to ascertain its position. Even modern atomic clocks ultimately rely on a macroscopic number of photons to read out the laser frequency (which is stabilized by feedback from atomic transitions). In all these cases, the underlying thermalization processes are stochastic and have some intrinsic rate Γ . The time passing between two events, relative to assumed smooth parameter time, is probabilistic and described by a probability density function. We refer to its average as μ and its uncertainty as σ . We define the *resolution* ν of the clock to be the inverse average time between two ticks $\nu = 1/\mu$, and the *accuracy* N as the average number of times the clock ticks until it is off by one tick. For a sequence of independent and identically distributed ticks, the accuracy equals the signal-to-noise ratio $N = \mu^2/\sigma^2$ [9]. A clock usually modulates the temporal distribution of when the thermalization events occur as to improve its accuracy. For most practical clocks, the stability of this temporal modulation (for instance the laser frequency in atomic clocks) is the limiting factor to the accuracy. We show that even if this temporal modulation is perfectly stable, the stochastic process underlying the tick generation limits the accuracy at given resolution. In other words, the need to irreversibly generate a signal bounds the clock's performance even if it uses an eternally stable frequency reference.

Theorem (Accuracy-resolution trade-off). *Clocks with elementary ticking events generated by a memoryless stochastic process at rate Γ obey the trade-off relation*

$$N \leq \frac{\Gamma^2}{\nu^2}. \quad (1)$$

This trade-off complements other established results in the field considering restrictions imposed on the clock through entropy production [9, 17] or Hilbertspace dimension [18, 19].

Temporal probability concentration. The most primitive clock consists of two out-of-equilibrium thermal reservoirs in contact with each other. Counting the individual stochastic thermalization events as ticks can serve as a way to measure time, we call this the *thermal reference clock*. Sufficiently large reservoirs are memoryless, therefore, such stochastic jumps at equal rate are exponentially distributed [20–22]. The coupling of the two baths defines a characteristic thermalization rate Γ . For an exponential distribution the standard deviation equals the average, which leads to $\mu = \sigma = \Gamma^{-1}$. Consequently, such a clock has unit accuracy $N = 1$ and only by averaging over many of these ticks are we able to achieve higher accuracy, but at the expense of resolution. Averaging over M independent and identically distributed (i.i.d.) such events increases the variance and mean of the tick time M -fold, leading to a resolution $\nu \propto 1/M$ and accuracy $N \propto M$. This gives an inverse proportional accuracy-resolution scaling

$$N = \frac{\Gamma}{\nu}, \quad (2)$$

quadratically smaller than the upper bound in eq. (1). Upon closer inspection, we find a thermodynamic cost associated to this increase in accuracy: instead of having a single irreversible event producing a tick, now M irreversible events are required.

A natural question to ask is whether it is possible to increase the accuracy beyond what is achievable through averaging in eq. (2), while still using the same underlying stochastic tick generating process. The answer lies in the observation that all ticking clocks known to us use a combination of two processes to tell time:

- (a) irreversible processes that generate ticks, and
- (b) a filter process, *temporal probability concentration*, which modulates the probability of the irreversible ticking events to occur.

By means of temporal probability concentration (TPC), a (sensible) clock centers the probability distribution of the stochastic events with a periodic process such that the ticks occur closely around well-defined instants in time. Exponential decay of an unstable two-level system is an example for a clock without TPC where the decay defines the tick. The tick probability density conditioned on the tick not yet having occurred equals Γ , the decay constant which is time-independent. Such a clock has accuracy $N = 1$. Clockworks modulate this probability by, for example, driving the two-level system from the ground state into the excited state. This gives an effective time-dependency of the ticking probability as illustrated in Fig. 1, and this driving is what we call TPC and in general, it gives an accuracy $N > 1$. The key is to perform the driving autonomously without external temporal control. In the following we introduce a model for this with more details on the formalism of TPC in Sec. A of the Appendix.

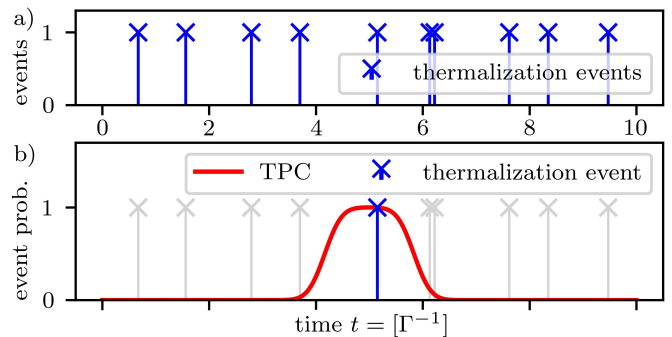


FIG. 1. The two plots qualitatively show the ticking behavior of two clocks with respect to parameter time t (horizontal axis). The ticks of such a clock are generated by individual thermalization events at rate Γ (vertical stripes). Figure a) sketches a generic example, where these events are Poisson distributed. In b), temporal probability concentration is shown, another process through which the probability of an thermalization event can be suppressed at times (shown by the hat-shaped curve). As a result, the probability density for a tick can be concentrated around a desired average value, here $5\Gamma^{-1}$, with the tick time uncertainty of order Γ^{-1} bounded by the width of the TPC window.




FIG. 2. We illustrate the oversampling regime with a pendulum in a weakly lit environment. The two sources of entropy production for this clock are, firstly, the friction within the clockwork itself and, secondly, the light-matter interaction to track the position of the pendulum. While the entropy from friction can in principle be vanishingly small, the one from observation is fundamental and can not be made zero without losing the ability to measure time. The plot shows the elementary ticking events of this clock as a function of time, i.e., the photons reflected off the pendulum when it is close to its maximum deflection, and the pendulum oscillation (TPC). In the oversampling regime, the photon rate is much higher than the frequency of the TPC, and this allows to use the TPC period to define ticks.

Examples of clocks following this scheme are given in [8–10, 18, 23–25], with a particularly illustrative one given in the first mentioned reference: a three-qubit system where one of the qubits couples dissipatively to the electromagnetic field with strength Γ , emitting photons when it decays, and these events are then counted as ticks. Alone, the qubit would undergo exponential decay and give accuracy $N = 1$, but here it is autonomously driven by the other two qubits, which themselves are coupled to out-of-equilibrium heat baths. The Hamiltonian coupling the three qubits is the periodic process which makes the effective decay probability of this clock time-

dependent, and leads to an enhanced accuracy $N > 1$ due to TPC. More details are in Sec. A 1 of the Appendix. Many macroscopically sized clocks, be it pendulums or atomic clocks, operate in an oversampling regime, where multiple irreversible events occur per TPC cycle (see Figure 2). In this regime, the thermodynamic cost of a clockwork often becomes obscure, as both the dissipation due to the macroscopic number of irreversible events and the TPC have to be accounted for. Atomic clocks, for example, do not count photons as a way to tell time, rather they use the oversampled coherent oscillations of a maser tuned some stable reference atomic transition to estimate the TPC frequency. The time-scale Γ of the fundamental ticking events appearing in eq. (1) is therefore not the limiting factor to the accuracy of atomic clocks, the stability of the coherent oscillation of the electromagnetic field is, i.e., the TPC stability. In atomic clocks, accuracy is examined using Allan Variance which captures the stability of the TPC oscillation over many different time-scales [26–28]. Quantum projection noise, thermal noise but also natural drifts in the experimental setup are what affect the stability of atomic clocks [29, 30]. A summary with some key references on the working principle of atomic clocks can be found in Sec. B of the Appendix.

So far, we have established that every clock is subject to irreversible processes and that through TPC, they can increase their accuracy. In the following, we introduce a mathematical model to describe clocks on a quantum scale and where the two contributions (a) irreversible ticks and (b) temporal probability concentrations have an explicit representation in the equations of motion. Eventually, this framework allows us to formulate the fundamental trade-off between the accuracy and resolution of clocks, dictated by thermodynamics.

Model. Quantum clocks [8–10, 18, 23–25] only weakly coupled to a memoryless environment (sufficiently large thermal baths are one such environment) can be described by a Lindblad master-equation [31]. If we are again talking about ticking clocks, their state can be Fourier-decomposed by introducing a free counting field χ [32–34],

$$\rho(t, \chi) = \sum_n \rho^{(n)}(t) e^{in\chi}. \quad (3)$$

Each non-normalized density matrix $\rho^{(n)}(t)$ can be thought of as the system’s state conditioned on n ticks having already occurred. In a setting where the interactions with the environment are memoryless, the ticks are produced by linear jump operators J_j , the generators of the process (a). Aside from this, the clock is subject to a general open quantum system’s evolution with Lindblad operator $\mathcal{L}_{\text{no tick}}$ which is the generator of TPC, i.e., process (b). We are interested in the statistics of the time T between any two successive ticks n and $n + 1$. These statistics may differ from each tick to the next one because the initial state changes, and this also means that the accuracy N and frequency ν can change with each tick n . However, the trade-off theorem is agnostic to the

value n , and holds for all pairs of N and ν . Without loss of generality, we can therefore assume that the n th tick happened at time $t = 0$ and that $n = 0$, and we can look at the evolution of the conditional state $\rho^{(0)}(t)$ that exactly $n = 0$ ticks have occurred. Given the initial state $\rho^{(0)}(0)$ of the clock from eq. (3), the evolution is entirely determined by

$$\dot{\rho}^{(0)}(t) = \mathcal{L}_{\text{no tick}}[\rho^{(0)}(t)] - \frac{1}{2} \sum_j \left\{ J_j^\dagger J_j, \rho^{(0)}(t) \right\}. \quad (4)$$

The right-most term of eq. (4) produces the ticks and is therefore responsible for the process (a) while $\mathcal{L}_{\text{no tick}}$ generates the TPC, i.e., process (b). In this sense, one may attempt to separate (a) and (b) into two independent processes, as in [10]; however the back-action of the tick channel always affects the clock evolution which is an inherent feature of clocks that operate on a quantum scale.

Ticking statistics. Together with the initial state, eq. (4) defines the evolved state $\rho^{(0)}(t)$ and therefore entirely determines how the tick time random variable T is distributed, because the cumulative probability $P[t \leq T]$ that no tick has occurred up to time t equals the trace $P[t \leq T] = \text{tr} \rho^{(0)}$. To highlight the influence of TPC, we can equivalently express the cumulative tick probability as a function of the conditional tick rate $P[T = t | T \geq t]$, defined as the instantaneous tick probability at time t conditioned on the tick not having happened before. We can obtain this rate by working with the clock’s state conditioned on not having decayed $\rho^{\text{no tick}} = \rho^{(0)} / \text{tr} \rho^{(0)}$. Then, the conditional tick rate equals the trace $\text{tr}(V \rho^{\text{no tick}}(t'))$, where $V = \sum_j J_j^\dagger J_j$ is the positive operator generating the clock’s ticks [35], and we find the relationship

$$P[t \leq T] = \exp\left(-\int_0^t dt' \text{tr}(V \rho^{\text{no tick}}(t'))\right), \quad (5)$$

which we further outline in the Appendix A 2.

The master-equation description of quantum ticking clocks allows us to formalize the statement made in the accuracy-resolution trade-off theorem. The tick channel is governed by the positive operator V , whose spectral decomposition reveals the time-scales involved in the decay. We define the rate Γ as the fastest one of them,

$$\Gamma := \|V\|_{\max} = \max_{\rho \in \mathcal{S}(\mathcal{H}_C)} \text{tr} V \rho, \quad (6)$$

where the maximum is taken over all possible clock states $\rho \in \mathcal{S}(\mathcal{H}_C)$. This is consistent with the special case of exponential decay (as used in [9, 17]) with rate Γ , since there, the tick generator is given by a single jump operator $J = \sqrt{\Gamma} |0\rangle\langle 1|$. If a clock is described by eq. (4) and produces its ticks by means of the generators J_j , then the accuracy is limited by the resolution, regardless of how well a possible clockwork in the background works. The resulting bound is eq. (1) from the accuracy-resolution




FIG. 3. In this figure, we plot the accuracy N on the vertical axis versus the resolution ν on the horizontal axis. The red curve highlights the upper bound for the accuracy $N = \Gamma^2/\nu^2$ given by the accuracy-resolution trade-off theorem. Below, we have the blue curve, indicating the accuracy and resolution that can be achieved by averaging the thermal reference clock with fundamental tick rate Γ . The shaded area between the two curves contains all the accuracy-resolution tuples that cannot be achieved by classically averaging a rate Γ stochastic process, but which are still allowed by the upper bound. How close one can get to the red curve with quantum clocks and whether classical clocks are constrained to be below the blue line are open questions.

trade-off theorem which we recall here for completeness,

$$N \leq \frac{\Gamma^2}{\nu^2},$$

and prove in the following.

Proof. The key observation is that no clock can on average tick faster than the decay process that mediates the ticks allows for. If the elementary ticks are generated by an ensemble of jump operators J_j , then the fastest such rate is given by Γ (see eq. (6)). A single such channel produces ticks that are exponentially distributed and as a consequence, we can reduce the generic form from eq. (5) to the special case of exponential decay, where the exponent reduces to $-\Gamma \int dt' p(t')$, and $p(t)$ is a measure of the clock's state population that can decay. Regardless of how the TPC modulates $p(t)$, the variance σ^2 of the tick can never be smaller than Γ^{-2} , the variance of exponential decay. This is a manifestation of the fact that no clock can tick faster than its underlying decay process. Now that the variance σ^2 of the tick time distribution is bounded from below by Γ^{-2} , the main theorem follows. We refer the reader to Sec. C of the Appendix for a detailed account of the proof. \square

In this letter, we have analyzed quantum timekeeping devices through a thermodynamic lens, where their tick generation can be decomposed into two processes, (a) a stochastic process which irreversibly produces the ticks and (b) temporal probability concentration (TPC) through a clockwork which controls when these elementary ticks occur. Then, we asserted that there is a fundamental trade-off between the clock's accuracy and resolution (see trade-off theorem but also Fig. 3), stating that the number of times a clock can tick until it goes wrong by one tick is universally bounded by the inverse of the clock's resolution squared.

Atomic and optical clocks. Timekeeping devices fundamentally require stochastic thermalization events to measure time, and we have shown that for the class of clocks using those events directly to define ticks, the trade-off from eq. (1) applies. Macroscopic clocks such as atomic and optical ones work in a different regime, where the TPC is sampled by irreversible events and ticks are defined not by event numbers directly, but rather those are used to estimate the oscillatory TPC process frequency. This frequency is usually downsampled to a lower frequency, like for example the 10 MHz, or the 1 Hz standard that is then used to generate an electrical signal to read out the ticks at said resolution [36]. Our work points out some of the challenges that have to be overcome for optical clocks being used to produce ticks at THz resolution: for one, ultra-fast electronics that can generate an electric signal using only a single photon per oscillation of the e.m. field. For another, this gives an estimate of the power required from a laser to create such a photon flux. While for time-standards whose main goal it is to provide long-time stability, power consumption may not be the primary concern, there are other uses for clocks where energy-efficiency matters. For example, quantum technologies require accurate high-resolution timers that do not disturb the fragile state of the quantum system through heat dissipation [37–40]. For building noise-robust quantum devices, it may thus become unavoidable to account for the thermodynamic resources consumed by clocks, which is where we expect quantum clocks to outperform their macroscopic classical counterparts.

Achievability. Good clocks excel by resolving time well while at the same time being highly accurate, i.e., maximizing both ν and N ; ideally, they do so at optimal thermodynamic costs far away from the oversampled regime and only with a single thermalization event per tick. As a foundational question, we may ask, does there exist an appropriate clock Lindblad operator \mathcal{L} , through which it is possible to saturate the bound $N = \Gamma^2/\nu^2$? And the answer is: No, at least not for finite systems. As it turns out, saturating this bound amounts to a time-dependent Lindblad-operator, which instantaneously rotates the clock state from a subspace of states where it can't tick onto a state from which it decays at rate Γ (see proof in Sec. C of the Appendix). The TPC in this case would be ideal as it concentrates the decay event to the most narrow time-window possible, the one given by the underlying stochastic process, but quantum speed limits prohibit such an instantaneous state rotation for systems finite in energy and dimension [41–43]. A weaker question we can ask is whether at least the scaling $N = O(\Gamma^2/\nu^2)$ can be reached, and which resources are required for this. Without imposing any restriction on the clock's Lindbladian aside from time-independence, there are quantum clocks which asymptotically reach the squared scaling of N for large dimensions of the clock Hilbert space [23]. To do so, they require highly coherent states, whose generation using only thermal resources is technically in-

finitely expensive from an entropic perspective [40] and the approximation with finite resources is an open problem. Of interest in the field of thermodynamics are the autonomous quantum clocks which only require thermal resources to run [9, 10, 12, 24, 35], in particular no external control but also no coherence in the initial state. It is ongoing research, whether it is possible for such clocks to approach the optimal accuracy-resolution scaling.

Questions about fundamental precision limits are generally of great interest. The accuracy of clocks falls into this category and closely related problems have been examined in the field of (quantum) stochastic thermodynamics under the name of thermodynamic uncertainty relations (TUR) [44–46] and kinematic uncertainty relations (KUR) [47–49]. Future work exploring connections between timekeeping and the TUR / KUR has to reveal whether a quantum thermodynamic advantage close to the optimal accuracy-resolution bound $N = \Gamma^2/\nu^2$ is in principle possible. Both superconducting circuits and optomechanical systems are promising platforms to test the achievability of the optimal accuracy-resolution relation

while accounting for the resources and ensuring that they do not introduce a hidden clock through a backdoor.

Acknowledgements. We wish to acknowledge fruitful discussions and feedback from Ralph Silva, Nuriya Nurgalieva, Renato Renner, Maximilian Edward Lock and Jake Xuereb. F.M., P.E. and M.H. acknowledge the financial support from the European Commission under the Horizon Europe project ASPECTS (Grant No. 101080167) and the ERC consolidator grant COCOQUEST (Grant No. 101043705). F.M. acknowledges the SEMP scholarship from Movetia for his research stay at Atominstut, Technische Universität Wien. E.S. acknowledges the support from the Austrian Science Fund (FWF) through the START project Y879-N27, the ESQ Discovery grant “Emergence of physical laws: From mathematical foundations to applications in many body physics” and the ERC consolidator grant COCOQUEST (Grant No. 101043705). P.E. and M.H. further acknowledge funds from the FQXi (FQXi-IAF19-03-S2) within the project “Fueling quantum field machines with information”.

* florianmeier256@gmail.com

† e.schwarzahans@gmx.at

‡ paul.erker@tuwien.ac.at

§ marcus.huber@tuwien.ac.at

- [1] L. Essen and J. V. L. Parry, An Atomic Standard of Frequency and Time Interval: A Cæsium Resonator, *Nature* **176**, 280 (1955).
- [2] A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, and P. Schmidt, Optical atomic clocks, *Reviews of Modern Physics* **87**, 637 (2015).
- [3] F. Riehle, Towards a redefinition of the second based on optical atomic clocks, *Comptes Rendus Physique* **16**, 506 (2015), the measurement of time / La mesure du temps.
- [4] T. Bothwell, C. J. Kennedy, A. Aeppli, D. Kedar, J. M. Robinson, E. Oelker, A. Staron, and J. Ye, Resolving the gravitational redshift across a millimetre-scale atomic sample, *Nature* **602**, 420 (2022).
- [5] S. Ranković, Y.-C. Liang, and R. Renner, Quantum clocks and their synchronisation - the Alternate Ticks Game, [arXiv:1506.01373 \[quant-ph\]](https://arxiv.org/abs/1506.01373) (2015).
- [6] R. Renner and S. Stupar, eds., *Time in Physics* (Springer International Publishing, 2017).
- [7] S. Stupar, C. Klumpp, R. Renner, and N. Gisin, Performance of stochastic clocks in the Alternate Ticks Game, [arXiv:1806.08812 \[quant-ph\]](https://arxiv.org/abs/1806.08812) (2018).
- [8] M. P. Woods, Autonomous Ticking Clocks from Axiomatic Principles, *Quantum* **5**, 381 (2021).
- [9] P. Erker, M. T. Mitchison, R. Silva, M. P. Woods, N. Brunner, and M. Huber, Autonomous Quantum Clocks: Does Thermodynamics Limit Our Ability to Measure Time?, *Phys. Rev. X* **7**, 031022 (2017).
- [10] E. Schwarzahans, M. P. E. Lock, P. Erker, N. Friis, and M. Huber, Autonomous Temporal Probability Concentration: Clockworks and the Second Law of Thermodynamics, *Phys. Rev. X* **11**, 011046 (2021).
- [11] G. J. Milburn, The thermodynamics of clocks, *Contemporary Physics* **61**, 69 (2020).
- [12] S. K. Manikandan, Autonomous quantum clocks using athermal resources, *Phys. Rev. Res.* **5**, 043013 (2023).
- [13] D. Cilluffo, Statistical time-domain characterization of non-periodic optical clocks, *Quantum* **6**, 764 (2022).
- [14] A. N. Pearson, Y. Guryanova, P. Erker, E. A. Laird, G. A. D. Briggs, M. Huber, and N. Ares, Measuring the Thermodynamic Cost of Timekeeping, *Phys. Rev. X* **11**, 021029 (2021).
- [15] M. Berholts, R. Knut, R. Stefanuik, H. Wikmark, S. Saha, and J. Söderström, Quantum watch and its intrinsic proof of accuracy, *Phys. Rev. Research* **4**, 043041 (2022).
- [16] X. He, P. Pakkiam, A. A. Gangat, M. J. Kewming, G. J. Milburn, and A. Fedorov, Effect of measurement back-action on quantum clock precision studied with a superconducting circuit, *Phys. Rev. Appl.* **20**, 034038 (2023).
- [17] A. C. Barato and U. Seifert, Cost and precision of brownian clocks, *Phys. Rev. X* **6**, 041053 (2016).
- [18] M. P. Woods, R. Silva, G. Pütz, S. Stupar, and R. Renner, Quantum Clocks are More Accurate Than Classical Ones, *PRX Quantum* **3**, 010319 (2022).
- [19] Y. Yang and R. Renner, Ultimate limit on time signal generation, [arXiv:2004.07857 \[quant-ph\]](https://arxiv.org/abs/2004.07857) (2020).
- [20] R. S. Ingarden, A. Kossakowski, and M. Ohya, *Information Dynamics and Open Systems* (Springer Netherlands, 1997).
- [21] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, 2007).
- [22] A. Klenke, *Probability Theory* (Springer International Publishing, 2020).
- [23] M. P. Woods, R. Silva, and J. Oppenheim, Autonomous Quantum Machines and Finite-Sized Clocks, *Annales Henri Poincaré* **20**, 125 (2019).

- [24] P. Erker, [The quantum hourglass. approaching time measurement with quantum information theory](#), Master Thesis, ETH Zürich (2014).
- [25] A. P. T. Dost and M. P. Woods, Quantum advantages in timekeeping: dimensional advantage, entropic advantage and how to realise them via berry phases and ultra-regular spontaneous emission, [2303.10029 \[quant-ph\]](#) (2023).
- [26] D. W. Allan, Statistics of atomic frequency standards, [Proceedings of the IEEE](#) **54**, 221 (1966).
- [27] D. W. Allan, N. Ashby, and C. C. Hodge, [The Science of Timekeeping Application Note 1289](#), Hewlett Packard (1997).
- [28] W. Riley and D. Howe, [Handbook of Frequency Stability Analysis](#) (2008).
- [29] N. Shiga, M. Mizuno, K. Kido, P. Phoonthong, and K. Okada, Accelerating the averaging rate of atomic ensemble clock stability using atomic phase lock, [New Journal of Physics](#) **16**, 073029 (2014).
- [30] M. Schulte, C. Lisdat, P. O. Schmidt, U. Sterr, and K. Hammerer, Prospects and challenges for squeezing-enhanced optical atomic clocks, [Nature Communications](#) **11**, 5955 (2020).
- [31] G. Lindblad, On the generators of quantum dynamical semigroups, [Communications in Mathematical Physics](#) **48**, 119 (1976).
- [32] G. Schaller, [Open Quantum Systems Far from Equilibrium](#) (Springer International Publishing, 2014).
- [33] M. Bruderer, L. D. Contreras-Pulido, M. Thaller, L. Sironi, D. Obreschkow, and M. B. Plenio, Inverse counting statistics for stochastic and open quantum systems: the characteristic polynomial approach, [New Journal of Physics](#) **16**, 033030 (2014).
- [34] G. T. Landi, M. J. Kewming, M. T. Mitchison, and P. P. Potts, Current fluctuations in open quantum systems: Bridging the gap between quantum continuous measurements and full counting statistics, [arXiv:2303.04270 \[quant-ph\]](#) (2023).
- [35] F. Meier, [Performance limits of decay clocks: Fundamental accuracy and resolution limits of quantum clocks with exponential ticking mechanism](#), Master Thesis, ETH Zürich (2022).
- [36] B. L. S. Marlow and D. R. Scherer, A Review of Commercial and Emerging Atomic Frequency Standards, [IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control](#) **68**, 2007 (2021).
- [37] H. Ball, W. D. Oliver, and M. J. Biercuk, The role of master clock stability in quantum information processing, [npj Quantum Information](#) **2**, 16033 (2016).
- [38] L. Buffoni, S. Gherardini, E. Zambrini Cruzeiro, and Y. Omar, Third law of thermodynamics and the scaling of quantum computers, [Phys. Rev. Lett.](#) **129**, 150602 (2022).
- [39] M. O. Tholén, R. Borgani, G. R. Di Carlo, A. Bengtsson, C. Krizan, M. Kudra, G. Tancredi, J. Bylander, P. Delsing, S. Gasparinetti, and D. B. Haviland, Measurement and control of a superconducting quantum processor with a fully integrated radio-frequency system on a chip, [Review of Scientific Instruments](#) **93**, 104711 (2022).
- [40] P. Taranto, F. Bakhshinezhad, A. Bluhm, R. Silva, N. Friis, M. P. Lock, G. Vitagliano, F. C. Binder, T. Debarba, E. Schwarzhans, F. Clivaz, and M. Huber, Landauer Versus Nernst: What is the True Cost of Cooling a Quantum System?, [PRX Quantum](#) **4**, 010332 (2023).
- [41] L. Mandelstam, I. Tamm, and I. E. Tamm, The Uncertainty Relation Between Energy and Time in Non-relativistic Quantum Mechanics, in [Selected Papers](#), edited by B. M. Bolotovskii, V. Y. Frenkel, and R. Peierls (Springer Berlin Heidelberg, Berlin, Heidelberg, 1991) pp. 115–123.
- [42] N. Margolus and L. B. Levitin, The maximum speed of dynamical evolution, [Physica D: Nonlinear Phenomena](#) **120**, 188 (1998), proceedings of the Fourth Workshop on Physics and Consumption.
- [43] S. Deffner and S. Campbell, Quantum speed limits: from heisenberg’s uncertainty principle to optimal quantum control, [Journal of Physics A: Mathematical and Theoretical](#) **50**, 453001 (2017).
- [44] A. C. Barato and U. Seifert, Thermodynamic Uncertainty Relation for Biomolecular Processes, [Phys. Rev. Lett.](#) **114**, 158101 (2015).
- [45] J. M. Horowitz and T. R. Gingrich, Thermodynamic uncertainty relations constrain non-equilibrium fluctuations, [Nature Physics](#) **16**, 15 (2020).
- [46] T. V. Vu and K. Saito, Thermodynamics of Precision in Markovian Open Quantum Dynamics, [Phys. Rev. Lett.](#) **128**, 140602 (2022).
- [47] I. D. Terlizzi and M. Baiesi, Kinetic uncertainty relation, [Journal of Physics A: Mathematical and Theoretical](#) **52**, 02LT03 (2018).
- [48] V. T. Vo, T. V. Vu, and Y. Hasegawa, Unified thermodynamic–kinetic uncertainty relation, [Journal of Physics A: Mathematical and Theoretical](#) **55**, 405004 (2022).
- [49] K. Prech, P. Johansson, E. Nyholm, G. T. Landi, C. Verdozzi, P. Samuelsson, and P. P. Potts, Entanglement and thermokinetic uncertainty relations in coherent mesoscopic transport, [Phys. Rev. Research](#) **5**, 023155 (2023).
- [50] H. Carmichael, [An Open Systems Approach to Quantum Optics](#) (Springer Berlin Heidelberg, 1993).
- [51] M. Eichhorn, [Laser Physics](#) (Springer International Publishing, 2014).
- [52] F. G. Major, [The Quantum Beat](#) (Springer New York, 2007).
- [53] S. Bize, The unit of time: Present and future directions, [Comptes Rendus Physique](#) **20**, 153 (2019).

APPENDICES

Appendix A: Temporal probability concentration

In this section, we elaborate on the technical details regarding the concept of temporal probability concentration (TPC) which we have introduced in the main text of this letter. This concept has first been explicitly mentioned in [10], but implicitly, it has been applied in a variety of works [9, 18, 23, 25]. In the main text, a distinction of tick generation into two separate processes was made: a) the irreversible process whose occurrence *defines* the tick, and b) an additional filter mechanism, which temporally concentrates the probability of when a tick, i.e., process a) occurs. We refer to the latter as TPC.

1. An example

As an illustration, let us take a fixed exponential decay process, i.e., a physical system comprising two states, say an excited one and the ground state, and the excited state decays with a fixed rate Γ into the ground state. If we were to initialize this system in the excited state, its tick rate would be constantly equals to Γ and the resulting tick probability density $P[T = t] = \Gamma e^{-\Gamma t}$ would be exponential. Not unsurprisingly, this clock does not excell in accuracy, which manifests itself in the accuracy $N = 1$, as already elaborated in the main text. Let us now introduce an additional dynamical process, for example the driving by an autonomous thermal machine as in [9]. Such an interaction could first rotate the clock state from the ground state from which it can not decay into the excited state, from which it can decay, leading to an effective decay rate not always equals Γ , but now modulated dynamically by the autonomous thermal machine. This time-dependency is the TPC of the tick process, and in the example of [9] it leads to an accuracy $N > 1$.

We here summarize the clock example from [9] to illustrate the general model for clocks introduced with the equations of motion in eq. (4) in the main text. In general, the two-level system in the clock described above can be extended to a d -dimensional ladder which allows to delay the decay process even further. To be more specific, the clock model comprises the ladder \mathcal{H}_L spanned by the states $|0\rangle_L, |1\rangle_L, \dots, |d-1\rangle_L$, and the thermal machine. The thermal machine itself is made up of two qubits, a *cold* one \mathcal{H}_C with the states $|0\rangle_C, |1\rangle_C$ and a *hot* one \mathcal{H}_H with states $|0\rangle_H, |1\rangle_H$. The interactions between the ladder and the two qubits is described by the Hamiltonian $H = H_0 + H_{\text{int}}$, where

$$H_0 = \omega_C |1\rangle\langle 1|_C + \omega_H |1\rangle\langle 1|_H + \sum_{n=0}^{d-1} n\omega_L |n\rangle\langle n|_L, \quad (\text{A1})$$

is the system's free Hamiltonian, with ω_C the energy-splitting of the cold qubit, ω_H that of the hot qubit, and

ω_L that of the ladder. The term

$$H_{\text{int}} = g \sum_{n=0}^{d-1} \left(|10\rangle\langle 01|_{CH} \otimes |n+1\rangle\langle n|_L + \text{h.c.} \right) \quad (\text{A2})$$

describes the population exchange between the two-qubit thermal machine and the ladder system. Furthermore, there are the thermal dissipators \mathcal{L}_C , and \mathcal{L}_H modelling the interactions of the cold qubit with a cold bath at inverse temperature β_C and those of the hot qubit with its hot bath at inverse temperature β_H . The dissipators are of the form $\mathcal{L}_C = n_C \gamma_C \mathcal{D}[|1\rangle\langle 0|_C] + (1+n_C) \gamma_C \mathcal{D}[|0\rangle\langle 1|_C]$, where n_C is the photon number of the cold bath given by Bose-Einstein statistics, $n_C = (e^{\beta_C \omega_C} - 1)^{-1}$, and γ_C is the coupling rate of the cold bath to the cold qubit. The dissipator for the hot bath is analogous but with the subscript H instead of C . The terms \mathcal{D} are as usually defined as $\mathcal{D}[L] = L \circ L^\dagger - \frac{1}{2} \{L^\dagger L, \circ\}$. All these expressions together give rise to the Lindbladian part $\mathcal{L}_{\text{no tick}}$ responsible for TPC,

$$\mathcal{L}_{\text{no tick}} = -i[H, \circ] + \mathcal{L}_C + \mathcal{L}_H. \quad (\text{A3})$$

The tick generation then comes from the ladder decaying from its top level state into the ground state through the tick generating operator J , defined by

$$J = \sqrt{\Gamma} |0\rangle\langle d-1|_L, \quad (\text{A4})$$

where Γ is the coupling strength of the ladder's transition $|d-1\rangle_L \rightarrow |0\rangle_L$ to the environment, i.e., the decay rate. If we go through the calculations provided in reference [9], we find that this clock's accuracy in the regime $\Gamma, g \ll \gamma_H, \gamma_C$, and $g \ll \omega_C, \omega_L$, and under the resonance condition $\omega_C + \omega_L = \omega_H$ grows with the dimension d but the resolution decreases. In the asymptotic limit, as cold bath temperatures goes to absolute zero, $\beta_C^{-1} \rightarrow 0$, the accuracy grows linearly in the dimension, $N \propto d$, and the resolution inverse linearly, $\nu \propto d^{-1}$. The resulting accuracy-resolution relationship is also inverse linear for this clock, $N \propto \nu^{-1}$; as the authors in [9] state, the clock in this regime behaves essentially classically, hence, not unexpectedly, we find the classical trade-off from eq. (2) in the main text.

2. General formalism

In the following we approach the concept of TPC from two different perspectives, firstly from the equations of motion by starting from eq. (4), describing explicitly the evolution of a given tick state, and secondly from a formal probability theoretic paradigm.

TPC as an emergent property from the equations of motion. For convenience, let us recall the equation of motion (4) from the main text,

$$\dot{\rho}^{(0)}(t) = \mathcal{L}_{\text{no tick}}[\rho^{(0)}(t)] - \frac{1}{2} \sum_j \left\{ J_j^\dagger J_j, \rho^{(0)}(t) \right\}. \quad (\text{A5})$$

Ticks, by definition, are the transitions generated by the operators J_j , and this is the process a) in our list from before. This means, if the clock system undergoes any one of the stochastic transitions J_j , this is counted as a tick.

We can take the trace of the equation of motion (A5), and we obtain on the left-hand side the derivative of the cumulative tick probability density,

$$\text{tr} \left\{ \dot{\rho}^{(0)}(t) \right\} = \frac{d}{dt} P[t \leq T]. \quad (\text{A6})$$

For a well-behaved, here, continuously differentiable, cumulative probability $P[t \leq T]$, the derivative equals the probability density function (PDF) $P[T = t]$. This allows us to identify the trace of the right-hand side of eq. (A5) as the probability density,

$$P[t = T] = \text{tr} \left\{ V \rho^{(0)}(t) \right\}, \quad (\text{A7})$$

where $V = \sum_j J_j^\dagger J_j$ as defined in the main text. This probability distribution completely describes the statistics of the tick considered. Average time between ticks μ and variance of the time between ticks σ^2 are defined with respect to this PDF. An related quantity, which is of particular relevance from an operational perspective is the conditional tick probability density: the probability density $P[T = t|T \geq t]$ that the tick occurs at time t conditioned on the fact that it has not ticked before that time t . We can calculate the conditional PDF using

$$P[T = t|T \geq t] = \frac{P[T = t \wedge T \geq t]}{P[T \geq t]} \quad (\text{A8})$$

$$= \frac{P[T = t]}{P[T \geq t]}, \quad (\text{A9})$$

which reveals that the conditional tick PDF equals the trace that appears in eq. (5) in the main text,

$$P[T = t|T \geq t] = \text{tr}(V \rho^{\text{no tick}}(t')). \quad (\text{A10})$$

In this form, TPC becomes particularly apparent: for a clock without any internal dynamics, i.e., $\mathcal{L}_{\text{no tick}} \equiv 0$ where ticks are generated through exponential decay, the conditional tick PDF is constant, $P[T = t] = \Gamma$, where Γ is the rate of the exponential decay. This is the example, where the stochastic thermalization event that generates the tick is not temporally concentrated, hence the time-independence of the conditional tick PDF. In contrast, for a non-trivial clockwork as for example presented in the references [9, 18, 23, 25], the conditional tick PDF has time-dependency, which eventually leads to a non-exponential tick probability with accuracy $N > 1$.

TPC as a conditional probability density. The previous paragraph introduced the notion of the conditional tick PDF $P[T = t|T \geq t]$ that a tick happens at time t conditioned on the tick not having happened before. Equation (5) in the main text relates this expression to

the tick cumulative tick probability density $P[T \geq t]$ via the more general identity,

$$P[T \geq t] = \exp \left(- \int_0^t d\tau P[T = t|T \geq t] \right). \quad (\text{A11})$$

Here we would like to derive this expression and elaborate. Definition of cumulative probability distributions (CDFs) and PDFs ensure, that

$$\frac{d}{dt} P[T \geq t] = -P[T = t], \quad (\text{A12})$$

so long as the derivative is well-defined and continuous. If we then use the conditional probability law from eq. (A8) of this supplemental material, we find that

$$\frac{d}{dt} P[T \geq t] = -P[T \geq t] P[T = t|T \geq t]. \quad (\text{A13})$$

The solution of this differential equation is the exponential expression as generally given in eq. (A11). This shows the relation in eq. (5).

Appendix B: Working principle of atomic clocks

Atomic clocks are the technological state-of-the-art when it comes to timekeeping and in the letter, we discuss the relevance of the trade-off theorem to atomic clocks. In this section, we give the interested reader a brief description of the working principle of atomic clocks and refer them to more specialized references.

Frequency estimation vs. tick generation. Atomic clocks use the period of the coherent electromagnetic field oscillation in a laser to measure time and they use a reference frequency, for example electronic transitions in an atom, to stabilize the laser's frequency. Commonly, atomic clocks are used to generate a stable frequency reference. The coherent light field of the laser which is the basis of every atomic clock is generated using stimulated emission (see e.g. [50] for an open quantum system's approach or [51], for an applied textbook). In practice, the laser frequency is inevitably subject to drifts and noise due to interactions with the environment and inherent imperfectness of the constituents (e.g. Brownian motion). To correct for this, elaborate feedback techniques have been developed, e.g., the hyperfine ground-state transition of an Caesium-133 atom is used as a frequency reference to stabilize the laser [1, 36, 52, 53]. In our model introduced in the main text, all frequencies are idealized to be perfectly stable; in particular also the TPC's oscillatory frequency is assumed to be unchanging. We reveal that despite these idealizations, such a clock can not be perfect, due to the underlying stochasticity of the thermal processes. Formally, we show this with the accuracy-resolution trade-off theorem.

However, frequency stabilization is only part of the story required for accurate timekeeping, as an additional

process is required to produce ticks. In the oversampling regime atomic clocks work in, ticks are not uniquely defined but one may choose to define ticks as the zero-crossings or the maxima of the oscillation. In this regime, the elementary stochastic events (here: photoemission) are not directly used as a tick definition anymore but rather as a means to sample the TPC oscillation, which then defines ticks as for example one period of the oscillation. Only in the regime where such a clock is run with a single photon per oscillation of the electromagnetic field, and the photon defines a tick, does the trade-off theorem impose a practical restriction on the clock's performance.

Appendix C: Proof of the accuracy-resolution trade-off

Before we get started with the proof of the accuracy-resolution trade-off theorem (which is a generalization of the one shown in [35]), some preliminary notation has to be established. When it comes to the cumulative non-tick probability $P[t < T]$ as in eq. (5) of the main text, the trace $\text{tr } V\rho_C(t)$ can be rewritten as

$$\text{tr}(V\rho_C(t)) = \Gamma p(t), \quad (\text{C1})$$

where $p : \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ is a smooth function that takes values between 0 and 1. This comes from the fact that Γ is the maximum of the expression over all states ρ and the function can only take positive values because V is a positive operator. This expression puts us into the position, where we can write

$$P[t \leq T] = \exp\left(-\Gamma \int_0^t dt' p(t')\right), \quad (\text{C2})$$

which is formally the exact same expression that one would obtain in the case where the clock's elementary ticks are exponential decay. We will from now on refer to $p(t)$ as the *top-level population*. The remaining sections are an adapted version of the proof of the accuracy-resolution theorem originally derived in [35].

1. The Heaviside- Θ population

Definition 1 (Heaviside populations). *The family of Heaviside populations comprises all the functions*

$$p^\Theta : \mathbb{R}_{\geq 0} \rightarrow [0, 1] \quad (\text{C3})$$

$$t \mapsto \Theta(t - t_0), \quad (\text{C4})$$

where $t_0 \geq 0$.

Properties of the Heaviside population. Such a top-level population is unphysical because of its discontinuity at time $t = t_0$. This is not an obstacle, though, as we only use the Heaviside population to prove an upper bound for the accuracy of a decay clock. No claim

is made whether we can physically obtain the Heaviside population. The non-tick probability $P_\Theta[t \leq T]$ and tick probability density $p_{\text{tick}}^\Theta(t)$ associated to this population are given by (see eq. (C2))

$$P_\Theta[t \leq T] = \begin{cases} 1 & t \leq t_0 \\ e^{-\Gamma(t-t_0)} & t \geq t_0, \end{cases} \quad (\text{C5})$$

and

$$p_{\text{tick}}^\Theta(t) = \begin{cases} 0 & t \leq t_0 \\ \Gamma e^{-\Gamma(t-t_0)} & t \geq t_0. \end{cases} \quad (\text{C6})$$

See Fig. 4 for a visualization of these functions. Using the analytic solution of the integral over an exponential, $\int_0^\infty dx e^{-ax} = 1/a$, we can calculate the accuracy and resolution for clocks working the Heaviside population.

Lemma 2. *The Heaviside population $p^\Theta(t) = \Theta(t - t_0)$ has an accuracy and resolution given by*

$$N = (1 + \Gamma t_0)^2, \text{ and } \nu = \frac{1}{t_0 + \frac{1}{\Gamma}}, \quad (\text{C7})$$

leading to an accuracy-resolution relation

$$N = \frac{\Gamma^2}{\nu^2}. \quad (\text{C8})$$

Before we prove Lemma 2, let us put the Heaviside population into perspective. Given the accuracy-resolution trade-off theorem the equality $N = \Gamma^2/\nu^2$ for clocks with Heaviside population tells us that this top-level population profile is ideal in the accuracy-resolution sense. This is no coincidence: a Heaviside population with offset at t_0 requires a clock in the background which perfectly knows t_0 . The population then describes a ladder conditioned on not having decayed, whose top-level is populated precisely at $t = t_0$, and stays there. This would require instantaneous external driving at $t = t_0$,




FIG. 4. This graph visualizes the ticking statistics for the Heaviside population. The solid line in the above figure visualizes the tick probability density, whereas the dashed line stands for the cumulative non-tick probability. The time axis scales in inverse units of the decay rate Γ and the population in the example is chosen such that on average the decay occurs at $\mu = 4.5\Gamma$.

i.e., a perfect background clock. The exponential decay channel coupled to the top-level then smears out the Heaviside profile of $p^\ominus(t)$ into an exponential decaying tick probability density $p_{\text{tick}}^\ominus(t)$ of width $1/\Gamma$, giving rise to the results of Lemma 2.

Proof. The average tick time is given by an integral over $tp_{\text{tick}}^\ominus(t)$ and splits into two parts

$$\mu = \int_0^\infty dt tp_{\text{tick}}^\ominus(t) = 0 + \int_{t_0}^\infty dt te^{-\Gamma(t-t_0)} \quad (\text{C9})$$

$$= t_0 + \frac{1}{\Gamma}. \quad (\text{C10})$$

Thus, the resolution is given by $\nu = (t_0 + \frac{1}{\Gamma})^{-1}$. For the accuracy, which is defined as $N = (\mu/\sigma)^2$, we only need to calculate the variance σ^2 . The variance, however, is invariant under translations of the tick probability density and, thus, one can calculate the variance without loss of generality for $t_0 = 0$,

$$\sigma^2 = \int_0^\infty dt t^2 e^{-\Gamma t} - \mu^2 \quad (\text{C11})$$

$$= \frac{2}{\Gamma^2} - \frac{1}{\Gamma^2} = \frac{1}{\Gamma^2}. \quad (\text{C12})$$

Expressing the accuracy in terms of t_0 , we find $N = (1 + \Gamma t_0)^2$, that is, the accuracy increases with higher offset t_0 . There is a tradeoff, though: The greater t_0 , the lower the resolution. Eliminating the t_0 -dependency in the equations, we can establish an accuracy-resolution relation for the family of Heaviside populations given by $N = \Gamma^2/\nu^2$. \square

Analytic preliminaries. The tick probability density $p_{\text{tick}}(t)$ reveals how likely it is for a tick to happen during a given time interval. By the fundamental theorem of calculus, the cumulative non-tick probability $P[t \leq T]$ (whose negative derivative is $p_{\text{tick}}(t)$) contains the same information, and for proving the accuracy-resolution trade-off, the latter function turns out to be useful. Let us collect some general identities for the non-tick probability $P[t \leq T]$ which for some function $p(t)$ (be reminded, that the superscript *no tick* is suppressed) is given by

$$P[t \leq T] = e^{-\Gamma \int_0^t d\tau p(\tau)}. \quad (\text{C13})$$

There is a general constraint (Lemma 3) on how fast the non-tick probability decays. The fact that $p(t) \in [0, 1]$ ensures that $P[t \leq T]$ is monotonically decreasing (because the integral $\int_0^t d\tau p(\tau)$ is monotonically increasing) but does this not faster than exponentially.

Lemma 3. *For all $t > 0$ and $s > 0$, the cumulative non-tick probability satisfies the following inequalities:*

$$P[t \leq T]e^{-\Gamma s} \leq P[t + s \leq T] \leq P[t \leq T]. \quad (\text{C14})$$

Proof. This is a consequence of eq. (C13) and the fact that $0 \leq p(t) \leq 1$. \square

The cumulative non-tick probability can be used to calculate the moments of the probability density $p_{\text{tick}}(t)$ (see Definition 4 and Lemma 5) via the relation $p_{\text{tick}}(t) = -\dot{P}[t \leq T]$ and partial integration. It is important to note, that this discussion only makes sense for clocks that tick with certainty, i.e., $\lim_{t \rightarrow \infty} P[t \leq T] = 0$. This need not be true for all clocks, however, whenever we talk about a *tick probability density*, we implicitly assume that we have a properly normalized probability density in the probability theoretic sense [22].

Definition 4 (Moments). *For a given tick probability density $p_{\text{tick}}(t)$, define its k -th moment as*

$$t_k := \int_0^\infty dt t^k p_{\text{tick}}(t). \quad (\text{C15})$$

Lemma 5 (Tick probability moments). *The k -th moment of $p_{\text{tick}}(t)$ is related to $P[t \leq T]$ by the integral*

$$t_k = k \int_0^\infty dt t^{k-1} P[t \leq T]. \quad (\text{C16})$$

Proof. This is partial integration and for the boundary conditions, we use the assumption that $\lim_{t \rightarrow \infty} P[t \leq T] = 0$. \square

In particular, we can apply this result to $\mu = t_1$ and $\sigma = t_2 - \mu^2$, two expressions that are essential in Section C2, where we prove the upper bound. The average tick time $\mu = t_1$ can be written as the area below the graph of $P[t \leq T]$,

$$\mu = t_1 = \int_0^\infty dt tp_{\text{tick}}(t) = \int_0^\infty dt P[t \leq T]. \quad (\text{C17})$$

The second moment t_2 , on the other hand, is related to the center of mass of the graph for $P[t \leq T]$, up to normalization,

$$t_2 = \int_0^\infty dt t^2 p_{\text{tick}}(t) = 2 \int_0^\infty dt t P[t \leq T]. \quad (\text{C18})$$

2. Proof construction

The introduction of the Heaviside population in Definition 1 together with the result in Lemma 2, that these populations achieved the (claimed) optimal accuracy resolution relation, leads us to the following approach in proving the inequality $N \leq \Gamma^2/\nu^2$: we try to show that all clocks are worse than the one with a Heaviside population, which has essentially a perfect background clock. In that sense, what we are showing is that no clock is better than the one that is already perfect. The only premise we have for the proof is that the clocks we consider eventually tick (i.e., $p_{\text{tick}}(t)$ is a valid probability density) and that their evolution is continuous,¹ we call

¹

this *well-behaved*. That being said, our strategy to prove the upper bound consists of the following three steps:

- (i) We show that for any well-behaved top-level population $p(t)$, there exists a Heaviside- Θ type top-level population with the same average tick time.
- (ii) Then, we argue that the variance of the tick probability density coming from the Heaviside- Θ top-level population lower-bounds the variance coming from the generic population $p(t)$.
- (iii) We conclude that any well-behaved top-level population $p(t)$ must have an accuracy upper bounded by that of the Heaviside population and by using Lemma 2 on the properties of the Heaviside population, we have $N \leq \Gamma^2/\nu^2$. Once we are there, we have proven the accuracy-resolution trade-off theorem.

Step (i). Begin with a generic, but well behaved top-level population $p(t)$. Let μ be its average tick time and define $t_0 := \mu - 1/\Gamma$. The top-level population

$$p^\Theta(t) = \Theta(t - t_0) \quad (\text{C19})$$

has the same first moment $t_1 = \mu$ as the generic top-level population.² For the results from Lemma 2 to carry over, we need $t_0 > 0$. This is generally true as Lemma 6 guarantees. In fact, it tells us (see Fig. 5) that the average tick time of any well-behaved top-level probability can not be smaller than $1/\Gamma$ and that the best resolution is achieved by the top-level population which is constantly one.

Lemma 6 (Resolution upper bound). *For any well-behaved top-level population $p : \mathbb{R}_{\geq 0} \rightarrow [0, 1]$, the induced resolution ν cannot be greater than Γ .*

Proof. Equivalently to the statement in the Lemma, we can prove that $\mu \geq 1/\Gamma$. For this matter, use eq. (C17), to estimate the average tick time

$$\mu = \int_0^\infty dt P[t \leq T] \geq \int_0^\infty dt e^{-\Gamma t} = 1/\Gamma. \quad (\text{C20})$$

This concludes the proof. \square

¹ This ensures we are doing a fair comparison, otherwise, we'd have to discuss how to compare clocks that possibly never tick to ones that always tick. We reserve that discussion for future work.

² The average tick time is given by $\mu^\Theta = t_0 + 1/\Gamma$, with t_0 coming from the time translation of the Heaviside- Θ function and $1/\Gamma$ coming from the exponential decay.




FIG. 5. The sketch shows the time-dependency of a generic top-level population $p(t)$ together with the associated cumulative non-tick probability $P[t \leq T]$ compared to the non-tick probability $e^{-\Gamma t}$ of the top-level population which is constantly 1. A top-level population smaller than unity leads to a slower decay of the non-tick probability, which is responsible for a lower bound on the average tick-time given by $1/\Gamma$.

Step (ii). If $p(t) = p^\Theta(t)$, there is nothing to be shown, because $p^\Theta(t)$ achieves the minimal tick time variance $\sigma^2 = 1/\Gamma^2$ (see eq. (C12)). Hence, we assume from now on $p(t) \neq p^\Theta(t)$ to avoid this pathological case. Proposition 8 in Section C3 further discusses, that there exists a unique $t_* > t_0 = \mu - 1/\Gamma$, such that $P[t_* \leq T] = e^{-\Gamma(t_* - t_0)}$ (see Fig. 6). For all times $t < t_*$, the non-tick probability for the generic top-level population is smaller than that of the Heaviside population, i.e., $P[t \leq T] \leq P_\Theta[t \leq T]$. In the generic case, a tick before t_* is therefore more likely than for the Heaviside case. On the other hand, for all $t > t_*$, we have $P[t \leq T] \geq P_\Theta[t \leq T]$. That is, in the generic case, it is also more likely that after t_* the tick did not happen. Formally, the variance of the tick signal is bigger for a generic population than for the Heaviside one (Proposition 7).

Proposition 7. *For any well-behaved top-level population $p : \mathbb{R}_{\geq 0} \rightarrow [0, 1]$, the variance of the tick probability density $p_{\text{tick}}(t)$ is lower-bounded by that of $p_{\text{tick}}^\Theta(t)$, the tick probability density coming from the Heaviside top-level population $p^\Theta(t)$ defined in eq. (C19).*

Proof. By construction both $p_{\text{tick}}(t)$ and $p_{\text{tick}}^\Theta(t)$ have the same average tick time μ . Because the variance of the tick time is given by the difference $\sigma^2 = t_2 - \mu^2$, the following two statements are equivalent,

$$\sigma^2 \geq \sigma_\Theta^2 \Leftrightarrow t_2 \geq t_2^\Theta. \quad (\text{C21})$$

The second moments on the right-hand side can be calculated as the center of mass of the respective non-tick probabilities $P[t \leq T]$ and $P_\Theta[t \leq T]$ (up to constant prefactors, see also eq. (C18)). Therefore, the question we have to answer is, whether the graph of $P[t \leq T]$ has a center of mass at larger values of t then that of $P_\Theta[t \leq T]$? The two areas A_1 and A_2 in Fig. 6 are equal and they correspond to the difference in the area below the graphs of $P[t \leq T]$ and $P_\Theta[t \leq T]$. The graph of $P[t \leq T]$ has less ‘mass’ (i.e., area) on the interval $0 \geq t \geq t_*$ than




FIG. 6. The dashed lines show the top-level populations for the generic (blue) and the Heaviside case (green). The induced cumulative non-tick probabilities $P[t \leq T]$ and $P_\Theta[t \leq T]$ are drawn in the same color but as solid lines. The time t_* marks the unique crossing point of the two non-tick probabilities. On its right and left, the two areas A_1 and A_2 are defined by the difference between the two non-tick probabilities.

$P_\Theta[t \leq T]$. For $t_* \geq t \geq \infty$, however, $P[t \leq T]$ has more ‘mass’, resulting in a center of mass further to the right for $P[t \leq T]$ than for $P_\Theta[t \leq T]$. Formally,

$$\frac{t_2}{2} = \int_0^\infty dt t P[t \leq T] \quad (\text{C22})$$

$$= \int_0^\infty dt t P_\Theta[t \leq T] + \int_0^\infty dt t \underbrace{(P[t \leq T] - P_\Theta[t \leq T])}_{:=\Delta P(t)} \quad (\text{C23})$$

$$\geq \int_0^\infty dt t P_\Theta[t \leq T] = \frac{t_2^\Theta}{2}, \quad (\text{C24})$$

where the inequality comes about because the integral over $\Delta P(t)$ is greater equal than zero,

$$\begin{aligned} \int_0^\infty dt t \Delta P(t) &= \int_0^{t_*} dt t \Delta P(t) + \int_{t_*}^\infty dt t \Delta P(t) \quad (\text{C25}) \\ &\geq \int_0^{t_*} dt t_* \Delta P(t) + \int_{t_*}^\infty dt t_* \Delta P(t) = 0. \end{aligned} \quad (\text{C26})$$

All in all, this shows that $t_2 \geq t_2^\Theta$, which, by the initial remark, proves the proposition. \square

Step (iii). Combining all the results from the previous two steps, we proof the Theorem.

Proof. The accuracy N is defined as $N = (\mu/\sigma)^2$, and the resolution ν as $\nu = 1/\mu$. We can therefore express

$$N = \left(\frac{\mu}{\sigma}\right)^2 = \frac{1}{\nu^2 \sigma^2}. \quad (\text{C27})$$

From Proposition 7 we know that $\sigma^2 \geq \sigma_\Theta^2$, but $\sigma_\Theta^2 = 1/\Gamma^2$, as we derived in Lemma 2 (see eq. (C12) in the proof). In conclusion,

$$N = \frac{1}{\nu^2 \sigma^2} \leq \frac{1}{\nu^2 \sigma_\Theta^2} = \frac{\Gamma^2}{\nu^2}, \quad (\text{C28})$$




FIG. 7. This sketch visualizes the first step of the the proof of Proposition 8. The solid green line indicates the non-tick probability for the Heaviside population. We see that for the red dashed line which does not fall below 1 before t_0 , the total integral is greater than that of $P_\Theta[t \leq T]$, contradicting the construction, which ensures that both functions have the same integral. The blue, dashed line shows a plausible function for $P[t \leq T]$.

which proves the claim, and thereby the optimality of the family of Heaviside- Θ top-level populations. \square

3. Details on the optimality of the Heaviside population

In Step (ii) of Section C 2, the claim is made that there exists a unique $t_* \geq t_0 = \mu - 1/\Gamma$, such that the two non-tick probabilities coincide $P[t_* \leq T] = e^{-\Gamma(t_* - t_0)}$. To be more precise, we formulate

Proposition 8. *The set $\mathcal{S} = \{t \geq 0 : P[t \leq T] = P_\Theta[t \leq T]\}$ is of the form*

$$\mathcal{S} = [0, a] \sqcup \{t_*\}, \quad (\text{C29})$$

where $t_* > t_0$, and $[0, a]$ is an interval with upper bound $a < t_0$.

Proof. We divide the proof of this proposition into three steps (a) - (c), with visual aids in Figures 7 and 8.

Step (a). We claim that *there exists a $0 < t < t_0$, such that $P[t \leq T] < 1$* . Suppose that this was not the case, then $P[t \leq T] = 1$ for all $t \leq t_0$ (see Fig. 7). By requirement, there must exist a t such that $p(t) \neq p^\Theta(t)$, and because by our (contradictory) assumption t cannot be smaller than t_0 , it must be bigger than t_0 . Since, $P[t \leq T] \geq e^{-\Gamma(t - t_0)}$, for all $t > t_0$, there must exist a $t > t_0$, such that $P[t \leq T] > e^{-\Gamma(t - t_0)}$. But this contradicts the assumption that $\int_0^\infty dt P[t \leq T] = t_0 + 1/\Gamma$. Therefore, it must be that there is a $t < t_0$ with $P[t \leq T] < 1$.

Step (b). Our next claim is, that *there exists an $0 < a < t_0$ such that on $[0, a]$, we have $P[t \leq T] = 1$ and for all other $a < t \leq t_0$, $P[t \leq T] < 1$* . By our previous claim (1), there exists $0 < t < t_0$ with $P[t \leq T] < 1$. Let a be the infimum of all such t .³ Then, $[0, a] \subset \mathcal{S}$ (continuity




FIG. 8. Here, we visualize step two of the proof of Proposition 8. Again, the solid, green line indicates the non-tick probability for the Heaviside population. To satisfy the construction that $P[t \leq T]$ and $P_{\Theta}[t \leq T]$ have the same integral on $\mathbb{R}_{\geq 0}$, the area below the graph of $P[t \leq T]$ for $t > t_0$ must be bigger than that of $P_{\Theta}[t \leq T]$. Thus, at some point after t_0 , $P[t \leq T]$ must be bigger than $P_{\Theta}[t \leq T]$ (see the blue area denoted by A_{good}), and not always smaller (see red area denoted by A_{bad}).

of P). Moreover, by the properties of an infimum, for all $\varepsilon > 0$, there exists a $\delta > 0$, such that $P[a + \delta \leq T] < 1$. But $P[t \leq T]$ is monotonically decreasing, hence, this is true for all $\delta > 0$, such that $a + \delta \leq t_0$, proving the claim.

Step (c). In this last step, we show *the existence of a unique $t_* > t_0$ such that $t_* \in \mathcal{S}$* . For a visualization of this step, see Fig. 8. The integrals of $P[t \leq T]$ and $P_{\Theta}[t \leq T]$ on $\mathbb{R}_{\geq 0}$ coincide. However, we have just argued that on a non-empty interval between 0 and t_0 , $P[t \leq T]$ is strictly smaller than $P_{\Theta}[t \leq T]$. This leads to the inequality $\int_0^{t_0} dt P[t \leq T] < \int_0^{t_0} dt P_{\Theta}[t \leq T]$. In order to ensure that the integrals over all of $\mathbb{R}_{\geq 0}$ are equal, there must exist a $t' > t_0$ such that $P[t' \leq T] > P_{\Theta}[t' \leq T]$. For all $t > t'$, this inequality must be true too, because $P[t \leq T]$ can not drop faster than $e^{-\Gamma(t-t_0)}$, due to boundedness of the top-level population $p(t) \leq 1$. Set t_* to be the infimum of all such t' . Note that t_* must be strictly greater than t_0 due to continuity of $P[t \leq T]$. This concludes the proof of the last step and therefore of the proposition. \square

³ The set $\{t : P[t \leq T] < 1\}$ is lower bounded by 0 and non-empty by the claim made in step (a), thus, the infimum exists.